

Taboo and Collaborative Knowledge Production: Evidence from Wikipedia

KAYLEA CHAMPION, University of Washington, United States

BENJAMIN MAKO HILL, University of Washington, United States

By definition, people are reticent or even unwilling to talk about taboo subjects. Because subjects like sexuality, health, and violence are taboo in most cultures, important information on each of these subjects can be difficult to obtain. Are peer produced knowledge bases like Wikipedia a promising approach for providing people with information on taboo subjects? With its reliance on volunteers who might also be averse to taboo, can the peer production model produce high-quality information on taboo subjects? In this paper, we seek to understand the role of taboo in knowledge bases produced by volunteers. We do so by developing a novel computational approach to identify taboo subjects and by using this method to identify a set of articles on taboo subjects in English Wikipedia. We find that articles on taboo subjects are more popular than non-taboo articles and that they are frequently vandalized. Despite frequent vandalism attacks, we also find that taboo articles are higher quality than non-taboo articles. We hypothesize that stigmatizing societal attitudes will lead contributors to taboo subjects to seek to be less identifiable. Although our results are consistent with this proposal in several ways, we surprisingly find that contributors make themselves more identifiable in others.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**; **Collaborative and social computing theory, concepts and paradigms**; *Empirical studies in HCI*; **Wikis**; **Collaborative content creation**; • **Security and privacy** → **Social aspects of security and privacy**.

Additional Key Words and Phrases: peer production; Wikipedia; online communities; taboo; privacy; anonymity;

ACM Reference Format:

Kaylea Champion and Benjamin Mako Hill. 2023. Taboo and Collaborative Knowledge Production: Evidence from Wikipedia. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 299 (October 2023), 25 pages. <https://doi.org/10.1145/3610090>

1 INTRODUCTION

Taboos are behavioral prohibitions characterized by the notion that violations of the prohibition create symbolic uncleanness or pollution [2, 23]. Although taboos vary enormously across cultures, they exist in virtually all societies [2, 23]. People try to avoid being polluted by taboo in a range of ways—frequently by not talking about or even mentioning taboo subjects. When taboo subjects are discussed, people refer to them indirectly or vaguely. As a result, access to high-quality information on taboo subjects is frequently difficult. This is problematic because taboo subjects often include important subjects such as mental health, reproduction, menstruation, abuse, extremism, and human rights violations.

The growth of the Internet and the development of free high-quality knowledge bases like Wikipedia have made a wealth of information widely available. Are Internet knowledge bases a

Authors' addresses: Kaylea Champion, University of Washington, Seattle, WA, United States, kaylea@uw.edu; Benjamin Mako Hill, University of Washington, Seattle, WA, United States, makohill@uw.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2023 Copyright held by the owner/author(s).
2573-0142/2023/10-ART299
<https://doi.org/10.1145/3610090>

promising approach to providing people with information on taboo subjects? In that they provide opportunities for privacy in information consumption, social computing systems have been cited as promising avenues for accessing information on a range of taboo subjects like menstruation [3, 18, 76, 86], sexual abuse, and harassment [6, 57]. On the other hand, many online knowledge bases rely on volunteers who choose their own tasks and subject areas, and it seems reasonable to assume that at least some volunteers may be reticent to contribute to resources about taboo subjects. In either case, there is reason to believe that taboo shapes knowledge production in social computing systems. That said, we know of no work that explores the role that taboo plays in online knowledge production.

One challenge to studying taboo as a more general phenomenon is to develop a way to identify taboos systematically. Examining taboo from a systematic perspective allows us to develop evidence about taboo as a broad social phenomenon. To do so, we draw inspiration from work in linguistics to develop a novel computational approach [16]. Our method uses a supervised machine learning classifier trained on words in dictionary definitions associated with euphemisms—i.e., subjects that speakers go out of their way to describe indirectly, often at the cost of clarity. To understand how taboo might be shaping activity in social computing, we use these words from dictionary definitions to identify a set of articles on taboo subjects in English Wikipedia as well as a comparison set of otherwise similar articles. For example, since the term “passed away” is a euphemism that would include the word “dead” or “death” in its definition, we might identify the Wikipedia article on “death” as taboo.

Using detailed digital trace data on articles and contributors from Wikipedia, we test five hypotheses derived from theory about how taboo will change the way that Wikipedia articles are consumed and produced. We find support for hypotheses that articles on taboo subjects are more popular than non-taboo articles and that taboo articles are more frequently vandalized than non-taboo articles. Despite these frequent attacks, we also find that taboo articles are of higher quality than non-taboo articles, contrary to our hypothesis. We hypothesize that societal attitudes against contact or association with these subjects would lead contributors to taboo subjects to seek to be less identifiable, and although this is sustained in part by our results, it is also contradicted in part. These surprising results suggest that contributors to the public production of taboo knowledge navigate more complex privacy trade-offs than previously theorized.

The remaining sections of the paper are structured as follows. In §2, we examine theories of peer production, taboo, social norms, and identifiability/anonymity, including what these theories suggest about the peer production of information about taboo subjects. In §3, we describe our empirical setting Wikipedia. We then describe the novel technique we develop to identify taboo and the construction of our analytic sample in §4, and our analytical process in §5. Section §6 reports results from our hypothesis tests. We discuss the significance and implications of our findings in §7, some important limitations of this work in §8, and conclude in §9.

2 BACKGROUND

2.1 Commons-Based Peer Production

As Benkler [10] describes, one of the technological and social advances of the last few decades is the creation of a novel form of production: commons-based peer production. Individuals participating in commons-based peer production activities online are volunteers who self-organize around a goal, select their own tasks, and engage in the creation or maintenance of information goods. These volunteers are sometimes motivated by generosity and altruism, or simply pursuit of their own passions—which may be in contradiction to dominant social and market logics. Often these information goods are also public goods, as in the case of Wikimedia projects such as Wikipedia

and Wiktionary, the mapping project Open Street Map, and free/libre open source software such as GNU/Linux and the Python programming language. These valuable projects are critical parts of our digital infrastructure: delivering knowledge that answers our searches online, the software that organizes and displays those answers, and the servers and protocols that convey them from distant data centers to our fingertips. Although these projects are rightfully hailed as transformative, their results—and the peer production process that produces them—are far from perfect.

Studies of peer production projects have pointed out numerous examples of bias and neglect. The participant pool of many prominent examples of peer production is generally male, including Wikipedia [21, 25, 38], Linux [59, 66], and Open Street Map [80]. The products developed through peer production may also be biased or neglect important materials. In Wikipedia, previous work has found neglect of articles about women [85], non-English languages [36, 46, 56], and countries, religion, and LGBTQ subjects [88]. In Open Street Map, research has found a lack of map information about the global South [80]. These biases in peer-produced resources are troubling because they reflect and may serve to perpetuate existing societal biases, inequalities, and hegemonic structures. Despite these flaws, however, peer production has such advantages as self-organization and self-selection of tasks.

Examining the question of taboo knowledge in peer production offers us insight into the extent to which this novel organizational form is simply serving to reproduce and magnify existing features of society, or whether participants may perhaps be finding ways to resist cultural norms.

2.2 Taboo—and Taboo in HCI

Taboo has been the subject of an enormous amount of scholarship in anthropology, sociology, and linguistics [2, 23]. Taboos demarcate the forbidden and unspeakable parts of existence from those that are recognized as sanctified, worthy, or simply acceptable. A taboo acts as a behavioral prohibition and is characterized by the notion that violations of the taboo create a sense of uncleanness or pollution. The uncleanness can be literal (e.g., contact with something covered in germs may transfer those germs) or may be symbolic (e.g., make one's prayers unacceptable in the eyes of God). Because making contact with something taboo makes a person unclean, taboo spreads its uncleanness through interaction. Although the taboo itself may make someone feel personally uncomfortable or embarrassed even if they are alone, taboo is often enacted socially. Although speakers may use euphemism to protect the sensibilities of others when discussing taboo subjects, McGlone and Batchelor [54] found through an experimental manipulation that speakers are more likely to use euphemism when discussing taboo subjects if they expect to be identified to the listener, concluding that the reticence to speak about taboo subjects is more a matter of a speaker trying to save face than protecting others. Taboo may manifest in complex ways in social computing settings. We may not only have a sense of being visible to friends, but also imagine a broader audience, which might be profoundly public [50, 52]. We are further visible to our technology—not only in the form of history and device logs, but also in the ways our behaviors train the algorithms that curate our experience—any of which may serve to embarrass us, restrain us, and drive us to seek privacy [82].

Mary Douglas's work in *Purity and Danger: An Analysis of the Concepts of Pollution and Taboo* points out that when “pollution rules” that constitute taboo are examined with care, social order is constantly implicated in what we consider clean or unclean, acceptable or forbidden. Ultimately, she says, we should understand taboo as part of a symbolic system in which “uncleanness is matter out of place” [23, p. 41] and a violation of “cherished classifications” (*ibid*, p. 37). Hence, to the extent that people interact with subjects that are taboo, they are in some way challenging social order and by extension society itself. As a result, reference materials associated with taboo subjects may be restricted, censored, or banned outright [2].

The English word “taboo” originated in the voyages of Captain Cook who derived the word from the Tongan word “tapu” meaning sacred or forbidden. The term came to be used in colonial and anthropological accounts of non-Western cultures, often as part of characterizing those cultures as “uncivilized” [2]. Despite this history, taboo has come to be understood as existing in virtually all cultures and societies throughout history [2, 22, 23]. Although every member of every society has some internal and personal sense of taboo, taboos vary by culture, religion, relative levels of privilege, and more. Like early colonial anthropologists, we may fail to recognize our own taboos because they seem natural or even objective to us [23]. Seeking out a more systematic way to identify taboo may therefore serve to de-center and de-naturalize our own cultural position.

