How AI Reshapes Human Content Creation: The Case of Grokipedia vs Wikipedia*

Tin Cheuk Leung† Wake Forest University

Koleman Strumpf‡ Wake Forest University

Preliminary, Do Not Cite
December 2025

Abstract

Generative AI models rely on large volumes of user-generated content (UGC), yet we know little about how AI tools affect the production of that underlying content. We study this question using the introduction of Grokipedia, an AI-generated online encyclopedia operated by xAI, which provides automated summaries that could either substitute for human editing or expand visibility and draw in new contributors. Using a new panel dataset covering 1.4 million Wikipedia pages of notable individuals—linking edit histories, traffic data, article characteristics, and Grokipedia coverage—we estimate the causal effect of AI summaries on subsequent human contributions by constructing matched samples of treated and untreated pages within occupational fields. We find a consistent and surprising result: *the availability of Grokipedia increases human editing activity.* Page views also rise, suggesting that AI summaries act as an attention amplifier that ultimately stimulates UGC. We assess how Grokipedia interacts with existing content by constructing a semantic similarity measure between each Grokipedia summary and its corresponding Wikipedia page. Similarity scores are concentrated near one—about 23% of summaries are nearly identical—but exhibit substantial variation. We use this variation to test whether more novel AI summaries elicit greater human engagement, and find that pages with *lower* similarity experience *larger* increases in editing after Grokipedia's launch. Our findings show that embedded AI tools can complement—rather than crowd out—voluntary human contributions, highlighting a positive feedback loop between AI features and the data ecosystem that sustains modern AI models.

JEL Classification: D83, L82, L86

Keywords: Wikipedia, Digital Platforms, Artificial Intelligence, User-Generated Content

*We thank `@lthlnkso` for assistance with Grokipedia data. The usual disclaimer applies.

†Contact: Department of Economics, Kirby Hall, Wake Forest University, Box 7505, Winston-Salem, NC 27109. Email: leungtc@wfu.edu

‡Email: strumpks@wfu.edu

1 Introduction

Artificial intelligence (AI) systems rely fundamentally on vast amounts of user-generated content (UGC), yet we know surprisingly little about how the rapid rise of AI tools feeds back into the production of the very content on which they depend. While existing research has focused on how AI affects labor markets, productivity, misinformation, and online behavior, much less attention has been paid to how AI tools influence human contributions to large-scale, volunteer-driven knowledge platforms. This gap is consequential: if AI-generated content discourages or displaces human contributors, the long-run sustainability of the data ecosystem that trains modern AI models may be undermined.[1]

We address this question by studying the launch of Grokipedia, an AI-generated encyclopedia introduced by Elon Musk in late 2025 as a direct competitor and potential substitute for Wikipedia. Grokipedia produces concise AI-written summaries for notable individuals and topics using proprietary and publicly available data, offering users an alternative source of information outside the Wikipedia ecosystem.

Theoretically, the effect of Grokipedia on Wikipedia's UGC is ambiguous. AI may act as a substitute: if users rely on Grokipedia summaries instead of Wikipedia, the reduced traffic and diminished perceived value of contributing may depress human editing activity. But AI may also act as a complement: users may draw information from Grokipedia to improve or update Wikipedia pages, or the publicity surrounding a new AI platform may direct attention toward existing Wikipedia entries, leading to more edits. The competition for viewers may also elicit greater effort from Wikipedia contributors. Which force dominates is ultimately an empirical question.

An anecdote from Grokipedia's launch illustrates this ambiguity—and hints at a complementary channel. On October 27, 2025—the day Grokipedia went live—the initial Grokipedia summary of Donald Trump notably included the sentence: "Family dynamics, including Fred Sr.'s pragmatic engagement with local political machines for development approvals, further embedded a patriotic self-interest wary of overreaching authority, though direct *anti-communist* transmission remains less documented beyond the era's

[1]The senior director of product at the Wikimedia Foundation argues that generative AI summaries on search engines like Google have led to an 8% reduction in year-over-year Wikipedia page views in Fall 2025 (Miller, 2025).

conservative milieu." Although this line was later revised, human editors on Wikipedia reacted almost immediately. On October 30, a Wikipedia editor added Trump's page to the category "American Anti-Communists," referencing content analogous to what appeared in the Grokipedia summary.[2] This episode provides anecdotal evidence that AI-generated summaries on an external platform may inform or prompt human editing behavior on Wikipedia.

To study this, we assemble a new panel dataset of 1.4 million Wikipedia pages of notable people across five various occupations (Academia, Culture, Leaders, Politics, and Sports) for which Grokipedia summaries became available on October 27, 2025. Our dataset combines (i) detailed page-level edit histories from Wikipedia; (ii) traffic data capturing daily and weekly readership patterns; (iii) rich article-level features such as length, age, topical categories, and historical attention; and (iv) metadata linking pages to Grokipedia's coverage. Our empirical design follows the matching framework developed in Section 5. Because Grokipedia disproportionately selects highly viewed pages, and visibility varies systematically across occupations, we construct a matched sample using Mahalanobis-distance nearest-neighbor matching within each field. Treated pages (those with a Grokipedia entry) are paired with the closest untreated page based on pre-treatment views, edits, long-run readership, and page characteristics, approximating the counterfactual trajectory each treated page would have followed absent Grokipedia. We then estimate average treatment effects on the treated (ATT) by comparing post-launch changes in views and edits between each treated page and its matched control.

Across both approaches, our results reveal a surprising and robust pattern: the availability of Grokipedia summaries leads to an increase, rather than a decrease, in human-generated edits. Pages that gain an AI summary experience significantly more editing activity—both in total edits and in the probability of receiving any edit—relative to matched control pages. We also document a substantial rise in readership following Grokipedia's introduction, suggesting that AI-generated summaries may function as an entry point that expands audience reach and ultimately spurs human engagement.

