

**Wikipedia @ 20**

# **An Encyclopedia with Breaking News**

**Brian Keegan**

**Published on:** Jun 20, 2019

**Updated on:** Sep 05, 2019

**License:** [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Image credit: [Branden Harvey on Unsplash](#)

---

## Introduction

The web was a very different place for news in the United States between 2001 and 2006. The hanging chads from the 2000 Presidential election, the spectacular calamity of 9/11, the unrepentant lies around Operation Iraqi Freedom, and campy reality television featuring Donald Trump were all of this time. The burst of the dot-com bubble and corporate malfeasance of companies like Enron dampened entrepreneurial spirits, news publishers were optimistically sharing their stories online without paywalls, and blogging was heralded as the future of technology-mediated accountability and participatory democracy. “You” was TIME Magazine's Person of the Year in 2006 because “Web 2.0” platforms like YouTube, MySpace, and Second Life had become tools for “bringing together the small contributions of millions of people and making them matter” [1].

Wikipedia was a part of this primordial soup predating newsfeed-mediated engagement, recommender-driven polarization, politicized content moderation, and geopolitical disinformation campaigns. From very early in its history, Wikipedia leveraged the supply and demand for information about breaking news and current events into strategies that continue to sustain this radical experiment in online peer production. This chapter will explore Wikipedia's earliest efforts to cover breaking news events, common features of these unique collaborations, and how these features may serve as a model for other social platforms grappling with problems like disinformation.

I first encountered Wikipedia as an undergraduate student around 2004. My introduction to Wikipedia was likely a product of the socio-technical coupling between Google and Wikipedia during this era. Google helped Wikipedia because Google's ranking algorithms privileged Wikipedia's highly interlinked articles, which brought an influx of users, some (tiny) fraction of whom became contributing editors like myself. Wikipedia also helped Google because Wikipedia could reliably generate both general interest and up-to-date content that satisfied its users' information seeking needs, which brought users back to Google rather than its competitors. The aftermath of natural disaster, death of a celebrity, or a new pop culture sensation are all occasions for people to seek out background information to help them make sense of these events. Traditional journalistic offerings provide incremental updates about the immediate subject but often lack context or background: Why are there earthquakes in Indonesia? Who is Saddam Hussein? What is Eurovision? The availability and timeliness of Wikipedia content around topics of general interest would prove to be critical for its own sustainability in addition to complementing other platforms' need to serve relevant and up-to-date content.

Wikipedia also entered the popular awareness of undergraduates like myself through the pitiless warnings from instructors and librarians about its lack of reliability as a citation. Popular coverage of Wikipedia amplified these anxieties into an eye-rolling moral panic: the hapless youth and unwashed masses were accessing—even contributing to!—knowledge without adequate supervision. Much of the earliest research on Wikipedia were motivated by similar misgivings about its quality and reliability and compared its topical coverage, choice of citations, and factual errors to standards like *Britannica*. While these anxieties were largely reversed through empirical research and changes in professional culture, they also missed the forest for the trees: the value and authority of Wikipedia was not in any single article's quality, but in its network of hyperlinked articles. The "Wikipedia Game" where users attempt to traverse the shortest path between two random articles captured the common user experience of wandering Wikipedia's knowledge graph until you forget how you arrived at article like "Succession to the British throne." More than synthesizing knowledge as a tertiary source like traditional encyclopedias, Wikipedia's hyperlink network invited users to follow their interests, dive deeper on topics, introduce missing connections, and create new articles where none existed. Where the decentralized web created a fragmented user experience requiring directories (*e.g.*, Yahoo!) and search engines (*e.g.*, Google) for navigation, Wikipedia's hyperlinked articles foreshadowed an era of centralized web platforms that sustain user engagement with a consistent experience and "bottomless" content to consume and engage.

I began editing Wikipedia in 2005 on an autobiographical trajectory. My first contributions were to the articles about the environs of my hometown, and I moved on to editing articles related to my university. Like many editors, I worked to improve the quality of these articles because of both intrinsic motivations, such as boosting topics I was invested in, as well as extrinsic motivations, such as having the page become "Today's Featured Article" on Wikipedia's homepage. But I learned there are many ways to get to Wikipedia's front page. Immediately to the right of Today's Featured Article is the "In The News" (ITN) box featuring "articles that have been substantially updated to reflect recent or current events of wide interest" [2]. The presence of news-like content in an encyclopedia is uncanny. On one hand, encyclopedias are supposed to be stable references of historical knowledge rather than dynamic accounts of current events. On the other hand, there is a long history of encyclopedia editors grappling with how to incorporate new knowledge and encyclopedia publishers competing to be the most up-to-date [3]. Wikipedia's choice to privilege content related to current events via the ITN is also shrewd: it simultaneously is a shortcut to content users may already be searching for, it showcases the dynamism and quality of Wikipedia articles, and it invites users to consume and contribute to content outside of their primary interests.