Despite the rich literature in other disciplines, we believe that ours is the first article in social computing to approach the study of taboo in general, instead of focusing on specific taboo subjects. To what extent is “taboo” a subject of concern for research at the intersection of society, technology, and design? Although a search of the over 2.9 million articles in the ACM digital library returns 1,849 results for ‘taboo’, most of these articles refer to an algorithmic search strategy in which potential solutions are marked as forbidden or to a party game in which players try to guess a word when prompted with clues that cannot include the word itself. A closer examination of the 174 research articles published in SIGCHI venues revealed that most of the remaining articles mentioned the term “taboo” only once and in passing.

Although none of these articles sought to understand or measure the effect of taboo in general, our search did reveal a number of articles on systems and designs that are intended to counter or transcend the harmful effects of taboo on information environments in specific cases such as menstruation [3, 18, 76], sexuality [45, 63], toilet training [37], and AIDS education [73]. Another line of research seeks to understand how people use existing social computing systems in ways that are shaped by specific taboos—e.g., sharing information about menstrual health [86] and menopause [49], or describing experiences such as sexual abuse [6], sexual harassment [57] and pregnancy loss [5].

One related body of work is the study of stigma, which is a connected but distinct concept. A taboo is a behavioral prohibition, whereas stigma refers to an identity state, perhaps but not always associated with violating taboos. One may be stigmatized for a congenital condition or by violating a range of different types of social norms. In *Stigma: Notes on the Management of Spoiled Identity*, Erwin Goffman describes stigma as “an undesired differentness from what is expected” [28, p. 5] and, like Douglas, invokes the role of social order. For Goffman, a stigmatized identity means that one’s social identity as perceived externally is in some way different from what is perceived internally. For example, a formerly incarcerated person may be stigmatized—they see themselves as ‘like everyone else’ and trying to live a normal life, but to their neighbors they may forever be discredited as criminal. Goffman observes that stigma in society is transferred via association: the child of a person known to be a criminal may also be the subject of suspicion.

Although none of this work takes up the concept of taboo in general, these papers point to the fact that taboo is an important feature of online information seeking and production in at least several specific cases. Taking a systematic approach is complementary to these studies that focus on *a priori* taboos by supporting comparative analysis across contexts and across a range of taboos, as well as potentially expanding the set of taboos explored through various methods.

2.3 Taboo Content Consumption

Avoidance of taboo is rarely universal. We may simply need information related to taboo subjects, such as those related to such everyday activities as urination, defecation, menstruation, and sexual intercourse. Some people deliberately violate taboos and seek out community in doing so. Others are struggling with taboo issues or have survived taboo experiences and are seeking supportive

environments. Further, the forbidden nature of taboos may lead some people to feel drawn toward the subject—they may titillate or fascinate or offer an avenue to rebel against society’s norms.

Privacy may also play a role in taboo violation. For example, Malinowski [51] famously observed in his 1926 account of the people of Trobriand Island that a taboo may be widely violated in private and the subject of gossip and humor, yet still have severe consequences when its violation is turned into a public spectacle. Likewise, support groups often have rules about maintaining participant anonymity, and some sites have norms and affordances that support anonymized interactions about taboo subjects [4].

Web browsing may feel relatively private, especially when we are only consuming content. We may feel less inhibited in an online environment and more willing to engage with material that we would be reluctant to be seen reading about or heard speaking about publicly [74]. One might expect that feelings of safety to be higher in websites—like Wikipedia—that are accessible without paywalls or registration requirements which signal to a reader that consumption is being tracked. Given that information seekers may have relatively few alternative sources for high-quality information about taboo subjects and the relative privacy of reading Wikipedia, we suggest **H1: peer produced resources about taboo subjects will receive higher readership than peer produced resources about other comparable subjects.**

2.4 Taboo, Public Knowledge, and Identifiability

Although consuming content hosted in knowledge bases like Wikipedia may feel private, producing knowledge resources is often a publicly visible act. In Wikipedia, every article has a history tab containing each past version of itself. A “diff” view allows anyone to examine the letter-by-letter impact of every change made. Beyond the contribution itself, the account or IP address of every contributor is tracked and made visible. Both the article’s revision history and an individual contributor’s contribution history can be reviewed by anyone. Contributors can offer additional information about themselves by customizing their user page, by registering an email address that can be used to contact them, and by disclosing their gender. Although the effort required to identify contributors varies, Wikipedia’s high level of transparency means that every contributor is identifiable in some dimension. Users of tools to protect personal privacy, including anonymity-preserving proxies such as the popular Tor browser, are banned from contributing to Wikipedia [83].

On Wikipedia, people edit only what they want to edit [11]. Because taboos are so powerful and pervasive, they may become internalized. As a result, contributors may seek to avoid taboo subjects even in private settings. A high level of public visibility and the voluntary nature of peer produced projects suggest that people may be reluctant to choose to contribute to taboo subjects out of a reluctance to contaminate themselves. Therefore, we propose **H2: peer produced resources about taboo subjects will receive fewer contributions than other comparable peer produced subjects.**

Online knowledge bases are typically secondary or tertiary sources of information that depend on the availability of other sources to establish credibility [24]. Taboo subjects are often the targets of censorship efforts [2]. For example, materials about taboo subjects like violence, sexuality, race, and religion are regular targets of book banning campaigns in K-12 schools in the United States [9]. Blocking of political, religious, and social norm-violating content (including pornography, drugs, LGBTQ content, and online dating) is widespread [90]. At the state level, for example, China suppresses information on human rights abuses and collective action, and Norway blocks many pornography and gambling websites [47, 75, 90]. Although social pressure and a lack of available secondary sources might lead to lower-quality contributions, the nature of taboo subjects themselves may tend to draw in or inspire bad faith contributions—e.g., writing sexual slurs on

pages about sexual subjects. Given societal and governmental opposition to materials on taboo subjects, the suppression of primary and secondary sources that might otherwise be used to build knowledge bases, and the role of troll, griefers, and vandals, we propose **H3: peer produced resources about taboo subjects will receive lower-quality contributions than peer produced resources about other comparable subjects.**

Some theorists of open content production have made direct connections between the quality of a resource and the volume of contributors. Raymond [65] famously described this process as “Linus’s Law:” “given enough eyeballs, all bugs are shallow.” Empirically, Haklay et al. [31] found that positional accuracy in open mapping projects increases with the number of participants. Greenstein and Zhu [30] found that neutrality in point of view in Wikipedia articles improved with increases in the number of contributors. Synthesizing these observations—and as a consequence of H2 (fewer overall contributions) and H3 (lower-quality contributions)—we propose **H4: peer produced resources about taboo subjects will be lower quality than peer produced resources about other comparable subjects.**

Finally, we consider the impact of taboo on the willingness of contributors to identify themselves. Possessing knowledge about a taboo subject may lead to stigma [28]. Social embarrassment and opprobrium can result from being known to have violated taboos leading to loss of reputation and position, or in extreme cases, to incarceration and violence. For example, Wikipedia editors have been harassed for their contributions both on and off Wikipedia [14, 42, 60]. In some cases, editors have even been prosecuted for their edits, as in the case of the Belarusian government arresting a Wikipedia editor for writing that Russia has invaded Ukraine [72]. Hence, people contributing knowledge about a taboo subject to online knowledge bases may engage in a range of strategies to control information about their identity and behavior [70]. For example, the Russian language edition of Wikipedia has suppressed the public edit logs for pages associated with the Russian invasion of Ukraine and encouraged editors to switch to alternate, less identifiable accounts as a means to protect their safety [1, 13].

If users participate at all, they may only engage with taboo when they feel they have control over information about themselves: by interacting anonymously or developing a new or special-purpose account or profile [12, 26, 44, 55, 77]. Because all these factors suggest the desirability of partially controlling or fully avoiding personal association with taboos, we propose **H5: peer produced resources about taboo subjects will be more likely to receive contributions from less identifiable personas than those about other comparable subjects.**

3 EMPIRICAL SETTING

Wikipedia is a vital source of information worldwide, with 318 language editions as of 2023.¹ In December 2022 alone, Wikipedia received more than 24 million contributions and served up more than 23 billion page views across all projects, with 255,255 new user accounts created.² The largest language edition, English, contains more than 7.5 million articles as of January 2023 and averages 559 new articles created each day.³ Wikipedia entries are often the first result for a query on Google; text from relevant Wikipedia pages may be surfaced in search engine sidebars without the user needing to dig any further [87].

¹https://en.wikipedia.org/wiki/List_of_Wikipedias

²<https://stats.wikimedia.org>

³<https://en.wikipedia.org/wiki/Wikipedia:Statistics>

4 METHODS

4.1 Technique for Identifying Taboo Subjects

Our research questions concern taboo subjects in general. Therefore, we seek to minimize the assumptions we make about what is and is not taboo based on our own cultural context. Work in linguistics concerning taboo subjects has described how people use figurative language and euphemism to deal with taboo [22]. For example, we may say “passed on” instead of “died.” Burridge [16] describes euphemisms as “a verbal escape hatch in response to taboos.”

To build our dataset of taboo articles in English Wikipedia, we made use of English Wiktionary.⁴ Wiktionary is a sister project to Wikipedia, building a peer produced dictionary. Wiktionary has received comparatively less attention from researchers when compared to the encyclopedia-building Wikipedia project.⁵ Because Wiktionary entries are somewhat free-form, we used the *wiktextextract* Wiktionary parser and associated dataset to separate entries into key components for analysis: word, definition, and tags.⁶ In Wiktionary, definitions and tags are associated with each “sense” of a word to manage polysemy. Our corpus is composed of definitions remaining after we filtered the dataset to remove non-definitions (e.g., redirects to other words, definitions stating only that a word is a synonym or initialism of another word) and non-English definitions when the entry was supplying a translation. This filtration of 1,099,350 dictionary entries from the September 1, 2021 version of Wiktionary left us with 404,304 unique entries, each composed of a word, a definition, and, in some cases, associated tags.

Many dictionaries tag particular word usages as euphemistic. This makes dictionaries a source of general data on taboo because *definitions* marked as euphemism contain a description of subjects considered taboo. For example, the term “member” has multiple definitions, one of which is marked as “euphemistic” in English Wiktionary (see Figure 1). The term being defined (i.e., “member”) is not taboo—that is precisely why it is effective at creating distance from the taboo concept in the definition. We were inspired in this approach by Buscaldi and Hernandez-Farias [17] who used the fact that Italian Wiktionary contributors have explicitly tagged some words as “taboo” to train a classifier for sentiment analysis.