To understand how Grokipedia interacts with existing content, we construct a seman-

[2]Revision link: `https://en.wikipedia.org/w/index.php?title=Donald_Trump&diff=prev&oldid=1319468163`

tic similarity measure between Grokipedia summaries and their corresponding Wikipedia pages. The distribution is skewed toward high overlap—about 23% of summaries are nearly identical to Wikipedia text—yet exhibits meaningful variation, with most scores between 0.5 and 1 and a mode around 0.85. We then use this variation to test a mechanism suggested by early anecdotal evidence: AI-generated summaries may introduce new frames or details that prompt editors to revisit and update articles. Consistent with this hypothesis, pages with *lower* similarity between Grokipedia and Wikipedia experience *larger* increases in human editing after Grokipedia's launch. This relationship remains negative and statistically significant across specifications controlling for field categories and page-level attributes.

Taken together, our findings offer novel evidence that the integration of platform-embedded AI tools can complement—rather than substitute for—voluntary human contributions in online knowledge production. These results contribute to a nascent but crucial literature on generative AI's impact on UGC and highlight an important feedback loop: AI tools built on crowd-sourced information can, under certain conditions, stimulate further contributions, thereby reinforcing the data infrastructure on which AI itself depends. An important facet of this environment is the private and uncompensated provision of a public good, and generative AI may enhance some pro-social explanations for such provision.

The remainder of the paper is as follows. Section 2 discuss how our work connects to and advances several active research areas, and Section 3 provides the background of Wikipedia and the launch of Grokipedia. We describe the data in Section 4. We then discuss our empirical strategy and results in Section 5. The last section concludes.

2 Literature

Our study speaks directly to a new but quickly expanding body of work examining how generative AI affects user content creation. A central question in this literature is whether AI-produced summaries function as substitutes for human contributions, thereby weakening users' incentives to generate original material and ultimately reducing the volume of user-generated content (UGC). Early empirical evidence from online knowledge com-

munities lends support to this concern: the introduction of tools such as ChatGPT has been associated with noticeable declines in human-authored posts, edits, and other forms of participation (Lyu et al., 2025; Burtch, Lee, and Chen, 2024; Glickman and Zhang, 2024; del Rio-Chanona, Laurentsyeva, and Wachs, 2023).

In contrast to these findings, our preliminary results indicate that the emergence of generative AI—in our setting, Grokipedia—has a positive effect on subsequent human editing activity on Wikipedia.[3] Although we are still working to unpack the mechanisms underlying this unexpected pattern, the evidence suggests that AI tools may, under certain conditions, complement rather than crowd out human contributions. By documenting this heterogeneity and providing one of the first empirical cases where AI appears to stimulate rather than suppress UGC, our paper adds nuance to the current debate and contributes a novel perspective to this rapidly evolving literature.

Our results contribute to the broader literature examining whether generative AI is a complement or substitute for human skills (Brynjolfsson, Li, and Raymond, 2025; Acemoglu et al., 2022). Our work differs from most applications in two ways. First, the output we focus on (Wikipedia entries) is given away freely rather than sold, and contributors are not paid. As discussed below, the incentives for such public good contributions is likely shaped by pro-social motives potentially including altruism. It is not clear estimates from the literature which cpnsider paid employment generalize when there is no compensation, a common case on the internet (for example, many social network posts). Second, we look at knowledge creation rather than employment outcomes like productivity or wages.

Our paper also adds to work identifying the factors influencing the human creation of digital knowledge-bases.[4] Since contributors to such projects often receive no direct benefit and provide free labor, this is related to why individuals make private contributions to public goods. Zhang and Zhu (2011) shows that Wikipedia activity from registered contributors falls when there is a reduction in the potential readership, consistent with

[3]There is a nascent literature examining the connection between Grokipedia and Wikipedia, but it has a different focus than our work. These papers measure the static textual differences and citation practices between Grokipedia and the associated Wikipedia entry (Triedman and Mantzarlis, 2025; Yasseri, 2025). Our work in contrast causally estimates the dynamic feedback effects between the two sources.

[4]Some Wikipedia entries appear to be AI/LLM-generated (Brooks, Eggert, and Peskoff, 2024) despite Wikipedia policies prohibiting this practice. These are a small number of primarily recently created articles so it is not yet possible to measure their impact.

contributors getting some benefit from other's consumption (Zhu, Walker, and Muchnik 2020 find that exogenous increases in Wikipedia content lead to more edits and editors). Gallus (2017) finds that even a symbolic rewards lead to greater Wikipedia editor retention and activity. And Ren, Chen, and Riedl (2016) show that Wikipedia projects in which contributors have greater diversity of interests in the various topics within the project have higher productivity and retention rates. A common them in these works is that a pro-social motivation underlies the contribution. This suggests a more muted response among Wikipedia contributors to Grokipedia's entry.

3 Background on Wikipedia and Grokipedia

The institutional backdrop for Wikipedia's interaction with generative AI is one of cautious permissiveness. Official guidelines state that the platform does not prohibit the use of AI tools for drafting or editing, but emphasizes that editors remain fully responsible for ensuring verifiability, factual accuracy, reliable sourcing, and adherence to neutrality norms (Wikipedia contributors, 2024). AI-generated text is therefore acceptable only to the extent that it undergoes substantive human review. This reflects a broader philosophy articulated in the Wikimedia Foundation's AI strategy, which frames generative models as tools meant to augment, rather than replace, volunteer contributors (Owen, 2025).

In practice, however, the community has grown increasingly wary of large volumes of low-quality, machine-generated material. Reports highlight common problems—hallucinated citations, generic "AI-sounding" prose, unverifiable claims—which have led editors to develop stronger safeguards against unreviewed AI content (Heater, 2025). These concerns culminated in the 2025 introduction of a speedy deletion mechanism for articles suspected of being unvetted AI output ("AI slop") (Hale, 2025). As a result, although Wikipedia formally allows AI-assisted editing, it effectively imposes high human-verification costs: editors must manually check every claim, repair citations, and often rewrite entire sections to meet community standards.