## September 11 and Wikipedia

To understand how Wikipedia's "In The News" template and its broader culture of breaking news collaborations came about, we have to return to the immediate aftermath of the September 11, 2001

attacks. Wikipedia was 10 months old at the time of the attacks and while it already surpassed its elder sibling Nupedia in the number of articles, it was far from certain that the project would ever reach a sustainable level of activity. Although a comprehensive accounting of the editing activity in the immediate aftermath of the events has been lost to a server migration, snapshots from the Internet Archive's Wayback Machine along with listserv discussions document the extent to which the Wikipedia community at the time went into overdrive in response to the attack. Far from being an idiosyncratic case of online collaboration, the decisions made by editors at the time to use Wikipedia's unique collaborative capacities to deeply cover the September 11 attacks would fundamentally change the direction, scope, and culture of Wikipedia as a project to the present day.

A Wayback Machine snapshot of the "September 11, 2001 terrorist attack" article from October 9, 2011 captures the remarkable breadth and depth of topics that were authored and organized together about the attacks. There were timelines, documentation of closings and cancellations, lists of casualties, links to donating blood and money, articles on political and economic effects, and newly created articles about the buildings, cities, flights, and perpetrators as well as topics like "terrorism", "box-cutter knife", and "collective trauma." Approximately 100 September 11-related articles were created in total (when Wikipedia as a whole had only 13,000 articles) but Wikipedia's content attracted links from other prominent web gateways like Yahoo! that brought in an influx of desperately-needed new users to the project.

The list of casualties, enumerating each of the nearly 3,000 victims, sorted by name and location, and categorized by civilian or first responder), became a source of tension in the weeks following the attack. Some editors argued this level of detailed coverage was unbecoming of the traditional encyclopedias Wikipedia was trying to emulate stylistically and supporters referenced the rule that "Wikipedia is not paper" to justify a goal of writing biographies for thousands of victims, survivors, and leaders. As the trauma-induced altruism continued to fade, Wikipedia editors continued to raise concerns about the quality, notability, and importance of these memorialization efforts given the other demands of writing an encyclopedia. By September 2002 the community reached a consensus decision to move the September 11-related recollections and non-notable pages to a "Memorial Wiki." The launch of the memorial wiki led to heated discussions about which September 11-related articles would get to stay on Wikipedia and which would be relegated to the Memorial Wiki. The Memorial Wiki ultimately failed to thrive: its stagnant content and lack of editing activity led to accumulating vandalism and it was effectively shuttered by September 2006. The creation, rejection, and disappearance of the 9/11 Memorial Wiki's content remains an under-appreciated cautionary tale about the presumed durability of peer-produced knowledge: this content only persists when it remains integrated to the larger common project rather than being relegated to a smaller and more specialized project. Wikipedia is not immune from "rich get richer" mechanisms.

The Wikipedia community's over-reaction to the September 11 attacks and the discussions about the memorial content led to reflexive rule-making about news that persists today. The "What Wikipedia is Not" (WP:NOT) policy predates the attacks and enumerates that Wikipedia is not a dictionary, manual, directory, or a variety of other reference genres. In the midst of the debates in 2002 about what to do with the September 11 memorial content, the WP:NOT policy was expanded to assert that Wikipedia is not "a news report". The revised policy attempted to thread the needle between the channeling of collaborative energy following current events against diluting the mission of writing an encyclopedia. The policy emphasized that "Wikipedia should not offer news reports on breaking stories" but conceded "creating encyclopedia articles on topics currently in the news is an excellent idea" as long as current events articles are written in an encyclopedic style. This "NOT NEWS" policy has persisted to the present and the policy now emphasizes that "Wikipedia should not offer first-hand news reports on breaking stores" and "newsworthy events do not [automatically] qualify for inclusion... breaking news should not be emphasized or treated differently from other information." Another change in identity that emerged as a result of the September 11 memorial content was the addition of "Memorials" to the WP:NOT policy. The policy, revised in 2004, now emphasizes that "Wikipedia is not the place to memorialize deceased friends, relatives, acquaintances, or others who do not meet such requirements." These normative guardrails remain in place today to channel the outpouring of pro-social collaborative energy and sensemaking in the aftermath of traumatic events.