Our approach involved collecting only the text of definitions, not the words being defined. Tags in English Wiktionary can be applied to any definition. In Figure 1, “(euphemistic)” and “(logic)” are tags. We treated all definitions tagged as “euphemistic” as indicative of taboo. For example, in Figure 1, the definitions of “member” that refer to group association, a part of a whole, animal parts, and logic are each marked as not indicative of taboo in our dataset. However, the definition that refers to genitalia is tagged “euphemistic” and is marked as indicative of taboo in our dataset. We removed numbers and stop words, using the English stop words list in the Python Natural Language Tool Kit (NLTK), from all definitions.⁷ We also removed the words ‘term’, ‘used’, ‘usually’, ‘particularly’, ‘etc’, ‘extremely’, ‘especially’, ‘one’, ‘en’, ‘something’, ‘often’, ‘synonym’, ‘like’, and ‘person’ because these words recur very frequently in Wiktionary entries but convey only intensity.

We used term frequency-inverse document frequency (TF-IDF) [64] and ridge regression from the scikit-learn python package [61] to assess which words, bigrams, and trigrams (“n-grams”) were more commonly present in the definitions tagged as euphemistic. We sorted these terms by their coefficient from ridge regression to develop a list of 500 n-grams indicative of taboo.

⁴https://en.wiktionary.org/wiki/Wiktionary:Main_Page

⁵For example, Wiktionary appears four times among the 1,972 examples of Wikipedia-related research listed in the Wikipedia page about Wikipedia research: https://en.wikipedia.org/wiki/Wikipedia:List_of_academic_studies_about_Wikipedia

⁶See <https://github.com/tatuylonen/wiktextextract> and <https://kaikki.org/dictionary/>

⁷<https://www.nltk.org/>



Fig. 1. The Wiktionary definition of “member”—which has meanings that range from group association to anatomical.

Table 1. A random selection of articles from our taboo set. We have omitted racial/ethnic slurs, explicit sexual acts, and profanity in this table but included them in our dataset.

Sex differences in human physiology	Genitourinary system	Love God	Cat meat
Gastrointestinal tract	Damnation	Alive and Dead	Old Age
Monosodium glutamate	Being Bobby Brown	Abdominal Obesity	Feral cat
Amazon Women on the Moon	Menstruation	Love and God	Layoff
Mental disorder	Unemployment	Head on Collision	Vulva

To test our hypotheses, it is not necessary that we detect all taboo articles in Wikipedia. Instead, we need only generate a dataset composed of two samples, one of which has a higher proportion of taboo articles. In machine learning terms, we sought precision at the cost of recall. As a result, we performed an exact matching between Wikipedia article titles (minus stop words) and the 500 n-grams we identified as most associated with taboo. This conservative approach left us with a relatively small but high-confidence sample. Examples of these articles are listed in Table 1.⁸

For the comparison set, we considered using a sample of random Wikipedia articles. However, in exploratory analyses, we realized that this sample is inappropriate because most article titles in Wikipedia are not n-grams that appear in dictionary entries and could never have been selected by our taboo article identification process. Although many Wikipedia articles are phrases, events, locations, and proper names, the articles in our taboo set have a “dictionary definition-esque” quality. To ensure a comparable sample, we limit our comparison set to a random selection from the population of 115,681 articles in English Wikipedia with titles (minus stop words) that match an n-gram found in English Wiktionary definitions.

We conducted additional processing on both the comparison set and the taboo set. Many Wikipedia pages serve only to disambiguate a term or provide a list of other articles. These non-articles were omitted. Further, some Wikipedia articles are “redirects”—they serve as largely behind-the-scenes connectors to other articles and require special handling (see Hill and Shaw [39].) When a term in our n-gram list was the title of an article serving as a redirect, we followed the redirect and extracted data for the target article. When the target of a redirect was a subsection of a longer page, we dropped the article from our sample because our measures cannot be easily

⁸Inclusion of the article on the reality show “Being Bobby Brown” in the taboo set initially surprised us. The article is included because the article “Hell to the no” (a phrase the article says was popularized by the show) is a redirect that lands on this article. The words “to”, “the” and “no” are stop words, leaving us with a match on the term “Hell.”




Fig. 2. Our analytical pipeline first extracts n-grams, labeling them taboo if they are drawn from definitions tagged as euphemistic. Our samples are drawn from those articles that match these n-grams.

constructed for article subsections. This left us with 74 taboo articles and 3,255 randomly selected comparison articles.

4.1.1 Validation. Evaluating our approach is difficult given the culturally varying nature of taboo and the lack of ground truth data. As a simple validity check, we test a hypothesis that articles in our taboo set will be more likely to have been assigned Wikipedia categories related to sex—a subject that appears frequently in taboos across many cultural contexts [6, 22, 23, 45, 51, 57]—than those in our comparison set.

Categories in Wikipedia are used for a wide range of purposes and are largely free form [81]. Past research by Chen et al. [19] and by Asthana and Halfaker [8] has shown the utility of categories associated with topically-focused WikiProjects [58] are a source of reliable information about an article’s subject. We used Wikimedia Foundation public APIs to obtain a list of 5,575 unique categories from the 3,329 articles in our datasets as well as their associated talk pages.

We then used a logistic regression model to assess the relationship between sexual topics and membership in the taboo or comparison set in our sample. Articles can, and often do, exist in multiple categories simultaneously. As a baseline, taboo articles comprise 2.2% of our sample. Given an article in our dataset that has been assigned to the scope of WikiProject Sexology and Sexuality, our model suggests that there is a 44% chance of it being in the much smaller taboo set ($\beta = 3.4$, $z = 8.43$, $p < 0.001$). This reflects very strong evidence that our taboo set is more likely to contain articles about sex and sexuality than is our comparison set.

4.2 Data and Measures

Having identified our taboo and comparison samples, we obtained the full text and metadata for revisions made to all articles in the samples by parsing the XML database dumps released by the Wikimedia Foundation.⁹ Given that our interest is in human behavior, we attempt to omit revisions made by bots. To identify edits made by bots, we used a list we scraped from a Wikipedia page listing current registered bots as well as a historical dataset produced and released as part of Geiger and Halfaker [27]. Our final article and revision datasets contained 177,974 revisions made to taboo articles and 2,052,209 revisions made to articles in the comparison set. We used these data to construct longitudinal measures for each article in our sample: measures of quality of each article each month, the contribution history for each article, each contribution as well as its contributor, and the per-month viewership of the article. We describe each in turn.

In **H1** we suggest that taboo subjects will receive higher readership than comparable subjects. We operationalize readership using aggregate article view count data published by the Wikimedia Foundation to obtain the total number of views for each article in our sample for each month. We

⁹<https://dumps.wikimedia.org>

obtained 541,913 article-month measures of viewership and then calculated the mean view rank for each article.

To test **H2** on the number of contributions, we count the number of contributions made to each article in our sample over the life of each article.

In **H3**, we take up the question of whether taboo subjects receive lower-quality contributions. We operationalize contribution quality using two measures. The first, *was reverted*, indicates that a given contribution was rejected after it was published. A contribution is said to have been reverted if someone restored an article to its state before the contribution in question. We looked 10 contributions forward to see if a contribution was reverted. The second measure we use, available through the ORES machine learning classifier's "revision" model, offers a prediction of whether or not a given contribution was "damaging" [89]. In Wikipedia, damaging contributions are defined as those contributions that may need to be removed because they do not conform to Wikipedia guidelines. Damaging contributions can range from vandalism and misinformation to naive statements or formatting errors. Although many damaging contributions are simply reverted, damaging contributions can instead be revised to be acceptable or be treated by community members as an opportunity to teach a new contributor [33].

Our hypothesis **H4** concerns the quality of an article. We measure quality using the assessments generated using ORES' quality classifier [32]. The classifier was trained using the work of Wikipedia contributors and classifies revisions based on several structural features of articles, including length, the presence of pictures, and the use of links to other sources.

In **H5**, we hypothesized that contributors to taboo subjects will be less identifiable. Because identifiability has many dimensions including the presence of unique personal identifiers, location, the presence or absence of a consistent pseudonym, and behavioral patterns [53], we operationalize identifiability in several ways derived from prior work about identifiability on Wikipedia: (H5A) editing without an account [7, 26], (H5B) editing using a new account or one with fewer contributions [26, 48, 67], (H5C) revealing relatively little information on one's profile [15, 67], (H5D) identifying one's gender [15], and (H5E) setting one's account as "emailable," which also requires a confirmed email address.

We measure whether a contribution was made with or without an account using metadata associated with each revision (H5A). We analyze contributor experience levels by counting the number of revisions each contributor has made to identify each revision as the contributor's "*n*th edit" (H5B). We use the XML database dumps to determine whether contributors had a user profile page at the time they were revising the articles in our samples (H5C), setting a dichotomous variable if they had a user page at any of these time points. Contributor gender and being emailable by others can be set by users in their Wikipedia setting page. We obtained this information via queries to the Wikimedia public API (H5D, H5E).

Whether a contribution is made by a user with or without an account is highly visible to other Wikipedia users. Previous editing by a user is less visible but is still readily accessible. These signals have been previously found to influence how work is received in Wikipedia, with non-accountholders and accounts with lower editing counts being moderated more strictly [79]. These two identifying signals are also recorded automatically by the Wikipedia website software. The other signals are "opt-in." User pages are free text profile pages and exist only if a user chooses to make one. Gender and emailability are only available via an API call, although emailability can be inferred by attempting to email the person via the website interface that is available to all logged-in users.

There is a potential confounder in our tests for H5A and H5B because some articles are "protected" so that they cannot be edited by contributors without accounts or by newcomers (those with

accounts less than 4 days old and fewer than 10 edits).¹⁰ In this way, page protection directly determines whether some anonymity seekers can participate.