This policy environment generates a particularly interesting context for studying the introduction of Grokipedia, an AI-native encyclopedia positioned as a competitor to Wikipedia, built in conjunction with xAI's "Grok" model. Launched in Oct 27th, 2025,

its stated purpose is to provide a faster-moving, AI-supported alternative to Wikipedia. According to public descriptions, Grokipedia was initially populated by automatically ingesting large parts of English-language Wikipedia, followed by additional AI-generated editorial curation. The project aims to maintain an up-to-date, searchable knowledge base closely integrated with the Grok chatbot, and its early growth has been rapid: by November 21 2025, Grokipedia already reported more than 1 million pages, although coverage remains uneven across topics.

In theory, a highly capable external AI system could substitute for human contributions, reducing editors' incentives to write or update articles on Wikipedia. Yet, given Wikipedia's insistence on human oversight, Grokipedia may instead function as a complementary upstream information source, lowering the cognitive cost of generating ideas or preliminary drafts that editors subsequently refine and integrate. The tension between Wikipedia's cautious approach to AI-generated text and the sudden availability of a competing AI platform thus creates fertile ground for unexpected positive spillovers: instead of discouraging participation, Grokipedia may stimulate human editing activity by serving as a catalyst for new edits, fact-checking, and article expansion.

4 Data

Our empirical analysis begins with the comprehensive list of notable individuals compiled by Laouenan et al. (2022). Their cross-verified database contains 2.29 million individuals, representing roughly one out of every 43,000 human beings who have ever lived, and includes a substantial share—approximately one-third—who do not have an English Wikipedia page. From this universe, we restrict attention to individuals who do have an English-language Wikipedia entry. We further narrow the sample to several primary occupational domains that are most relevant for our study: Politics, Sports, Leaders, Academia, and Culture, yielding approximately 1.4 million notable individuals with an English Wikipedia page in these five fields.

For each person in this filtered group, we scrape information from the official Wikipedia API covering both the period before the launch of Grokipedia on October 27, 2025 and the period afterward, with all data collection concluding on November 16, 2025. This scraping

includes the complete history of daily Wikipedia page views, the daily number of edits, and the full article text as it existed immediately before Grokipedia's introduction. We also collect the post-launch Wikipedia text so that any adjustments in content, length, or structure can be tracked in the weeks following the appearance of Grokipedia.

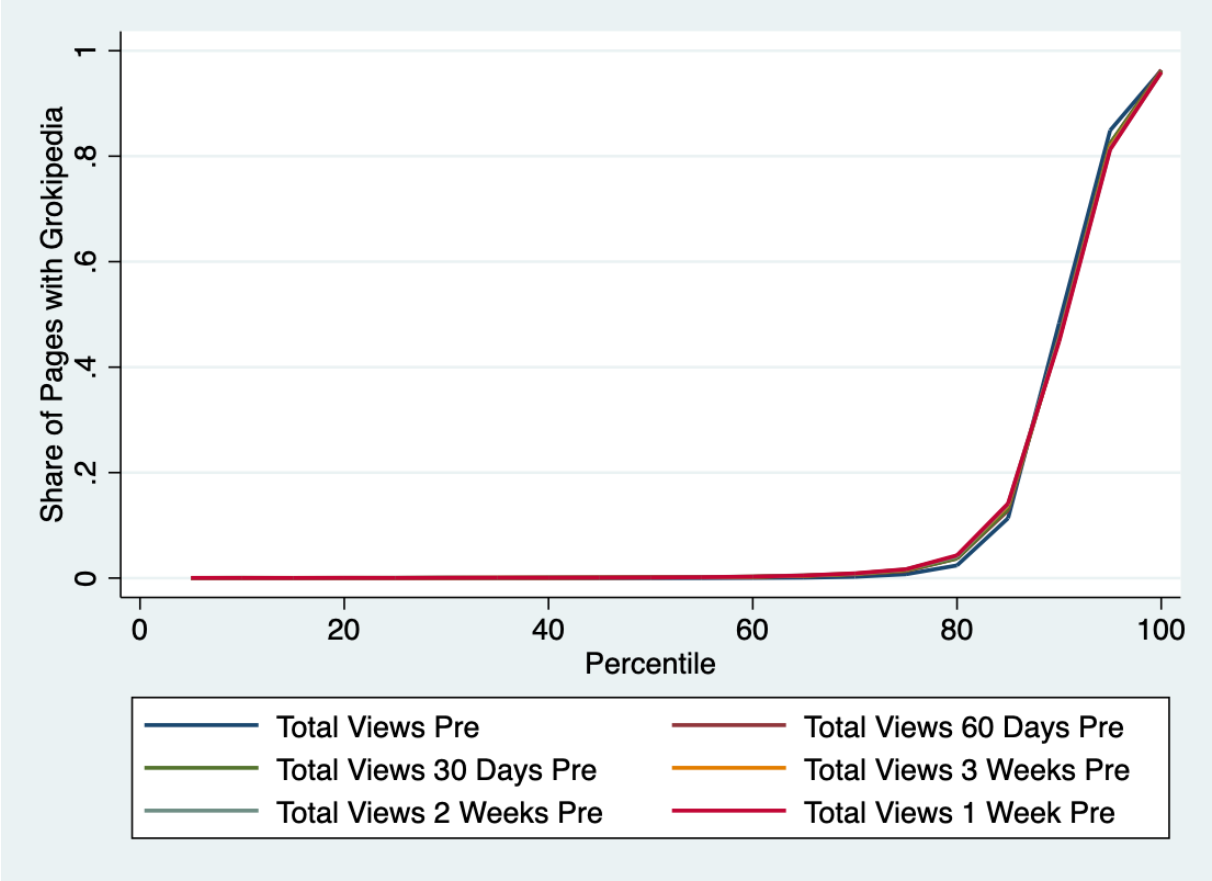
In parallel, for every individual in the notable-people list, we check whether a corresponding Grokipedia page exists. When one is present, we scrape the article text as well as any updates that have appeared up to our scraping cutoff date of November 16. It is important to note that Grokipedia itself continued to evolve beyond our data-collection window: the platform reported roughly 890,000 articles at the time our scraping ended, and only a few days later—on November 21, 2025—it announced an updated count of more than 1,016,241 articles. This indicates that Grokipedia was still growing rapidly during the period covered by our analysis. Combining all elements of the scraped data yields a unified person-level dataset that includes biographical identifiers, detailed Wikipedia traffic and edit histories, pre- and post-treatment Wikipedia content, and Grokipedia presence and text where available.