## WikiNews

Wikipedia's success in using peer production to rapidly generate high-quality encyclopedia articles invited efforts to extend the model to journalism itself. Unencumbered by the encyclopedia project's rules and policies, a peer-produced and wiki-mediated platform for journalism would unleash the community's demonstrated capacity for neutrality, timeliness, and accountability. Proposals for what would be Wikinews started in January 2003 and by December 2004, there were 20 language editions of Wikinews. Unlike Wikipedia, Wikinews encouraged reporting, interviews, and other "original research" while importing Wikipedia's neutral point of view as a bedrock value. The early signs were promising as the Indian Ocean earthquake and tsunami, Hurricane Katrina, and ongoing conflict in Iraq were complex and consequential events for producing and consuming popular content. Detractors criticized Wikinews proposals as a project fork that would deprive the parent Wikipedia project of the collaborative vigor that sustained its early success.

But Wikinews was not bereft of rules: the project made ultimately fatal decisions to adopt genres from journalism that undermined the ability for its editors to collaborate, collective attention to be focused, and network effects to be unlocked. Wikinews articles were subject to a formal editorial clearing process of marking articles as "In Development" until approved by an administrator for promotion to the front page. Rather than focusing popular and editorial attention around a central article, Wikinews encouraged editors to develop multiple parallel articles covering different angles. Once finalized by an

admin, the Wikinews article became immutable and new developments on the event required creating new articles. Lacking the will to draft a new article from scratch and being unable to revise existing articles, the pool of motivated contributors rapidly narrowed and featured article content remained stale. Finally, Wikipedia siphoned off the attention and production flowing to the articles surrounding an event from Wikinews because Wikinews referenced Wikipedia far more than Wikipedia referenced Wikinews. All of these forces conspired to strangle the growth of Wikinews, which peaked at 15 new articles per day in 2005 before falling below 5 new articles per day by 2011. The project limps along in 2019 but its coverage is biased by the interests of its remaining contributors and has only weak correspondence to traditional news values: a headline from mid-2019 reads “Wikinews attends Texas Haunters Convention.”

Wikipedia did not set out to innovate a new genre of journalism, but the failure of Wikinews to thrive illustrates the boundaries of documenting current events in an online peer production setting. The traditional news story, written as a sequential iteration updating but not replacing prior accounts of an event, is an artifact of the material limitations of writing on paper. For a major event like the 2004 Indian Ocean earthquake, Wikinews imported the norms of journalism and generated dozens of separate static articles, which ended up competing for scarce audience, contributors, and headline space. In contrast, Wikipedia offered a single, central article with a complete and dynamic account that could be linked from other article and privileged on prominent outlets like the In The News template on the homepage. Wikinews’s strong gatekeeping model requiring articles to be formally vetted before being “published” emulated the overzealous editorial model of Nupedia. In both cases, the content that ended up being published reflected the idiosyncrasies of the editors with the patience and skill to push their stories through the editorial pipeline rather than meeting the needs of a general audience. Finally, the Wikipedia model of a centralized account was much more legible to search engines like Google that boosted these articles’ authority and drove the virtuous feedback loops of more traffic, more contributors, more updates, and better content. In 2009, then Google vice president Marissa Mayer imagined a new web-oriented form of journalism where news stories did not compete against each other for authority or search engine results:

“how [might] the authoritativeness of news articles grow if an evolving story were published under a permanent, single URL as a living, changing, updating entity?” [\[4\]](#)

It is hard to imagine Ms. Mayer’s vision of the future of journalism was not influenced by the enormous volumes of traffic her search engine was referring to Wikipedia in the aftermath of current events.

## Features of breaking news collaborations

Even as an extremely active Wikipedia editor who made hundreds of revisions per month, I was always disappointed that I was never the first to create or update an article about a major current event. Wikipedia's editors had remarkable alacrity in revising content in response to current events: articles about deceased celebrities, political scandals, and natural disasters were all updated or created seemingly within minutes of the news breaking. My disappointment at being unable to author the first revisions shifted into curiosity, and I began to explore the revision histories of these breaking news articles. These explorations raised more questions about the emergent social behaviors and I switched my dissertation research project to exploring these breaking news collaborations. More than a decade later, Wikipedia's collaborations around breaking news continues to be a generative research context for myself and other researchers. When death, disaster, or conflict strikes in the world, I admit to some level of macabre excitement because there is now another case to add to my sample for research. Several general patterns have consistently emerged from my research over the past decade into Wikipedia's breaking news collaborations.