Following the method described in Hill and Shaw [40], we identify “protection spells”—periods wherein a given page was protected. As Hill and Shaw [40] describe, page protection data are left censored because of missing logs before 2008. As a result, we limit our assessment to the point at which reliable logs are available. We use these data to calculate the proportion of time each article in our dataset was protected since 2008. Articles varied in *protected proportion* from 0 (e.g., the articles *Abdominal Obesity* in our taboo set and *Abbot* in our comparison set that were unprotected for the entire period) to more than 0.99 (e.g. the articles *Hell* in our taboo set and *Messianic Judaism* in our comparison set that were protected for almost the entire period). We use protected proportion as a control variable in our analyses for H5A and H5B.

4.3 Ethics

This study was conducted entirely using publicly available data published by the Wikimedia Foundation and does not involve any interaction or intervention with human subjects. This type of research using these data has been reviewed by the IRB at our institution and has been determined to not be human subject research. However, this work removes public digital trace data from its original context. Additionally, computational approaches have the potential to reveal behavioral trends in ways that individuals may find uncomfortable, especially given the subject of this study. As a result, we have redacted account names and IP addresses of the individuals who contributed to the articles in our sample. Article view data were fully anonymized by the Wikimedia Foundation prior to release.

Finally, we recognize that the use of taboo language may have differing impacts on those reading our work in various contexts. To make our work as available as possible while minimizing harm, we have redacted or omitted the use of racial/ethnic slurs and profanity in Table 1 as well as in the text of our paper. However, all terms were included in the underlying analysis with no omissions. In the interest of accountability and open science, our full results and data are available in our dataset release and online supplement in an unabridged form.

5 ANALYTIC PLAN

With respect to **H1** on viewership, our dataset includes articles written at different times and Wikipedia itself has received varying levels of traffic in the last two decades. Therefore, we calculated the rank in terms of views of each article with respect to other articles in our samples. We compared the mean within-month view rank of each article in our taboo and comparison samples across all months and tested for statistical significance using a Mann-Whitney U-test. The unit of analysis is the article, with mean rank across months as the outcome variable and taboo as our key predictor.

To test **H2** about contribution quantity, we examine the median number of contributions in our two samples and tested statistical significance using a Mann-Whitney U-test. The unit of analysis is the article, with total contributions as the outcome variable and taboo as our key predictor.

To test **H3** on contribution quality, we compared the median rate of reverted contributions of the two samples (total reverted/total contributions) and the median rate of damaging contributions, and tested statistical significance using a Mann-Whitney U-test. The unit of analysis is the article, with total reverted (damaging) contributions as the outcome variable and taboo as our key predictor. As a robustness check, we conducted further analysis adding total contribution volume as a control variable in a linear regression.

¹⁰https://en.wikipedia.org/wiki/Wikipedia:Protection_policy

To test **H4** about article quality, we compared the median quality of the two samples and tested for statistical significance using a Mann-Whitney U-test. The unit of analysis is the article, with our measure of average article quality as the outcome variable and taboo as our key predictor.

To evaluate **H5A** on the probability of someone contributing without an account varying between contributions to taboo and non-taboo articles, we fit a logistic regression on account status with taboo and a predictor and with protection as a control. The unit of analysis for this test is the revision. Because we have repeated measures of articles, we fit a multilevel model with a random intercept term for each article. To evaluate **H5B** about the relationship of an article contribution being to a taboo article to contributor experience level, we fit a logistic regression on the logged average contributor experience with taboo as a predictor and page protection percentage as a control. The unit of analysis for this analysis is the article.

The identity affordances used as outcomes in tests for **H5C**, **H5D**, and **H5E** (user page, gender, emailability) are only available to accountholders, so we restrict these analyses to contributions made by accountholders. For **H5C**, **H5D**, and **H5E**, we used a user-level dataset. Our analysis of **H5C** (user page) uses a dichotomous variable based on whether the contributor had a user page while making a contribution to any article in our sample. We use a chi-squared test to evaluate the relationship between having a user page and ever editing a taboo article. Our analysis of **H5D** (gender) uses two measures: first, someone must opt in to the gender-specifying feature, and second, they must choose a gender from the menu (the options given by the interface are “male” and “female”). We tested these aspects separately. We evaluated the relationship between specifying a gender, specifying a gender that is female, and being emailable using a chi-squared test.

6 RESULTS

6.1 Public Interest in Taboo Subjects (H1)

Our analysis of page views provides evidence in support of **H1** that taboo articles are viewed more frequently than those in our comparison set. The median view rank for taboo articles is 11,525 and the median view rank for the baseline comparison set is 39,017 ($U = 189,375$, $p < 0.001$). These data are represented in the boxplots in top panel of Figure 3. Note that the highest possible rank in this analysis is 1. We observe that although both sets have long tails including very unpopular articles, the higher popularity of taboo articles is such that the interquartile ranges of the two sets overlap only slightly.

6.2 Contribution Quantity (H2)

Contrary to our expectations in **H2**, our analysis provides evidence that the quantity of contributions to taboo subjects is substantially higher than to the comparison set. The median number of contributions for taboo articles is 1,620, and the median number of contributions for the comparison set is 143 ($U = 53,010$, $p < 0.001$). These data are represented in the boxplots in the second panel of Figure 3.

6.3 Contribution Quality (H3)

Our analysis of contribution quality provides support for **H3**. We find that contributions to taboo articles are more frequently reverted than contributions to those in our comparison set. The median revert rate for taboo articles is 26.3% and the median revert rate for the baseline comparison set is 10.1% ($U = 48,566$, $p < 0.001$). These data are represented in the boxplots in the third panel of Figure 3.

We also measured contribution quality using the ORES model predicting “damaging” revisions and found that revisions to taboo articles are more frequently damaging than those to articles in our

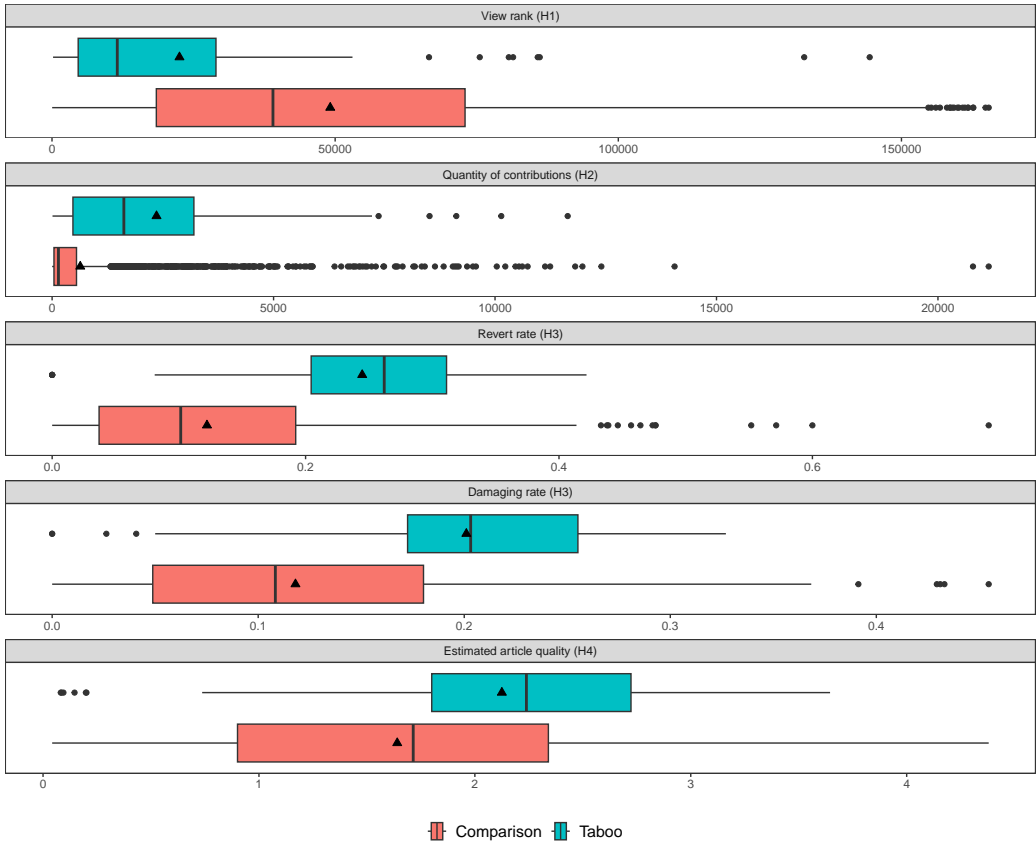


Fig. 3. Boxplots showing the distributions of article-level variables for five hypothesis tests. From top to bottom: (a) view rank of articles where view rank is calculated within a given month across all articles where the most viewed article would rank 1 (H1); (b) quantity of contributions (H2); (c) article-level revert rates (H3); (d) article-level damaging contribution rate (H3); and (e) quality of the articles (H4). Small vertical lines in the boxes indicate medians. Triangles are located at the mean.

comparison set. The median damaging contribution rate for taboo articles is 20.3% and the median damaging contribution rate for the baseline comparison set is 10.8%, ($U = 56,225, p < 0.001$). The distribution of the two sets are depicted in the fourth panel of Figure 3. These estimates are similar in size to our estimates for our first hypothesis test for H3.

Both of these results for H3 provide evidence in support of our hypothesis that taboo subjects are the target of lower-quality contributions. However, given that our test of H2 showed that the number of revisions to taboo articles is an order of magnitude higher than the number of revisions to articles in our comparison set, edit volume appears to be a possible confounder. Indeed, revert rate and overall quantity of revisions are moderately positively correlated ($\rho = 0.505, p < 0.001$). We report the results of our robustness check—a linear regression with a control for contribution count—in Table 2. We observe that our prior finding that taboo articles receive more damaging contributions than do non-taboo articles is robust to the control for overall contribution volume.

	Revert Count
(Intercept)	-27.3728*
	[-30.6155; -24.1300]
Taboo	49.5379*
	[29.9640; 69.1117]
Contribution Count	0.2791*
	[0.2771; 0.2812]
R ²	0.9590
Adj. R ²	0.9590
Num. obs.	3299

* 0 outside the confidence interval.