4.1 Who Were Selected to Have a Grokipedia Page?

Although the precise criteria Grokipedia uses to decide which Wikipedia pages receive a Grokipedia entry are not publicly documented, it is not unreasonable to assume that the platform prioritizes pages that already attract substantial attention on Wikipedia. This motivates our approach of examining how Grokipedia coverage correlates with underlying Wikipedia traffic and why we structure our data collection around measures of page popularity. To examine how the presence of a Grokipedia entry correlates with the popularity of a corresponding Wikipedia page, we compute for each of several pre-treatment view metrics (e.g., total views 30 days prior, 60 days prior, 1-3 weeks prior, and cumulative "wiki readers 2015–2018") the mean value of the Grokipedia indicator by 5-percentile bins of the view distribution. I then plot the mean indicator (y-axis) against percentile bins (x-axis from 5 to 100), resulting in Figure 1.

The figure displays a sharp and highly nonlinear pattern of coverage. Pages whose Wikipedia view counts lie below approximately the 80th percentile have virtually zero

Figure 1: Proportion of Grokipedia by 5-Percentile Bins of Pre-Treatment View Metrics



probability of being included in Grokipedia. In contrast, once the percentile exceeds around the 90th, the likelihood of having a Grokipedia page rises dramatically — climbing above 90 percent for the highest-view pages. This strongly suggests that Grokipedia's entry selection is heavily skewed toward high-traffic Wikipedia pages. For our propensity score matching design, the implication is that most of the treated (Grokipedia-covered) pages lie in a narrow band of high-view percentiles — roughly between the 85th and 95th percentiles — meaning the estimable variation in treatment is concentrated in that upper tail of the traffic distribution.

To better understand the factors associated with whether a Wikipedia page receives a corresponding Grokipedia entry, we estimate a series of linear probability models where the dependent variable is an indicator equal to one if the individual has a Grokipedia page. The key explanatory variable is the logarithm of Wikipedia page views in the 60 days preceding Grokipedia's launch, ln(Views 60 Pre), which captures the underlying vis-

ibility of the page immediately before treatment. We progressively enrich the specification by adding field fixed effects and their interactions with page views, and finally by controlling for observable Wikipedia-page characteristics such as article length, the individual's gender, and age.

Table 1 summarizes the results. Column (1) presents the simplest specification, where Grokipedia presence is regressed solely on ln(Views 60 Pre). The coefficient is large, positive, and highly statistically significant: a one-percent increase in pre-launch page views is associated with roughly a 0.15-percentage-point increase in the probability of having a Grokipedia page. Because this is a linear probability model, this implies that a ten-percent increase in views is associated with a 1.5-percentage-point increase in Grokipedia likelihood. More strikingly, ln(Views 60 Pre) alone explains almost half of the variation in Grokipedia coverage, with an R^2 of 0.477, highlighting how strongly page popularity predicts whether a page is included in Grokipedia.

Column (2) adds field indicators (Culture, Leaders, Politics, and Sports, with Academia as the omitted category). The coefficients show clear level differences across fields: even after controlling for pre-launch page views, pages belonging to Culture and Sports are each roughly six percentage points more likely to have a Grokipedia entry than pages in Academia, while pages in Leaders and Politics are about three to four percentage points more likely. The contribution of fields to explanatory power, however, is relatively minor. The R^2 rises only slightly, from 0.477 in column (1) to 0.482 in column (2), indicating that most of the variation in Grokipedia coverage is still accounted for by differences in pre-launch page views rather than occupational category.

Column (3) adds additional controls reflecting observable Wikipedia-page characteristics, including article length, gender, and the age of the biographical subject. The results show that longer Wikipedia pages are more likely to receive a Grokipedia entry, consistent with the idea that more developed pages attract greater attention. Gender differences are small but present: female individuals are roughly half a percentage point more likely than males to appear on Grokipedia after controlling for views and field categories. Age also exhibits a subtle pattern, with older individuals being slightly less likely to be selected. These controls contribute only modestly to explanatory power, raising

Table 1: OLS Regression on Grokipedia Presence

	Dependent variable: Grok indicator		
	(1)	(2)	(3)
ln(Views 60 Pre)	0.151***	0.151***	0.148***
	(0.000133)	(0.000136)	(0.000169)
Field: Culture		0.0662***	0.0664***
		(0.000744)	(0.000800)
Field: Leaders		0.0395***	0.0454***
		(0.000823)	(0.000896)
Field: Politics		0.0339***	0.0362***
		(0.000823)	(0.000881)
Field: Sports		0.0690***	0.0652***
		(0.000710)	(0.000764)
Wikipedia Word Count (10,000)			0.0321***
			(0.000324)
Gender: Male			-0.00648***
			(0.000570)
Gender: Other			0.0918***
			(0.0118)
Age			-0.0000370***
			(0.00000142)
Constant	-0.791***	-0.843***	-0.829***
	(0.000826)	(0.00103)	(0.00129)
Observations	1422093	1422093	1291333
R^2	0.477	0.482	0.502

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the R^2 from 0.482 in column (2) to 0.502 in column (3). Importantly, the coefficient on ln(Views 60 Pre) remains stable in magnitude and significance across specifications, reinforcing that pre-launch page views remain the dominant predictor of Grokipedia coverage even after accounting for field differences and page-level characteristics.