First, the contributors to breaking news articles are drawn from editors across the Wikipedia community, not introduced by a small set of "ambulance-chasing" editors who had specialized roles and routines of breaking news editing. This suggests the motivation and ability for editors to engage in breaking news collaborations is widely shared. This distributed collaborative capacity proved to be important throughout Wikipedia's history for mobilizing when multiple major events happened simultaneously. In March 2011 while the events of the Arab Spring demanded complex revisions across articles related to Tunisia, Egypt, Libya, and Syria, a 9.0-magnitude earthquake off the coast of Japan's Tōhoku province triggered a massive tsunami that ultimately killed more than 20,000 people and led to the most serious nuclear disaster since Chernobyl. Because of Wikipedians' distributed collaborative capacity, editors were able to process these major historical events in parallel when each event itself required a massive undertaking of synthesizing, coordinating, and deliberating across dozens of articles, talk pages, administrative processes, and language editions. Moreover, the contributors to breaking news article collaborations have diverse repertoires and roles on the project: an editor specializing in editing articles about Japanese boy bands shifted their focus to updating infrastructure damaged by the 2011 tsunami while another editor migrated their dispute resolution experience from Harry Potter articles to the Fukushima nuclear disaster article. Other topical areas that are proximate to breaking news have developed specialized routines for managing common coordination problems. When a new storm happens, members of the The WikiProject Tropical Cyclones shift to editing these articles and bring a wealth of experience for structure, style, references, and multimedia about storms to structure these collaborations. Pro-social responses in the aftermath of disaster and catastrophe are ubiquitous [\[5\]](#), but Wikipedia uniquely channels this energy into producing enduring and highly networked knowledge artifacts.



Second, no genre of breaking news events receives preferential collaborative treatment more than other events. Breaking news collaborations are regularly found around biographies of the recently deceased, natural disasters, accidents, conflicts and attacks, elections, controversies, sporting events, and popular culture. However, notability criteria still apply and are stringently enforced: while Wikipedia might have an article about every plane crash or hurricane, it does not have articles for every murder or political scandal. Systematic biases likewise creep into Wikipedia's coverage of breaking news events. For example, a minor earthquake like the 2011 Virginia earthquake that had no casualties has more contributors writing a longer article than a similar earthquake like the 2011 Myanmar earthquake that killed more than 150. Because breaking news collaborations emerge across topics, it has required Wikipedians to adopt policies and develop routines across the project. These include generalized warning templates that can be included to warn readers about the reliability of information on the rapidly-changing page and policies like "Wikipedia is not a newspaper" cautioning against employing original reporting or using journalistic writing styles. In extreme cases, editors use more complex affordances of MediaWiki: during the 2011 Tōhoku disasters, editors created a death toll template which can be updated in one place with the changes automatically propagating to all the other pages referencing the template rather than revising dozens of pages each time the death toll was revised.

Third, breaking news events are sites of large, rapid, and temporary collaborations that were otherwise rare on Wikipedia. The average Wikipedia article has accumulated fewer than ten unique editors and revisions over a span of years while breaking news articles can have hundreds of editors and revisions over a span of days. Examining the archival "zeitgeist" statistics for the English Wikipedia articles [6], the most actively-revised articles in any given month tend to be related to breaking news events or people in the news. In 2004, the articles with the most unique editors in a month included the "2004 Madrid train bombings" (112 editors in March), "Ronald Reagan" (114 editors in June), "2004 Summer Olympics" (92 editors in August), "2004 U.S. Presidential election timeline" (154 editors in November), and "2004 Indian Ocean earthquake and tsunami" (345 editors in December). The number and frequency of revisions to these articles was also extremely high: on major events, multiple revisions can be made in the same minute, complicating efforts for longer-form writing or copy-editing.

The MediaWiki software on which Wikipedia runs did not anticipate this kind of synchronous editing behavior, so editors revert to strategies for working around the limitations of the software such as making smaller and more frequent edits, merging in changes from a sandbox, or requesting an administrative lock on the article to incorporate requested changes. These collaborations are often temporary involving editors with disparate expertise and interests to come together to collaborate, most of them never having worked together before and with no expectations of collaborating again in the future. In the absence of social relationships to shape these emergent collaborations, editors are



guided by common interests and shared values around writing an encyclopedia. Even if most participants in breaking news collaborations return to editing their usual topics afterwards, these breaking news collaborations play a crucial role as "watering holes" where different groups' norms are reaffirmed and best practices are synthesized and then diffused back out through the rest of the project. Breaking news collaborations arguably play an important role in the viability of the broader Wikipedia project by engaging editors in challenging experiences, validating the investments of volunteer editors, and circulating innovations throughout the project.