Table 2. Linear regression model estimating effects for taboo on article revert count, with control for overall contribution count (H3).

6.4 Article Quality (H4)

Opposite to our proposal in **H4**, we find that articles addressing taboo subjects are higher quality than those in our comparison sample. The median quality level for taboo articles is 2.2 while the median quality for the comparison set is 1.7 ($U = 87,747$, $p < 0.001$). These data are visualized in the bottom panel of Figure 3.

6.5 Contributor Identifiability (H5)

In H5A-E, we proposed that contributors to taboo subjects would be less identifiable. The results for H5A about users with accounts are shown in Table 3 and suggest that a revision to a taboo article is more likely to be made by a user without an account than a revision to the comparison set. The effect is as we hypothesized. Although the size of this effect is small ($<1\%$), a 0.05% change in the 24 million contributions received in December 2022 is 120,000 contributions.

	Contributing Without an Account
(Intercept)	-0.9948*
	[-1.0141; -0.9755]
Taboo	0.1746*
	[0.0547; 0.2944]
Protection Level	0.0791*
	[0.0351; 0.1231]
AIC	2749327
BIC	2749377
Log Likelihood	-1374659
Num. obs.	2219242
Num. groups: encodedTitle	3293
Var: encodedTitle (Intercept)	0.2639

* 0 outside the confidence interval.

Table 3. Hierarchical logistic regression model estimating the effects for taboo on whether or not a revision is made by someone contributing without an account (H5A). Fixed effects for the article are not shown.

In terms of H5B about editor experience, we find that the average editor to a taboo article has made fewer edits (21,621) than the average editor to our comparison set (36,176) ($U = 171,617$; $p < 0.001$).

Contribution Count Model	
(Intercept)	10.6406*
	[10.6122; 10.6689]
Taboo	-0.2318*
	[-0.4143; -0.0492]
Protection Level	-1.2622*
	[-1.4679; -1.0565]
R ²	0.0477
Adj. R ²	0.0471
Num. obs.	3299

* 0 outside the confidence interval.

Table 4. Results of a linear model examining the relationship between number of contributions a contributor has made (log scale) and whether they are contributing to a taboo subject with 95% confidence intervals (H5B).

Further, in a model where we control for page protection (which blocks non-registered and low-experience contributors), contributors to taboo articles still tend to have fewer contributions. This result is surprising because one might expect that including this control would have reversed this relationship; blocking the least experienced contributors seems like it would be associated with contributors with more contributions rather than fewer, however this is not the case. Page protection is also associated with fewer contributions, as expected. Table 4 shows the results of a linear model with taboo as a predictor of (logged) contribution count and protection level as a control.

Our analysis contradicts what we proposed in H5C. Of the 181,597 accountholders about whom we have user page data, 35.3% of them had user pages at the point when they made any of the contributions in our sample. We found that 51.9% of the contributors who ever contributed to a taboo article had a user page during one of their overall contributions to our samples, whereas only 32.8% of those who did not contribute to a taboo article had user pages while contributing to a page in our sample, a statistically significant difference ($\chi^2 = 3267.2 : p < 0.001$).

These results provide evidence that accountholders who contributed to a taboo subject in our sample were more likely to have a user page than contributors in our sample who did not contribute to a taboo subject in our sample.

Next, we observe that of the contributors who had ever contributed to the taboo articles in our sample, 17.7% of them specified a gender, while of the contributors who did not contribute to the taboo articles, 10.0% specified a gender ($\chi^2 = 1045.9 : p < 0.001$). This result contradicts **H5D** and suggests that people who contribute to taboo articles using an account are more likely to specify their gender. Out of curiosity, we also examined the relationship between taboo and the reported gender of those 19,678 individuals who did choose to report it. We found that 8.6% of contributors to taboo articles who specify their gender specified female as compared to 8.9% in the comparison. We found that among account-holding contributors who choose to specify their gender, the relationship between reporting one's gender as female and contributing to a taboo subject is not statistically significant ($\chi^2 = 0.345 : p = 0.557$).

Finally, we consider emailability. We find that 42.6% of contributors to taboo articles have made themselves emailable, while 38.4% of contributors to the comparison set have done so. We found that the relationship between editing a taboo article and setting oneself to be emailable is statistically

significant ($\chi^2 = 149.8 : p < 0.001$). This result is contrary to **H5E**; contributors to taboo subjects are more likely to be emailable rather than less.

Of the five measures we used to operationalize different facets of identifiability, two were in the hypothesized direction: contributors to taboo subjects are less likely to have accounts and have less experience. We were surprised to find that accountholders who contribute to taboo subjects were more likely to have user pages, to reveal their gender and to make themselves emailable.

7 DISCUSSION

7.1 The Success of Taboo Articles on Wikipedia

Wikipedia articles on taboo subjects are very popular. Although one might suspect that societal opprobrium would suppress the quality of articles about taboo subjects, our results suggest the opposite is true. On Wikipedia, articles about taboo subjects receive more contributions and are higher quality than other similar articles. However, these articles also receive more low quality contributions than non-taboo articles do. We investigated whether techniques for limiting identifiability served as an important factor in the success of these articles and concluded that this relationship is not a simple one. Although contributors to taboo subjects were less likely to use accounts and tended to have less experience when compared to those who contribute to non-taboo subjects, we also found that among the accountholding editors in our sample, those editing taboo subjects were more likely to have user pages, to specify their gender, and to make themselves emailable than those who did not edit taboo subjects.

One potential explanation for our results with respect to identifiability may be the phenomena described in Menking et al. [55] who described sophisticated means that women contributing to Wikipedia employ in order to contribute safely. Although Menking et al. [55] describes “choosing what to edit” as a safety strategy, they also describe personal qualities that allow women to persist in the face of attacks. Contributors to taboo subjects may conceive of themselves as “people who can take it” in the words of one of Menking et al.’s [55] interviewees.

Or, in that the success of peer production projects may be tied to the diversity of motives from participants [11], it may be that contributors to taboo subjects are specifically motivated with respect to the taboo subjects to which they choose to contribute. For example, they may take an advocacy position with respect to women’s health information. Investigating this subject further might require additional methods such as participant interviews. Of course, an interview approach may struggle to include those who contribute without accounts and more casual contributors who may be harder to reach and recruit for interviews.

Taken as a whole, our findings introduce an empirical puzzle: How is it that taboo articles are higher quality than non-taboo articles despite the larger number and rate of poor quality and damaging contributions that these articles receive from less identifiable contributors? One possible explanation may be a mechanism Gorbatai [29] elaborates: perhaps these unhelpful contributions serve as a signal of public interest in a subject and draw in experienced editors to clean up the mess left by less experienced contributors and who also, perhaps, further improve the article while doing so. Our result provides further evidence that the end result of unhelpful contributions may depend on how the community responds, and reinforces the observation from Hill and Shaw [41] that increasing barriers to novice contributions (which may indeed be lower quality) may have deleterious effects.

7.2 Further Exploration of Article Quality

Given our surprising finding on the quality of articles about taboo subjects, we conducted additional analysis. Figure 4 visualizes the quality growth trajectory of taboo and comparison articles over




Fig. 4. Visualization of average article quality over time as predicted by the Wikimedia ORES API shown using generalized additive model (GAM) smoothers. We see that in the first several years of their existence, taboo subjects grow somewhat more quickly in quality, but that their quality growth over time begins to track more closely to the comparison set.




Fig. 5. Average quality of the first version of new articles over time shown using generalized additive model (GAM) smoothers. The rug along the axes identifies the areas with the greatest concentration of data.

time using local regression (LOESS) models. Quality is as predicted by the Wikimedia ORES model on a per-article-per-month level. Articles in these two sets are different early in their existence but generally follow similar trajectories after several years. Figure 5 shows the average quality of

articles when they are created. The initial quality of taboo articles and comparison articles in the early period of Wikipedia were fairly similar, but recently created articles are higher quality at the point of creation. Additionally, we have included a cross-tabulation of whether contributors were accountholders, which sample they appear in, and whether their contribution was reverted, in our online supplement. These results suggest that one potential explanation for the high quality of taboo articles is survival bias. For an article to survive long enough to enter our dataset, especially those articles created recently, it must have had sufficient quality when created in order to convey its value. Initial quality and growth place taboo articles further ahead in their development, even as maintenance and enhancement of both groups of articles eventually follow similar trajectories. One explanation for the preliminary differences we observe may be that individual editors exercising quality control in Wikipedia may scrutinize new taboo articles even more closely than they do other articles, perhaps drawn in by their own interest in taboo subjects or a sense that taboo subjects may draw in the creation of spurious articles.

Another part of the explanation for the quality of taboo articles may be the presence of health-related content in the taboo set. Wikipedia has specialized guidelines around health and medicine sourcing and is home to an organized effort to uphold those guidelines (WikiProject Medicine) [43, 84]. However, this is a question in need of additional research because many such organizations and content guidelines exist in Wikipedia, and because the existence of strict guidelines and the presence of a group enforcing those guidelines could serve to decrease the proportion of contributions ultimately deemed acceptable.

7.3 Implications For Design

Our results indicate that there is substantial reason for optimism with respect to the availability of knowledge about taboo subjects and that online communities like Wikipedia can, in some cases, resist societal norms in ways that support human thriving. This result is particularly striking because our setting is Wikipedia. As a tertiary source, an encyclopedia is dependent on the availability of secondary content in the form of publications from reliable sources. There may be much more we can learn from this success as we seek to replicate it in other settings.