Taken together, the results show that Grokipedia's early coverage decisions were overwhelmingly concentrated among individuals whose Wikipedia pages were already popular. Field variation and page characteristics matter at the margin, but the dominant factor remains the visibility of the underlying Wikipedia page.

4.2 Similarity Between Grokipedia and Wikipedia Text

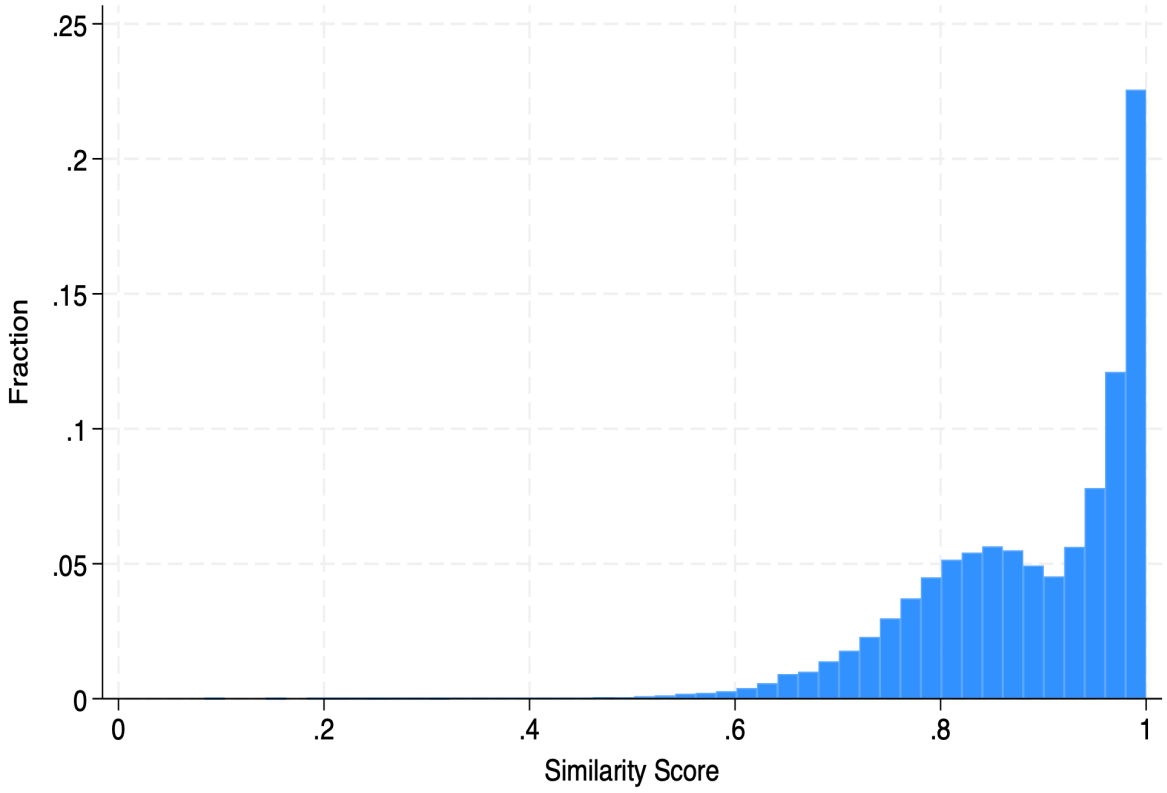
To assess how closely Grokipedia entries resemble their corresponding Wikipedia pages, we compute a semantic similarity score for all treated pages for which we currently have both Grokipedia text and the relevant historical version of the corresponding Wikipedia article.[5] At this stage, our dataset includes 46,467 matched pages, a subsample of the full treated sample of 138,989 pages. The remaining treated pages are those for which either the Grokipedia text or the pre–October 27, 2025 Wikipedia revision is not yet available; we are expanding our text corpus and expect to cover the entire treated sample in the next draft.

To compute similarity, we embed each Grokipedia article and its corresponding Wikipedia article using a transformer-based sentence embedding model (`all-MiniLM-L6-v2`), which maps each document into a high-dimensional semantic vector. We then calculate the cosine similarity between the two vectors, yielding a score between 0 and 1. A score of 1 indicates that the two texts are, from the model's perspective, essentially identical, while lower values reflect increasing divergence in wording, content, or emphasis. Because the embeddings capture semantic rather than purely lexical overlap, the measure is sensitive to paraphrasing and reorganization, not only verbatim copying.

Figure 2 plots the distribution of similarity scores across the matched sample. Several patterns stand out. First, a striking 23% of pages have a similarity score extremely close to 1, indicating that the Grokipedia text is essentially identical to the Wikipedia version. These cases reflect entries in which Grokipedia appears to replicate Wikipedia directly with minimal or no modification. Second, for the remaining observations, we observe a smooth and wide distribution of similarity values between roughly 0.5 and 1. The distribution has a clear mode around 0.85, suggesting that many Grokipedia entries follow Wikipedia closely while introducing modest but non-trivial deviations—whether through rephrasing, reordering, or selective elaboration. Very low similarity scores are rare, indicating that even when Grokipedia diverges, it typically retains substantial semantic overlap with the corresponding Wikipedia article.

[5]We thank `@lthlnkso` for generous assistance with the Grokipedia data.

Figure 2: Similarity Scores between Grokipedia and Wikipedia



5 Estimation and Results

To identify the causal effect of Grokipedia on subsequent Wikipedia activity, we construct a matched sample of treated and untreated individuals using a Mahalanobis-distance nearest-neighbor matching procedure. The treatment indicator equals one for individuals whose Wikipedia pages are linked to a dedicated Grokipedia entry and zero otherwise. Because individuals differ systematically in their baseline visibility across occupations, all matching is performed *within occupational fields*. In particular, individuals classified as "Academia" are matched exclusively to other academics, and individuals classified as "Politics" are matched exclusively to other politicians. This ensures that comparisons are drawn between individuals operating in similar visibility environments and avoids cross-field differences that would otherwise confound treatment effects.