Finally, breaking news articles are exceptionally high-quality compared to the median Wikipedia article: they tended to be longer, have more links to other Wikipedia articles, more references and citations, and more images, maps, and multimedia. Recent events make more "raw" material available in the form of reporting and social media content than historical events requiring archival research skills, which provides a richer set of inputs to generate better articles. But breaking news articles also benefit from "Linus's Law" [7] where a large number of diverse editors can accomplish tasks that would seem only possible for a small group of experts to accomplish. These articles also have a complex lifecycle of different cohorts of editors cycling through the collaboration over the course of days, weeks, and years. Biographical articles about the recently-deceased often go through a major rewrite to incorporate information from obituaries as well as a general reappraisal and standardization of structure and style rather than simply changing verb tenses and adding in the relevant information about the subject's death. Anniversaries have also become occasions for readers and editors to revisit an article and make new contributions. Wikipedia articles about current events provide a unique commons for emergent communities to gather to not only document and reappraise our understanding of the causes, contexts, and consequences of major and often traumatic events, but to support others' information-seeking and sensemaking as well.

All of these patterns reinforce the idiom that "Wikipedia works in practice, not in theory." Who are these editors that rapidly self-select and self-organize themselves in the absence of any formal coordination or delegation? Why have breaking news collaborations continued to employ generalists rather than developing a class of specialists? How did dozens of users synchronously edit a shared document using an asynchronous tool with none of features we take for granted in something like Google Docs? These remain open and vital questions for researchers 20 years after Wikipedia's launch.

## **Wikipedia in the age of disinformation**

Despite being the "encyclopedia that anyone can edit" and one of the most-trafficked websites in the world, Wikipedia did not show the same susceptibility to the coordinated disinformation campaigns that plagued social platforms like Facebook, YouTube, and Twitter around 2016. Although these platforms have made massive investments in human and automated moderation to improve users' experience, allay advertisers' concerns, and head-off regulator scrutiny, disinformation, harassment,

and other sociotechnical sludge remain endemic [8]. Provocateurs, outrage-mongers, and outright fascists have flocked to these virality engines' ability to actively recommend fringe ideas and compensate their creators while distributing them to audiences of millions. Platforms' attempts at common sense content moderation by removing or "demonetizing" the most egregious examples of hate speech and harassment have in turn led to accusations of their threats to "free speech" and "anti-conservative bias." What explains Wikipedia's apparent resilience to the sociotechnical sludge polluting other platforms?

The most obvious hypothesis is the difference in incentives between the user experience of advertising-driven engagement maximization and commons-based peer production models. Facebook, YouTube, and other popular social platforms generate billions of dollars in revenue by injecting personalized advertising alongside bottomless recommendations and newsfeeds managed by expensive engineers and infrastructures to engage users' attention. Every user's Facebook News Feed is personalized in response to their relationships, interests, and behavior. Content featuring novelty, humor, and outrage receives greater "engagement", so publishers and advertisers are locked in an arms race to produce ever more attention-grabbing content and target it for users' personalized feeds. Wikipedia has no newsfeed<sup>1</sup>, runs no advertising, and has a comparatively minuscule operating budget. But an overlooked and critical difference between Wikipedia and other social platforms is the absence of personalization in the user experience. Every English Wikipedia user's "Abraham Lincoln" article is the same regardless of their geography, gender, browsing history, or social graph. This common experience concentrates collective scrutiny and deliberative capacity rather than diffusing these accountability mechanisms across inscrutable and incommensurable personalized news feeds. Linus's Law—"given enough eyes, all bugs are shallow"—evidently holds for preserving the integrity of social information feeds.