Contributors to taboo subjects contradicted our expectations of low quality and maximal privacy. Our finding that contributors to taboo subjects used some privacy affordances but not others suggests that collaborative platforms should be wary of assuming that participants have a uniform perspective on their privacy. Indeed some may prefer to make themselves more identifiable despite the potential increased risk. However, policy changes that would increase requirements for identifiability may endanger this success, at least with respect to the contributors who appear in our data as non-accountholders or who have a lower overall edit count. Many peer production projects, including some language editions of Wikipedia, have policies requiring that participants make themselves even more identifiable than they are required to be on English Wikipedia. Our finding that contributors take a variable approach to privacy suggests designers exercise caution when considering privacy-sacrificing features because a lack of privacy may limit the participation of people seeking to engage with taboo subjects. Further, these results suggest that future studies of taboo in social computing systems should examine privacy-seeking behaviors among different roles (e.g., reader, contributor, administrator).

We also observe that identifiability is multifaceted. In this case, we find differences even in these relatively weak identifiers: these are not “real name” policies, user pages could contain content of any kind, none of the information being provided is verified, and although participants can opt into emailability, doing so does not expose their email address. Affordances that seek to make participants more identifiable on platforms have been found to have unwanted side effects that go beyond discouraging participation [41]. Additionally, diminishing contributor choices on what

information about them is collected and shared may diminish their willingness to contribute to subjects that may place them at risk. In turn, this could lead to lower-quality, neglected articles. Such an outcome would disproportionately harm information-seekers most in need of reliable information.

7.4 Implications for Future Research

Our methods have potential utility for future studies of taboo across HCI and across languages. Although we selected English corpora, Wiktionary exists in 186 different languages and Wikipedia in more than 300. Further, although there may be some advantage to the fact that both Wiktionary and Wikipedia operate under the same umbrella foundation and have overlapping communities, other dictionaries and sites of collaborative production could be examined in a similar manner.

Numerous other dictionaries use sense tagging to indicate that a definition corresponds to euphemistic usage. Further, there are many other sense tags in use, including the identification of slang, expletives, epithets, jargon, dialect, and archaic usage. Connecting linguistic scholarship about these forms of language to other social computing phenomena may be fruitful. Further, a computational approach of this kind is scalable across multiple settings, including those where human inspection is infeasible because of volume. We have made an anonymized dataset and our analytical code available for replication and reuse of our findings at <https://doi.org/10.7910/DVN/5OKEEO>.

8 LIMITATIONS

Our method of taboo identification through NLP techniques lacks important nuance and context. Because taboo varies across cultures, behaviors that are profoundly polluting for some may carry little or no taboo for others. We are researchers with relatively high levels of privilege with respect to freedom of inquiry; this positionality has inevitably influenced our results. Our profession, discipline, and employers act to legitimize our engagement with taboo subjects, which may in turn desensitize us to the risks of violating taboo.

Our ability to generalize these results to other contexts is limited by the fact that we only examined English language Wikipedia. This limitation reflects our language knowledge, given the deeply linguistic nature of the subject matter. Given that taboo is substantially influenced by local and cultural factors, our findings might be different in other language contexts. We hope that authors with a broader range of language backgrounds will take up similar questions in the future.

Additionally, contributions to Wiktionary, from which we drew our original list of key words indicating taboo, may be biased. We know that Wikipedia suffers from numerous participation gaps with respect to gender, skills, and race and we imagine these gaps might extend to Wiktionary as well in ways that might shape the kinds of taboo that is reflected in our dataset [34, 35, 38, 71]. Although we were unable to find data on participation gaps in Wiktionary, surveys conducted across all Wikimedia Foundation projects (which include Wikipedia and Wiktionary) have found gaps in participation.¹¹ Given that taboo functions as social control, insofar as similar contributor demographic gaps are present in Wiktionary, our identification of taboo subjects likely reflects the attitudes of the people who are more privileged and have more means to resist the consequences of taboo violations. Although this remains an important unaddressed threat to validity, we believe that it likely makes our method of identifying taboo conservative, if potentially biased.

Further, Wiktionary has limitations as a source of data on taboo in that its self-organized volunteer contributor nature and openness to non-experts may lead to inconsistencies; euphemistic definitions may go untagged while non-euphemisms may be erroneously tagged as such. For example, our taboo dataset includes the article *Super Bowl*. The bigram “super bowl” is in our taboo n-grams

¹¹https://meta.wikimedia.org/wiki/Community_Insights/Community_Insights_2021_Report

list because out of the five times it appears in definitions in Wiktionary, two of them (the terms “superb owl” and “Big Game”) have definitions that are tagged as euphemism. Although we would argue that this is not a correct usage of the euphemism tag (instead, these terms seem more like slang), we have included the Super Bowl article out of fidelity to the source material. Taking on this descriptivist approach can add noise, but can also expand our perspective. When we observed that our dataset included euphemisms around nuclear weapons and napalm, we realized that these topics were not what we had been expecting when we first began to explore taboo. Indeed, they might have escaped our notice. However, once sensitized to this taboo through our systematic approach, we were able to validate it and found that prior work has indeed explored the taboo surrounding nuclear and chemical weapons [20, 62, 69].

Our measures are limited as well. Views data only examines page hits—not whether the article was read. Contributions vary in their size, and their assessed quality is subject to community values and biases. The measure of article quality we have used, although common in Wikipedia research, only assesses articles based on observable features such as length, links, and the presence of images: it does not gauge content, and collapses quality down to a single continuous measure despite being derived from an ordinal system of quality classes that are not evenly spaced [78].

We used several measures to assess contributor identifiability. By constructing a user-level dataset with dichotomous variables for H5C, D, and E, we omit variation over time, and we do not consider the inequalities in participation between participants; some descriptive analysis of a revision-level dataset is in our online supplement. Although we use the presence of profiles as a sign of identifiability, people tell “privacy lies” by entering false information into online profiles—sometimes in order to avoid being stigmatized [68]. We make no effort to address whether user pages contain identifying information or whether the information provided is correct. Selecting a gender from a drop-down menu may or may not align with aspects of someone’s gender identity. Additionally, although we treat contributing without an account as a sign of lower identifiability, the IP addresses made visible as part of that contribution can be highly identifying. Finally, we only have information on whether these traits were present at the time we made our API queries in June 2022, and do not have historical information as to whether gender or emailability had been set at the moment of contribution. We attempt to mitigate these limitations through our use of several different measures to assess identifiability.

Finally, although our method seeks to minimize our assumptions about what is taboo, we made many choices throughout the process that may have influenced our results. We chose stop words. We relied on our own sense of what qualifies as taboo when checking our results for face validity during all points of conducting this research project. Although we believe that these ongoing informal assessments were consistent with the work of prior researchers, our own perspective inevitably entered into the conduct of this work.

9 CONCLUSION

Information-seekers turn to social computing systems to learn about subjects that society has gauged to be taboo, and Wikipedia contributors have successfully built resources to serve this need. This work makes three contributions: (1) we elaborate a computational technique for detecting taboo based on euphemistic dictionary definitions, (2) we illustrate this approach by applying it to English Wikipedia, and (3) we analyze distinctions in both how taboo subjects develop and who contributes to the development of information about taboo subjects. Future work should explore the relationships between identifiability and taboo subjects in other environments, continue to unpack the social dynamics of taboo, and consider the impacts of policy changes that diminish access to privacy-preserving technologies. Taboo subjects connect to some of the most fundamental pieces of human existence and both some of the most distressing and most uplifting parts of our

social experience. Although we as researchers may recoil from doing work on “not safe for work” subjects, facing taboos and tackling them as area of inquiry is worth the discomfort it may give us. Taboos shape our consumption and production behaviors across social computing systems. Exploring taboo subjects through the lens of HCI is a vital and transformative area of future work for our field.

ACKNOWLEDGMENTS

This work benefited substantively from the guidance of anonymous reviewers and the program committee at CSCW. The creation of our dataset was aided by the use of advanced computational, storage, and networking infrastructure provided by the Hyak supercomputer system at the University of Washington and high-performance computing systems supported by the Northwestern School of Communication and the National Science Foundation. This work was supported by the National Science Foundation (awards CNS-1703736 and CNS-1703049).

10 DATA AND CODE

A replication dataset including data, code, and other supplementary material has been placed in the Harvard Dataverse archive and is available at: <https://doi.org/10.7910/DVN/5OKEEO>

REFERENCES

- [1] 2022. Wikipedia:Wikipedia Signpost/2022-04-24/In focus. https://en.wikipedia.org/w/index.php?title=Wikipedia:Wikipedia_Signpost/2022-04-24/In_focus&oldid=1090495808
- [2] Keith Allan. 2019. Taboo words and language: An overview. In *The Oxford handbook of taboo words and language* (first edition ed.), Keith Allan (Ed.). Oxford University Press, Oxford ; New York, NY.
- [3] Teresa Almeida, Rob Comber, Gavin Wood, Dean Saraf, and Madeline Balaam. 2016. On Looking at the Vagina through Labella. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 1810–1821. <https://doi.org/10.1145/2858036.2858119>
- [4] Tawfiq Ammari, Sarita Schoenebeck, and Daniel Romero. 2019. Self-declared throwaway accounts on reddit: how platform affordances and shared norms enable parenting disclosure and support. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 135:1–135:30. <https://doi.org/10.1145/3359237>
- [5] Nazanin Andalibi. 2020. Disclosure, Privacy, and Stigma on Social Media: Examining Non-Disclosure of Distressing Experiences. *ACM Transactions on Computer-Human Interaction* 27, 3 (June 2020), 1–43. <https://doi.org/10.1145/3386600>
- [6] Nazanin Andalibi, Oliver L. Haimson, Munmun De Choudhury, and Andrea Forte. 2016. Understanding Social Media Disclosures of Sexual Abuse Through the Lenses of Support Seeking and Anonymity. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 3906–3918. <https://doi.org/10.1145/2858036.2858096>
- [7] Denise L. Anthony, Sean W. Smith, and Timothy Williamson. 2009. Reputation and reliability in collective goods: The case of the online encyclopedia Wikipedia. *Rationality and Society* 21, 3 (Aug. 2009), 283–306. <https://doi.org/10.1177/1043463109336804>
- [8] Sumit Asthana and Aaron Halfaker. 2018. With Few Eyes, All Hoaxes Are Deep. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (Nov. 2018), 21:1–21:18. <https://doi.org/10.1145/3274290>
- [9] Jessica Aucoin. 2021. Censorship in Libraries: A Retrospective Study of Banned and Challenged Books. *SLIS Connecting* 10, 2 (2021), 82–100. <https://doi.org/10.18785/slis.1002.09>
- [10] Yochai Benkler. 2002. Coase's Penguin, or, Linux and "The Nature of the Firm". *The Yale Law Journal* 112, 3 (Dec. 2002), 369. <https://doi.org/10.2307/1562247>
- [11] Yochai Benkler. 2006. *The wealth of networks: How social production transforms markets and freedom*. Yale University Press, New Haven, CT.
- [12] Jeremy Birnholtz, Nicholas Aaron Ross Merola, and Arindam Paul. 2015. "Is it Weird to Still Be a Virgin": Anonymous, Locally Targeted Questions on Facebook Confession Boards. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, Seoul Republic of Korea, 2613–2622. <https://doi.org/10.1145/2702123.2702410>
- [13] Bri and EpicPupper. 2022. Wikipedia:Wikipedia Signpost/2022-03-27/News and notes. https://en.wikipedia.org/w/index.php?title=Wikipedia:Wikipedia_Signpost/2022-03-27/News_and_notes&oldid=1084492254 Page Version ID: 1084492254.