The matching procedure leverages a rich set of pre-treatment covariates that capture both short-run and long-run patterns of Wikipedia activity. For each individual,

we observe total page views and edit counts in the one-, two-, and three-week windows preceding the introduction of Grokipedia, long-run readership over 2015–2018, and the total number of words on the page. Let X_i denote this vector of pre-treatment characteristics. Matching on X_i helps approximate the counterfactual post-treatment trajectory of treated individuals had their pages not been featured on Grokipedia.

Matching is implemented in Stata using the `psmatch2` command. In particular, the estimator computes the Mahalanobis distance

$$D(i,j) = \sqrt{(X_i - X_j)' S^{-1} (X_i - X_j)},$$

where S is the sample covariance matrix of the covariates within each occupational field. Each treated unit i is matched to the single untreated unit j that minimizes $D(i,j)$, subject to a caliper of 0.25, which restricts matches to lie within a maximum allowable distance. Treated units for which no acceptable control exists are excluded, reducing bias due to poor-quality matches.

After constructing the matched sample, we estimate the average treatment effect on the treated (ATT) by comparing post-treatment outcomes of each treated individual to those of its matched control. Outcomes include the total number of post-treatment page views and edits, as well as analogous measures computed over successive weekly windows following Grokipedia's introduction. Because matching is performed within fields and conditions on detailed pre-treatment activity, the resulting ATT estimates isolate differences in Wikipedia activity that can be attributed to Grokipedia exposure rather than to pre-existing heterogeneity in visibility or page characteristics.

Table 2 reports the summary statistics for the treated sample—Wikipedia pages with a corresponding Grokipedia entry—and the matched control sample, separately for each field category (Academia, Culture, Leaders, Politics, and Sports). For each field, columns are presented in treated–control pairs, allowing direct visual comparison of pre-treatment characteristics.

A first observation concerns the number of observations. In every field, the control group is smaller than the treated group. This pattern is mechanical and expected: our matching procedure allows matching with replacement. A single control observation (a

Table 2: Summary Statistics After Matching

Variable	Academia (T)	Academia (C)	Culture (T)	Culture (C)	Leaders (T)	Leaders (C)	Politics (T)	Politics (C)	Sports (T)	Sports (C)
# of Edits	0.12	0.10	0.15	0.17	0.15	0.16	0.12	0.12	0.17	0.16
(1 week pre)	(0.52)	(0.46)	(0.65)	(0.81)	(0.60)	(0.64)	(0.52)	(0.54)	(0.62)	(0.63)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]
# of Edits	0.12	0.13	0.14	0.15	0.09	0.10	0.14	0.15	0.17	0.21
(2 week pre)	(0.51)	(0.54)	(0.66)	(0.64)	(0.46)	(0.55)	(0.57)	(0.55)	(0.61)	(0.77)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]
# of Edits	0.10	0.11	0.15	0.13	0.10	0.11	0.12	0.12	0.13	0.14
(3 week pre)	(0.47)	(0.44)	(0.81)	(0.67)	(0.48)	(0.60)	(0.58)	(0.57)	(0.55)	(0.58)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]
# of Views	409.61	376.77	928.57	855.59	535.97	467.65	523.41	445.85	427.69	364.50
(1 week pre)	(219.84)	(211.68)	(1034.13)	(1003.49)	(402.54)	(374.58)	(379.57)	(344.92)	(373.57)	(330.61)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]
# of Views	403.89	367.68	903.87	804.26	517.91	443.84	509.94	434.19	423.80	375.37
(2 week pre)	(219.97)	(210.90)	(996.14)	(947.40)	(386.48)	(359.75)	(364.48)	(332.83)	(347.06)	(312.54)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]
# of Views	407.49	360.65	902.65	787.84	517.20	431.19	509.64	423.20	417.76	374.78
(3 week pre)	(233.32)	(216.25)	(1020.14)	(968.73)	(396.60)	(365.76)	(378.25)	(335.30)	(361.59)	(335.42)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]
Wikipedia Page	11507.19	11434.20	8557.63	8505.34	10733.55	10665.41	9735.57	9649.28	11647.81	11602.68
Length	(8749.85)	(8713.73)	(7510.86)	(7477.56)	(9649.77)	(9625.46)	(8362.39)	(8321.07)	(10487.36)	(10475.38)
	[6239]	[3016]	[69048]	[17639]	[18738]	[6541]	[12303]	[5012]	[32661]	[14244]

Standard deviations in parentheses. Number of observations in brackets.

Wikipedia page without a Grokipedia entry) may therefore serve as a match for multiple treated observations. As a result, the total number of unique control pages is necessarily lower even though each treated page is matched to its closest neighbor.

Turning to the pre-treatment variables, the treated and control samples appear highly comparable across fields. For all three pre-period windows—one week, two weeks, and three weeks prior to the launch of Grokipedia—the average number of Wikipedia edits is nearly identical between the treated and control groups. Differences are small in magnitude and well within one standard deviation. This pattern suggests that both sets of pages exhibit similar levels of editorial activity prior to treatment, an important condition for our identification strategy.

A similar picture emerges when comparing Wikipedia page length. Across all fields, the baseline page size of treated and control pages is extremely close, with mean differences typically under 100 characters on bases of 8,000–12,000. The tight similarity in length reinforces the notion that the matched controls are good counterfactuals for the treated pages.

The only variable where we observe consistent, though modest, differences is the number of pageviews. Across all fields and all pre-treatment windows, the treated pages display slightly higher mean daily views than the control pages. However, the differences are relatively small compared to the large standard deviations—particularly in high-interest categories such as Culture and Sports. While not suggestive of strong baseline imbalance, this gap is consistent with the idea that the individuals who eventually receive Grokipedia pages tend to be somewhat more visible or publicly recognizable. Importantly, after matching, these differences remain economically limited and do not raise concerns about large pre-treatment divergence between groups.