A second hypothesis explaining Wikipedia's resilience to the sociotechnical sludge is its strong editorial identity. Facebook, Youtube, and other popular social platforms have resisted implementing any editorial standards and emphasized they are neutral platforms for sharing content. This disposition stems from the protections they are afforded under Section 230 of the (otherwise defunct) Communications Decency Act that offers platforms immunity from the liabilities associated with its users' speech: Mark Zuckerberg does not have to worry about libel claims if I use Facebook to accuse Senator Ted Cruz of being the Zodiac Killer. Despite the enormous investments in content moderation in recent years to stem criticisms from users, producers, and advertisers, these constituencies have no formal say in the development of content moderation policies. Moreover, the uneven application of moderation policies is a regular source of embarrassment as platforms ban obvious parody content while unvarnished supremacists continue to operate in the open. In contrast, the Wikipedia community has developed an extensive—and often sclerotic—body of rules and policies governing the content, style, and processes of writing encyclopedia articles. Foremost among these policies is a

muscular commitment to neutrality by avoiding editorial bias and verifying claims against reliable sources. This “active” neutrality policy by the Wikipedia community is a stark contrast to the “passive” neutrality of platforms that resist exercising any editorial judgment out of fear of jeopardizing their Section 230 protections. Wikipedia’s strong editorial identity is reinforced by its interlocking rules as well as what is effectively a growing body of case law with Arbitration Committee precedents that can be invoked to shut down bad-faith deliberation and outright lies.

A third hypothesis explaining Wikipedia’s resilience to sociotechnical sludge is the absence of algorithmic amplification. The background above illustrates how Wikipedia articles can “trend” in response to current events and popular culture. However, Wikipedia’s editors exercise considerable “human in the loop” editorial discretion over both the substance of trending content as well as its amplification mechanisms, unlike the algorithms driving newsfeed-centered platforms like Facebook and YouTube that can be manipulated into privileging viral and outrageous content. The most common user experience of Wikipedia is arriving from a search engine and navigating to related articles via hyperlinks or follow-on searches rather than navigating in from a news feed or homepage. To the extent Wikipedia has mechanisms for amplifying content to users, they exist on the homepage as “Today’s Featured Article”, “In The News”, “Did You Know”, and “On This Day.” These mechanisms are all explicitly vetted by human editors following documented public policies and consensus-driven deliberation that still have remarkable alacrity in responding to current events. The responsiveness of Wikipedia editors to current events also provides an important counter-factual to claims from engineering culture that human-in-the-loop systems lack the scalability, speed, and accuracy of automated systems, despite accumulating evidence of automated systems’ multiple liabilities. Because the oversight and capacity to intervene in Wikipedia’s attention amplification mechanisms is delegated across hundreds of administrators and/or thousands of editors, they are substantially harder to compromise than algorithmic systems operating under “security through obscurity” strategies.

Social platforms confronting the limitations of their current engagement and moderation models are turning to Wikipedia. In October 2017, Facebook announced that it would provide “contextual information” about articles in the news feed that would include links to Wikipedia [9]. In March 2018, YouTube CEO Susan Wojcicki outlined a strategy wherein YouTube would connect videos containing conspiracies to corresponding Wikipedia articles in an effort to combat the spread of disinformation [10]. YouTube’s decision, in particular, came as a surprise to the Wikipedia community and the Wikimedia Foundation, who were given no forewarning that they would be indirectly policing YouTube’s toxic content. The fundamental risk was that the same dynamics that converted information-seeking Google search users into Wikipedia editors could also convert the conspiracists, ideologues, and culture warriors on these platforms into Wikipedia editors. These decisions to outsource content moderation to Wikipedia were deeply irresponsible: either YouTube failed to

comprehend the obvious risks of swamping the smaller volunteer project with their content moderation problems or they did not care.

YouTube's conduct in this case is a classic problem of governing what economists call "common goods" and its corresponding "tragedy of the commons." The knowledge produced—and more importantly, governed—by Wikipedia is "non-excludable" which means that it can still be used by people who have not contributed to it. However, the governance of this knowledge exhibits patterns of "rivalrousness" in which consumption by one actor reduces availability for others. In this case, YouTube contributed nothing to Wikipedia's governance but could still benefit from the credible content Wikipedia generated (non-excludability). But in outsourcing content moderation to Wikipedia editors and administrators, it was potentially reducing their capacity to attend to other content generation and moderation demands (rivalrousness). YouTube was effectively "over-fishing" the capacity of Wikipedia editors and administrators to handle sociotechnical sludge by requiring the volunteer Wikipedia community to do more of all of this work on behalf of an corporation who profits from not having to moderate its own content. But commons do not inevitably end up as tragedies; the research of Elinor Ostrom (which culminated in her 2009 Nobel Memorial Prize in Economic Sciences) charts out strategies for designing institutions for sustaining commons in the face of threats like over-use. Her 2006 edited volume with Charlotte Hess, *Knowledge as a Commons*, charts out prescient strategies for communities like Wikipedia to pursue to "define, protect, and build the knowledge commons in the digital age."