- [14] Taylor Kate Brown. 2016. Female scientist fights harassment with Wikipedia. *BBC News* (March 2016). <https://www.bbc.com/news/blogs-trending-35787730>
- [15] Sebastian Brückner, Florian Lemmerich, and Markus Strohmaier. 2021. Inferring Sociodemographic Attributes of Wikipedia Editors: State-of-the-art and Implications for Editor Privacy. In *Companion Proceedings of the Web Conference 2021*. ACM, Ljubljana Slovenia, 616–622. <https://doi.org/10.1145/3442442.3452350>
- [16] Kate Burridge. 2015. Taboo Words. In *The Oxford Handbook of the Word*, John R. Taylor (Ed.). <https://doi.org/10.1093/oxfordhb/9780199641604.013.017>
- [17] Davide Buscaldi and Irazú Hernandez-Farias. 2015. Sentiment Analysis on Microblogs for Natural Disasters Management: a Study on the 2014 Genoa Floodings. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15 Companion)*. Association for Computing Machinery, New York, NY, USA, 1185–1188. <https://doi.org/10.1145/2740908.2741727>
- [18] Nadia Campo Woytuk, Marie Louise Juul Søndergaard, Marianela Ciolfi Felice, and Madeline Balaam. 2020. Touching and Being in Touch with the Menstruating Body. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376471>
- [19] Jilin Chen, Yuqing Ren, and John Riedl. 2010. The effects of diversity on group productivity and member withdrawal in online volunteer groups. In *Proceedings of the 28th international conference on Human factors in computing systems (CHI '10)*. ACM, New York, NY, USA, 821–830. <https://doi.org/10.1145/1753326.1753447>
- [20] Leonard A. Cole. 1998. The Poison Weapons Taboo: Biology, Culture, and Policy. *Politics and the Life Sciences* 17, 2 (Sept. 1998), 119–132. <https://doi.org/10.1017/S0730938400012119>
- [21] Benjamin Collier and Julia Bear. 2012. Conflict, criticism, or confidence: an empirical examination of the gender gap in Wikipedia contributions. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work (CSCW '12)*. ACM, New York, NY, USA, 383–392. <https://doi.org/10.1145/2145204.2145265>
- [22] Eliecer Crespo Fernández. 2011. Conceptual metaphors in taboo-induced lexical variation. *Revista Alicantina de Estudios Ingleses* 24 (Nov. 2011), 53. <https://doi.org/10.14198/raei.2011.24.03>
- [23] Mary Douglas. 1978. *Purity and danger: an analysis of the concepts of pollution and taboo* (repr ed.). Routledge, London.
- [24] Heather Ford, Shilad Sen, David R. Musicant, and Nathaniel Miller. 2013. Getting to the Source: Where Does Wikipedia Get Its Information from?. In *Proceedings of the 9th International Symposium on Open Collaboration (WikiSym '13)*. ACM, New York, NY, USA, 9:1–9:10. <https://doi.org/10.1145/2491055.2491064>
- [25] Heather Ford and Judy Wajcman. 2017. ‘Anyone can edit’, not everyone does: Wikipedia’s infrastructure and the gender gap. *Social Studies of Science* 47, 4 (2017), 511–527. <https://doi.org/10.1177/0306312717692172>
- [26] 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>
- [27] R. Stuart Geiger and Aaron Halfaker. 2017. Operationalizing conflict and cooperation between automated software agents in wikipedia: a replication and expansion of ‘even good bots fight’. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (Dec. 2017), 1–33. <https://doi.org/10.1145/3134684>
- [28] Erving Goffman. 1963. *Stigma: Notes on the management of spoiled identity*. Simon and Schuster.
- [29] Andreea D. Gorbatai. 2014. *The Paradox of Novice Contributions to Collective Production: Evidence from Wikipedia*. Technical Report. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1949327
- [30] Shane Greenstein and Feng Zhu. 2016. Open Content, Linus’ Law, and Neutral Point of View. *Information Systems Research* 27, 3 (Sept. 2016), 618–635. <https://doi.org/10.1287/isre.2016.0643>
- [31] Mordechai (Muki) Haklay, Sofia Basiouka, Vyrion Antoniou, and Aamer Ather. 2010. How Many Volunteers Does it Take to Map an Area Well? The Validity of Linus’ Law to Volunteer Geographic Information. *The Cartographic Journal* 47, 4 (Nov. 2010), 315–322. <https://doi.org/10.1179/000870410X12911304958827>
- [32] Aaron Halfaker. 2017. Interpolating Quality Dynamics in Wikipedia and Demonstrating the Keilana Effect. ACM Press, 1–9. <https://doi.org/10.1145/3125433.3125475>
- [33] 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>
- [34] Eszter Hargittai and Kaitlin Jennrich. 2016. The Online Participation Divide. In *The Communication Crisis in America, And How to Fix It*, Mark Lloyd and Lewis A. Friedland (Eds.). Palgrave Macmillan US, New York, 199–213. https://doi.org/10.1057/978-1-349-94925-0_13
- [35] Eszter Hargittai and Aaron Shaw. 2015. Mind the skills gap: the role of Internet know-how and gender in differentiated contributions to Wikipedia. *Information, Communication & Society* 18, 4 (April 2015), 424–442. <https://doi.org/10.1080/1369118X.2014.957711>

- [36] Brent Hecht and Darren Gergle. 2010. The Tower of Babel meets Web 2.0: user-generated content and its applications in a multilingual context. In *Proceedings of the SIGCHI conference on human factors in computing systems (CHI '10)*. ACM, Atlanta, Georgia, USA, 291–300. <https://doi.org/10.1145/1753326.1753370>
- [37] Karey Helms. 2019. Do You Have to Pee?: A Design Space for Intimate and Somatic Data. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. ACM, San Diego CA USA, 1209–1222. <https://doi.org/10.1145/3322276.3322290>
- [38] Benjamin Mako Hill and Aaron Shaw. 2013. The Wikipedia gender gap revisited: characterizing survey response bias with propensity score estimation. *PLoS ONE* 8, 6 (June 2013), e65782. <https://doi.org/10.1371/journal.pone.0065782>
- [39] Benjamin Mako Hill and Aaron Shaw. 2014. Consider the redirect: a missing dimension of Wikipedia research. In *Proceedings of The International Symposium on Open Collaboration (OpenSym '14)*. ACM, New York, NY, USA, 28:1–28:4. <https://doi.org/10.1145/2641580.2641616>
- [40] Benjamin Mako Hill and Aaron Shaw. 2015. Page protection: another missing dimension of Wikipedia research. In *Proceedings of the 11th International Symposium on Open Collaboration (OpenSym '15)*. ACM, New York, NY, USA, 15:1–15:4. <https://doi.org/10.1145/2788993.2789846>
- [41] Benjamin Mako Hill and Aaron Shaw. 2021. The Hidden Costs of Requiring Accounts: Quasi-Experimental Evidence From Peer Production. *Communication Research* 48, 6 (Aug. 2021), 771–795. <https://doi.org/10.1177/0093650220910345>
- [42] Julia Jacobs. 2019. Wikipedia Isn't Officially a Social Network. But the Harassment Can Get Ugly. *The New York Times* (April 2019). <https://www.nytimes.com/2019/04/08/us/wikipedia-harassment-wikimedia-foundation.html>
- [43] Richard James. 2016. WikiProject Medicine: Creating Credibility in Consumer Health. *Journal of Hospital Librarianship* 16, 4 (Oct. 2016), 344–351. <https://doi.org/10.1080/15323269.2016.1221284>
- [44] 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)*. Association for Computing Machinery, New York, NY, USA, 2657–2666. <https://doi.org/10.1145/2470654.2481368>
- [45] Gopinaath Kannabiran, Jeffrey Bardzell, and Shaowen Bardzell. 2011. How HCI talks about sexuality: discursive strategies, blind spots, and opportunities for future research. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*. ACM Press, Vancouver, BC, Canada, 695. <https://doi.org/10.1145/1978942.1979043>
- [46] Sejal Khatri, Aaron Shaw, Sayamindu Dasgupta, and Benjamin Mako Hill. 2022. The social embeddedness of peer production: A comparative qualitative analysis of three Indian language Wikipedia editions. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–18. <https://doi.org/10.1145/3491102.3501832>
- [47] Gary King, Jennifer Pan, and Margaret E. Roberts. 2013. How censorship in China allows government criticism but silences collective expression. *American Political Science Review* 107, 2 (May 2013). <https://doi.org/10.1017/S0003055413000014>
- [48] Robert E. Kraut, Paul Resnick, and Sara Kiesler. 2012. *Building successful online communities: Evidence-based social design*. MIT Press, Cambridge, MA.
- [49] Amanda Lazar, Norman Makoto Su, Jeffrey Bardzell, and Shaowen Bardzell. 2019. Parting the Red Sea: Sociotechnical Systems and Lived Experiences of Menopause. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–16. <https://doi.org/10.1145/3290605.3300710>
- [50] Eden Litt. 2012. *Knock, Knock*. Who's There? The Imagined Audience. *Journal of Broadcasting & Electronic Media* 56, 3 (July 2012), 330–345. <https://doi.org/10.1080/08838151.2012.705195>
- [51] Bronislaw Malinowski. 1926. *Crime and custom in savage society*. Harcourt Brace & Company, New York.
- [52] A. E. Marwick and danah boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13, 1 (Feb. 2011), 114–133. <https://doi.org/10.1177/1461444810365313>
- [53] Gary T. Marx. 2016. *Windows into the soul: surveillance and society in an age of high technology*. The University of Chicago Press, Chicago ; London.
- [54] Matthew S. McGlone and Jennifer A. Batchelor. 2003. Looking Out for Number One: Euphemism and Face. *Journal of Communication* 53, 2 (June 2003), 251–264. <https://doi.org/10.1111/j.1460-2466.2003.tb02589.x>
- [55] Amanda Menking, Ingrid Erickson, and Wanda Pratt. 2019. People who can take it: how women Wikipedians negotiate and navigate safety. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, Glasgow, Scotland, UK, 472:1–472:14. <https://doi.org/10.1145/3290605.3300702>
- [56] Marc Miquel-Ribé and David Laniado. 2020. The Wikipedia Diversity Observatory: A Project to Identify and Bridge Content Gaps in Wikipedia. In *Proceedings of the 16th International Symposium on Open Collaboration*. ACM, Virtual conference Spain, 1–4. <https://doi.org/10.1145/3412569.3412866>
- [57] Aparna Moitra, Naemul Hassan, Manash Kumar Mandal, Mansurul Bhuiyan, and Syed Ishtiaque Ahmed. 2020. Understanding the Challenges for Bangladeshi Women to Participate in #MeToo Movement. *Proceedings of the ACM on Human-Computer Interaction* 4, GROUP (Jan. 2020), 1–25. <https://doi.org/10.1145/3375195>
- [58] Jonathan T. Morgan, Michael Gilbert, David W. McDonald, and Mark Zachry. 2013. Project talk: coordination work and group membership in WikiProjects. In *Proceedings of the 9th International Symposium on Open Collaboration*. ACM,