Taken together, the summary statistics indicate that the treated and control samples are well balanced along key pre-treatment dimensions, with very similar editing histories and page characteristics. The modest differences in pageviews highlight that treated subjects may attract slightly more public attention, but the magnitude of the gap is small relative to the underlying variance. Overall, the matching procedure produces a credible comparison set for our empirical analysis of how Grokipedia's introduction affects subsequent Wikipedia contributions.

5.1 Results

Table 3 reports the estimated average treatment effects on the treated (ATT) obtained from the matched sample. For each individual i, we evaluate the impact of being listed on Grokipedia after its launch on October 27, 2025 on subsequent Wikipedia activity. The outcome variables are the total number of page views and the total number of edits accumulated over the three weeks following launch (October 27–November 16). Formally, for each outcome Y_i, we estimate:

$$\widehat{\text{ATT}} = \frac{1}{N_T} \sum_{i:grok_i=1} \left(Y_i^{\text{post}} - Y_{j(i)}^{\text{post}} \right),$$

where $j(i)$ denotes the nearest matched control for treated unit i, constructed using Mahalanobis-distance matching on pre-treatment page activity and page characteristics within each occupational field.

Across both occupations, the results show a clear and economically meaningful increase

Table 3: ATT Estimates After Grokipedia (Total)

	Academia		Culture		Leaders		Politics		Sports	
	Views	Edits	Views	Edits	Views	Edits	Views	Edits	Views	Edits
ATT	235.50***	0.09***	372.74***	0.02	315.86***	0.08***	344.28***	0.16***	235.72***	0.05***
	(18.29)	(0.03)	(18.04)	(0.01)	(17.63)	(0.02)	(22.83)	(0.02)	(16.72)	(0.01)
N (treated matched)	6239	6239	69048	69048	18738	18738	12303	12303	32661	32661

Standard errors in parentheses.
* $p < .1$, ** $p < .05$, *** $p < .01$

in Wikipedia activity following inclusion in Grokipedia. Among academics, total post-period views rise by approximately 236, with a standard error of 18.3. This represents an increase of roughly 20–25 percent relative to the typical pre-period weekly volume observed in academia. The number of edits also increases, by about 0.09 edits over three weeks, and the estimate is statistically significant at the 1 percent level.

The effects are even larger for political figures. Politicians whose pages appear on Grokipedia experience an average increase of 344 post-period views, with a standard error of 22.8, corresponding to an increase of approximately 25–30 percent compared to their pre-period levels. The number of edits rises by 0.16 over the same period, again with strong statistical significance.

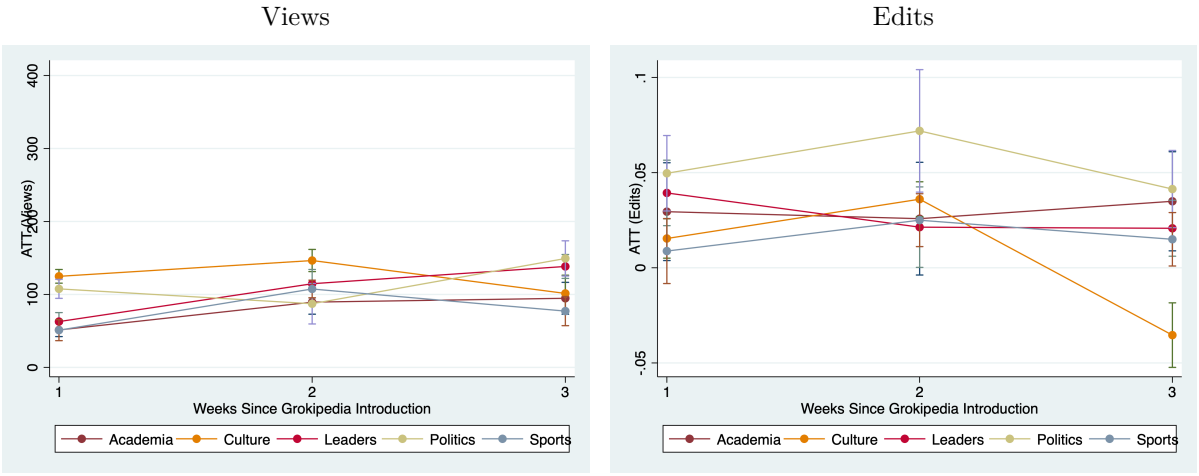
Taken together, the results indicate that the launch of Grokipedia meaningfully redirects user attention and editor activity toward the pages of treated individuals. The effects are statistically precise and substantively large, suggesting that Grokipedia functions as a powerful amplifier of traffic and engagement on Wikipedia within the fields it covers.

5.2 Dynamic Treatment Effects

Figure 3 presents the dynamic average treatment effects of Grokipedia's introduction on Wikipedia activity in the three weeks following the platform's launch on October 27, 2025. The left panel plots the estimated weekly effects on page views, while the right panel reports the corresponding effects on edit activity. Each point represents the ATT for a specific post-treatment week, together with its 95% confidence interval. These estimates are derived from the same matching procedure used in the main analysis, restricting comparisons to matched treated–control pairs within the same occupational field.

Several patterns emerge clearly. First, the weekly ATTs are economically meaningful

Figure 3: ATT Effects on Edits Over Time After Grokipedia



and quantitatively consistent with the main results reported in Table 3. In particular, the magnitudes of the week-specific estimates are roughly one-third of the cumulative three-week effect estimated in the main table, which is precisely what we would expect if the treatment effect is approximately constant throughout the three-week window. This serves as an internal consistency check for our design: the sum of the weekly effects lines up with the total effect estimated using the aggregated three-week outcome measure.

Second, the treatment effect appears relatively stable over time. For both views and edits, the weekly ATTs do not exhibit any notable attenuation or decay as the weeks progress. Instead, the estimates for week 1, week 2, and week 3 are statistically similar in magnitude, and in most cases we do not observe a significant decline that would suggest a rapid dissipation of the treatment effect. The only exception is pages in Culture, in which the ATT in the third week becomes negative (and statistically significant. This pattern is notable, given that many novelty-driven shocks tend to fade quickly; here, by contrast, Grokipedia's introduction seems to generate a persistent elevation in Wikipedia engagement at least within the immediate three-week horizon.