The case of Wikipedia content being re-deployed by unscrupulous platforms for their content moderation needs illustrates the risks associated with the "interoperability" of online platforms: content from Platform A can be plugged into Platform B but these connections can also cause blow-back as the bad behavior from Platform B moves to Platform A. Wikipedia's content is re-used in both visible and invisible ways across platforms: Google serves up Wikipedia content alongside its search results, Facebook uses it to populate information for its pages, and Apple's Siri or Amazon's Alexa will read summaries of articles. Wikipedia's content is also used in more invisible ways to train algorithms used for translation, image recognition, and concept similarity. These interoperable connections increase the prominence of Wikipedia's content, recruit new users to contribute, and highlight the need to preserve this commons, but every new interoperable link also introduces new threats. If a malicious actor wanted to undermine trust in these other major platforms, an under-realized vector is subtly compromising the quality of information from the Wikipedia and Wikidata content that they ingest.

Wikipedia's resilience to the strategic disinformation campaigns from 2016 should not be interpreted as intrinsic immunity to information manipulation: Wikipedia's most active editors are not representative of the population at large, which creates both biases in its content and blindspots in its

responses, which are then ingested and amplified through the web of interoperable dependencies outlined above. Wikipedia administrators botched its response to the GamerGate controversy in 2015 by acquiescing to a manipulative influence campaign and banning five editors who had been fending off extremist content [11]: this case illustrated how Wikipedia’s administrative procedures can be hijacked by bad-faith actors to target good-faith editors. On a lighter note, another illustration of the threats of interoperability is a case from October 2017. When users of Apple Siri asked “What is the national anthem of Bulgaria?”, they were served “Despacito”, a 2017 reggaeton pop hit, rather than the nineteenth century hymn “Mila Rodino” [12]. Somewhere deep in Apple’s knowledge graph, much of which is likely trained on Wikipedia and Wikidata, this erroneous pairing was introduced and never validated before being pushed out to millions of users.

Wikipedia and its increasingly important sister project Wikidata have been able to resist disinformation efforts because of the ability to match its supply of human-in-the-loop governance with demand for information: oversight follows the action. While it might be hard to embed disinformation into articles about candidates for an upcoming election because of this super-abundance of editorial attention, it might be trivial to persistently embed disinformation into provincial articles about distant historical events, specialized scientific topics, or marginal trivia about national anthems that lack sustained editorial oversight. While Wikipedia’s unique editorial model has shown greater resistance to the disinformation, harassment, and manipulation plaguing other social platforms to the point that its content is serving as a front-line defense, there are nevertheless growing precedents that Wikipedia’s content and governance has very real vulnerabilities that could easily and quickly propagate throughout a complex technical stack of interoperable technologies.

## Conclusion

Encyclopedists have always struggled with the limitations of synthesizing knowledge into paper documents because when the knowledge changes, so must the paper. Wikipedia was not the first encyclopedia to use the online medium to rapidly and inexpensively revise content in response to changes, but its unique “anyone can edit” model had the effect of entangling current events with the viability of the project.

The September 11 attacks were a critical moment in Wikipedia’s history. The events brought in an influx of new editors motivated to document the events, perpetrators, victims, and contexts and the outpouring of collaborative effort in the aftermath of the 9/11 attacks validated an under-appreciated strategy for growing the project. By simultaneously tapping into editors’ pro-social motivations following traumatic events as well as showcasing the quality and timeliness of the project’s content in a time of acute information seeking and sensemaking, Wikipedia could convert the large numbers of information seeking users into new contributors as well as increase popular trust in its radical editorial model. However, Wikipedia editors’ overzealous creation of 9/11-related content also

required the development of new rules and identities as guardrails that persist today about what the encyclopedia is and is not.

Wikipedia editors continue to invest enormous amounts of effort in covering breaking news and current events within the confines of these guardrails. Articles about the recently deceased, natural disasters, conflicts, and popular culture are sites of large and extremely dynamic collaborations involving dozens of editors making hundreds of revisions within hours. While Wikipedia's MediaWiki software was not designed with this use-case in mind, these high-tempo collaborations continue to serve crucial roles in sustaining the health of the broader project, close to 20 years after the early precedent of the 9/11 attacks: they bring in new users to the project, provide opportunities to disparate sub-communities to temporarily congregate, disseminate innovations and best practices into the rest of the community, and produce high-quality content hyperlinked to other relevant background.