- Hong Kong China, 1–10. <https://doi.org/10.1145/2491055.2491058>
- [59] Dawn Nafus. 2012. ‘Patches don’t have gender’: What is not open in open source software. *New Media & Society* 14, 4 (June 2012), 669–683. <https://doi.org/10.1177/1461444811422887>
- [60] Emma Paling. 2015. Wikipedia’s Hostility to Women. <https://www.theatlantic.com/technology/archive/2015/10/how-wikipedia-is-hostile-to-women/411619/> Section: Technology.
- [61] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research* 12, 85 (Oct. 2011), 2825–2830. <http://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html>
- [62] Richard M. Price. 2007. *The Chemical Weapons Taboo*. Cornell University Press. <https://www.degruyter.com/document/doi/10.7591/9781501729546/html>
- [63] Rifat Rahman, Md. Rishadur Rahman, Nafis Irtiza Tripto, Mohammed Eunus Ali, Sajid Hasan Apon, and Rifat Shahriyar. 2021. Adolescentbot: understanding opportunities for chatbots in combating adolescent sexual and reproductive health problems in bangladesh. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–15. <https://doi.org/10.1145/3411764.3445694>
- [64] Anand Rajaraman and Jeffrey David Ullman. 2011. *Mining of Massive Datasets*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9781139058452>
- [65] Eric Raymond. 2001. *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*. O’Reilly & Associates, Sebastopol, CA.
- [66] Joseph Reagle. 2013. “Free as in sexist?” Free culture and the gender gap. *First Monday* 18, 1 (2013). <http://firstmonday.org/ojs/index.php/fm/article/view/4291>
- [67] Marian-Andrei Rizoiu, Lexing Xie, Tiberio Caetano, and Manuel Cebrian. 2016. Evolution of Privacy Loss in Wikipedia. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. ACM, San Francisco California USA, 215–224. <https://doi.org/10.1145/2835776.2835798>
- [68] Shruti Sannon, Natalya N. Bazarova, and Dan Cosley. 2018. Privacy Lies: Understanding How, When, and Why People Lie to Protect Their Privacy in Multiple Online Contexts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–13. <https://doi.org/10.1145/3173574.3173626>
- [69] Thomas C. Schelling. 2007. The Nuclear Taboo - Spring 2007 - MIT International Review. <http://web.mit.edu/mitir/2007/spring/taboo.html>
- [70] Bryan Semaan, Bryan Dosono, and Lauren M. Britton. 2017. Impression Management in High Context Societies: ‘Saving Face’ with ICT. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, Portland Oregon USA, 712–725. <https://doi.org/10.1145/2998181.2998222>
- [71] Aaron Shaw and Eszter Hargittai. 2018. The Pipeline of Online Participation Inequalities: The Case of Wikipedia Editing. *Journal of Communication* 68, 1 (Feb. 2018), 143–168. <https://doi.org/10.1093/joc/jqx003>
- [72] Victoria Song. 2022. A top Wikipedia editor has been arrested in Belarus. <https://www.theverge.com/2022/3/11/22973293/wikipedia-editor-russia-belarus-ukraine>
- [73] Piya Sorcar, Benjamin Strauber, Prashant Loyalka, Neha Kumar, and Shelley Goldman. 2017. Sidestepping the Elephant in the Classroom: Using Culturally Localized Technology To Teach Around Taboos. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 2792–2804. <https://doi.org/10.1145/3025453.3025958>
- [74] John Suler. 2004. The online disinhibition effect. *CyberPsychology & Behavior* 7, 3 (June 2004), 321–326. <https://doi.org/10.1089/1094931041291295>
- [75] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. 2020. Censored Planet: An Internet-wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Virtual Event USA, 49–66. <https://doi.org/10.1145/3372297.3417883>
- [76] Marie Louise Juul Søndergaard, Marianela Ciolfi Felice, and Madeline Balaam. 2021. Designing Menstrual Technologies with Adolescents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–14. <https://doi.org/10.1145/3411764.3445471>
- [77] Lee Taber and Steve Whittaker. 2020. “On Finsta, I can say ‘Hail Satan’”: Being Authentic but Disagreeable on Instagram. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–14. <https://doi.org/10.1145/3313831.3376182>
- [78] Nathan TeBlunthuis. 2021. Measuring Wikipedia Article Quality in One Dimension by Extending ORES with Ordinal Regression. In *17th International Symposium on Open Collaboration*. ACM, Online, Spain. <https://doi.org/10.1145/3479986.3479991>
- [79] Nathan TeBlunthuis, Benjamin Mako Hill, and Aaron Halfaker. 2021. Effects of Algorithmic Flagging on Fairness: Quasi-experimental Evidence from Wikipedia. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 56:1–56:27. <https://doi.org/10.1145/3449130>

- [80] Jacob Thebault-Spieker, Brent Hecht, and Loren Terveen. 2018. Geographic Biases are ‘Born, not Made’: Exploring Contributors’ Spatiotemporal Behavior in OpenStreetMap. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork*. ACM Press, 71–82. <https://doi.org/10.1145/3148330.3148350>
- [81] Katherine Thornton and David W. McDonald. 2012. Tagging Wikipedia: collaboratively creating a category system. In *Proceedings of the 17th ACM international conference on Supporting group work - GROUP '12*. ACM Press, Sanibel Island, Florida, USA, 219. <https://doi.org/10.1145/2389176.2389210>
- [82] Eran Toch, Yang Wang, and Lorrie Faith Cranor. 2012. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* 22, 1-2 (April 2012), 203–220. <https://doi.org/10.1007/s11257-011-9110-z>
- [83] Chau Tran, Kaylea Champion, Andrea Forte, Benjamin Mako Hill, and Rachel Greenstadt. 2020. Are anonymity-seekers just like everybody else? An analysis of contributions to Wikipedia from Tor. In *2020 IEEE Symposium on Security and Privacy (SP)*, Vol. 1. IEEE Computer Society, San Francisco, California, 974–990. <https://doi.org/10.1109/SP40000.2020.00053>
- [84] L. Trevena. 2011. WikiProject Medicine. *BMJ* 342, jun08 3 (June 2011), d3387–d3387. <https://doi.org/10.1136/bmj.d3387>
- [85] Francesca Tripodi. 2021. Ms. Categorized: Gender, notability, and inequality on Wikipedia. *New Media & Society* (June 2021), 146144482110237. <https://doi.org/10.1177/14614448211023772>
- [86] Anupriya Tuli, Shaan Chopra, Neha Kumar, and Pushpendra Singh. 2018. Learning from and with Menstrupedia: Towards Menstrual Health Education in India. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 1–20. <https://doi.org/10.1145/3274443>
- [87] Nicholas Vincent, Isaac Johnson, and Brent Hecht. 2018. Examining Wikipedia With a Broader Lens: Quantifying the Value of Wikipedia’s Relationships with Other Large-Scale Online Communities. ACM Press, 1–13. <https://doi.org/10.1145/3173574.3174140>
- [88] Morten Warncke-Wang, Vivek Ranjan, Loren Terveen, and Brent Hecht. 2015. Misalignment between supply and demand of quality content in peer production communities. In *Proceedings of the Ninth International AAAI Conference on Web and Social Media (ICWSM '15)*. 493–502.
- [89] Diyi Yang, Aaron Halfaker, Robert Kraut, and Eduard Hovy. 2017. Identifying Semantic Edit Intentions from Revisions in Wikipedia. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (2017)*, 2000–2010.
- [90] Jonathan L. Zittrain, Robert Faris, Helmi Noman, Justin Clark, Casey Tilton, and Ryan Morrison-Westphal. 2017. The Shifting Landscape of Global Internet Censorship. <https://doi.org/10.2139/ssrn.2993485>

Received July 2022; revised January 2023; accepted March 2023