Wikipedia remains a product of a particular historical moment from the early 2000s, not only in terms of its adorably dated interface, but also the absence of advertising and engagement, news feeds and recommendation systems, and virality and polarization as central features that define so much of the user experience on other social platforms. Wikipedia's resilience to the disinformation that plagued Facebook, YouTube, and Google in 2016 would suggest this archaic user experience provided an important defense against actors who weaponized these attention amplification mechanisms on other platforms to malicious ends. But this story overlooks other explanations for Wikipedia's apparent resilience. Wikipedia's strong editorial identity emphasizing neutrality and verifiability provides a critical bulwark against the tide of sociotechnical sludge polluting these other platforms who only reluctantly moderate offensive and false content. But most critically, Wikipedia users and editors' attention is shared around common articles rather than distributed across personalized newsfeeds.

Does Wikipedia's success in covering breaking news and current events chart a path for other platforms to follow? Information seeking and sensemaking about current events drive enormous flows of online collective attention, which explains why "News feeds" and "Trending" topics are ubiquitous on social platforms. Whether and how Wikipedia can channel this demand for information likewise has been central to its ongoing identity, relevance, and sustainability. Wikipedia remains a valuable counterfactual for the potential of designing around information commons, human-in-the-loop decision-making, and strong editorial stances in the face of the Silicon Valley consensus emphasizing content personalization, automated moderation, and editorial indifference. The differences in how Wikipedia handles current event information may have insulated it from manipulation, but as platforms increasingly turn to Wikipedia for providing and moderating content, Wikipedia's very real vulnerabilities risk becoming a target.



## Footnotes

1. Active editors might quibble with this distinction since Wikipedia's Watchlist summarizes recent changes to articles they follow. There is also a list of Current Events articles. But neither of these are central to the median Wikipedia user's experience. [↵](#)

## Citations

1. Grossman, Lev. "You—Yes, You—Are TIME'S Person of the Year," *TIME*, December 25, 2006, <http://content.time.com/time/magazine/article/0,9171,1570810,00.html> [↵](#)
2. Wikipedia:In the News. (2019). Wikipedia.Wikipedia. "Wikipedia: In the news." Last modified, August 30, 2019. [https://en.wikipedia.org/wiki/Wikipedia:In\\_the\\_news](https://en.wikipedia.org/wiki/Wikipedia:In_the_news) [↵](#)
3. Loveland, J., & Reagle, J. (2013). Wikipedia and encyclopedic production. *New Media & Society*, 15(8), 1294–1311. <https://doi.org/10.1177/1461444812470428> [↵](#)
4. Seward, Zachary M. "Google News experimenting with links to Wikipedia on its homepage." Nieman Lab, June 9, 2009. <https://www.niemanlab.org/2009/06/google-news-experimenting-with-links-to-wikipedia-on-its-homepage/> [↵](#)
5. Solnit, Rebecca. *A Paradise Built in Hell: The Extraordinary Communities that Arise in Disaster*. New York: Penguin, 2010. [↵](#)
6.  
"Wikipedia Statistics English," Wikimedia Foundation, last modified August 31, 2019, <https://stats.wikimedia.org/EN/TablesWikipediaEN.htm> [↵](#)
7. Raymond, Eric S. *The Cathedral and the Bazaar*. Sebastopol, CA: O'Reilly Press, 1999. [↵](#)
8. Benkler, Yochai, Robert Faris, and Hal Roberts. *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford: Oxford University Press, 2018. [↵](#)
9.  
Constine, Josh. "Facebook tries fighting fake news with publisher info button on links." *TechCrunch*, October 5, 2017. <https://techcrunch.com/2017/10/05/facebook-article-information-button/> [↵](#)
- 10.



Matsakis, Louise. "YouTube Will Link Directly to Wikipedia to Fight Conspiracy Theories." *WIRED*, March 13, 2018.

<https://www.wired.com/story/youtube-will-link-directly-to-wikipedia-to-fight-conspiracies/> ↵

11. Williams, Lauren C. "The 'Five Horsemen' Of Wikipedia Paid The Price For Getting Between Trolls And Their Victims." *Think Progress*, March 6, 2015. <https://thinkprogress.org/the-five-horsemen-of-wikipedia-paid-the-price-for-getting-between-trolls-and-their-victims-9c835aeafdc8/> ↵

12.

Statt, Nick. "Why does Siri think the national anthem of Bulgaria is 'Despacito'?" *The Verge*, October 4, 2017.

<https://www.theverge.com/2017/10/4/16427146/apple-iphone-siri-bulgaria-national-anthem-despacito-incorrect> ↵