That said, we emphasize that the dynamic analysis currently captures only the short-run post-launch period. While the stability of the weekly ATTs is reassuring, it remains an open question whether the treatment effect may taper off, plateau, or even continue to grow over a longer horizon. As part of our revisions, we plan to extend the analysis to include additional weeks and months of post-treatment data as they become available.

This will allow us to evaluate whether the initial increase in Wikipedia activity induced by the emergence of Grokipedia represents a temporary shock or a more sustained shift in contributor and reader behavior.

5.3 Which Type of Wikipedia Pages Receive More Edits?

The anecdote from Grokipedia's launch already points to a potentially important mechanism behind our main findings: AI-generated summaries may sometimes operate as complements—rather than substitutes—to human editing activity on Wikipedia. On October 27, 2025, when Grokipedia first went live, its initial summary of Donald Trump included an unusual sentence describing how "Family dynamics, including Fred Sr.'s pragmatic engagement with local political machines for development approvals, further embedded a patriotic self-interest wary of overreaching authority, though direct *anti-communist* transmission remains less documented beyond the era's conservative milieu." Although this line was later removed in subsequent revisions, human editors on Wikipedia reacted almost immediately. Just three days later, on October 30, a Wikipedia editor added Trump's biography to the category "American Anti-Communists," a change that echoed themes appearing in the original Grokipedia summary. This episode highlights how AI-generated content might introduce new frames or specific details that prompt human editors to revisit Wikipedia entries, adding or reorganizing content in response.

Guided by this intuition, we examine whether pages for which Grokipedia provides more novel information—operationalized as a lower textual similarity between Grokipedia's summary and the pre-launch Wikipedia article—experience larger increases in editing activity following Grokipedia's debut. To do so, we estimate an OLS regression in which the dependent variable is the difference in total edits between treated pages and their matched controls during the post-launch period. The key independent variable is the similarity score computed between the Grokipedia and Wikipedia texts of the treated page. A lower similarity score indicates that Grokipedia offers information or framing not already present in the existing Wikipedia article. If the anecdotal Trump case is indicative of a broader pattern, we expect a negative coefficient on this variable: the more distinct Grokipedia's content is, the greater the subsequent human editing on Wikipedia.

The specification begins with a regression including only the similarity score. We then add field dummies—Culture, Leaders, Politics, and Sports—to capture systematic variation in editing patterns across subject areas. Finally, we include additional page-level controls, such as word count, gender, and age of the individual, to test whether the similarity mechanism is robust to these sources of heterogeneity. Table 4 presents the results.

Table 4: OLS Regression of $\triangle$ in Edits between Treated and Control Pages

	(1)	(2)	(3)
Similarity Score	-0.427***	-0.340*	-0.332*
	(0.128)	(0.137)	(0.143)
Field: Culture		-0.0900	-0.0931
		(0.0629)	(0.0664)
Field: Leaders		-0.0249	-0.0140
		(0.0693)	(0.0733)
Field: Politics		0.0729	0.0784
		(0.0742)	(0.0777)
Field: Sports		-0.0167	-0.0226
		(0.0683)	(0.0716)
Wikipedia Word Count (10,000)			0.0600***
			(0.0168)
Gender: Male			-0.0284
			(0.0340)
Gender: Other			0.142
			(0.424)
Age			-0.0000799
			(0.0000711)
Constant	0.427***	0.396**	0.362*
	(0.114)	(0.131)	(0.143)
Observations	46467	46467	43793
R^2	0.000	0.001	0.001

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Across all three columns, the coefficient on the similarity score is negative and statistically significant. In Column (1), the estimate is –0.427 and highly significant. Incorporating field fixed effects in Column (2) reduces the magnitude to –0.340, though the coefficient remains statistically significant at the 5 percent level. After adding the full set of page-specific controls in Column (3), the coefficient remains negative (–0.332) and significant. While the overall explanatory power of the regression is expectedly small—editorial activity is inherently noisy and influenced by many unobserved factors—the stability of the

similarity coefficient across specifications indicates a consistent relationship: Wikipedia pages that differ more from Grokipedia's summaries experience larger increases in human editing after Grokipedia's introduction.

It is important to note that the estimates in Table 4 are based on a subsample of 46,467 treated observations, rather than the full treated sample of 138,989 pages. This restriction arises because we currently lack complete Grokipedia and Wikipedia text for all treated pages, and therefore cannot compute similarity scores for the full dataset. Expanding the similarity-score construction to cover the complete set of treated biographies is a priority for the next draft. Doing so will allow us to test the mechanism more comprehensively and assess whether the relationship documented here holds across the broader sample.

Taken together, these results support the idea that generative AI can stimulate rather than suppress user-generated content. When Grokipedia introduces perspectives, narrative structures, or factual details that are not present in Wikipedia, editors appear more likely to respond. This complements the earlier anecdote and suggests that, in practice, AI and human contributors may interact in more complex and mutually reinforcing ways than a simple substitution framework would imply.

6 Conclusion

We find novel evidence about the complementarity of generative AI and human knowledge creation. One important aspect about our application–Wikipedia pages – is that contributors are not compensated for their labor. As with other private provision of public goods, the incentives may involve pro-social behavior (such as altruism). It will be interesting to see if our results extend to other examples such activity on the internet, such as social media and Creative Commons licensed software or data (these cases all potentially can be impacted by generative AI).

In our next revision of the paper, we will be able to look at longer terms impact of Grokipedia on Wikipedia. In addition to seeing if the patterns we have documented continue to hold over longer horizons, we will see if there is a different impact to Grokipedia revisions (as opposed to its initial creation). We will also be able to more closely look at one of the central research topics on Wikipedia, whether it is an unbiased source

(Greenstein and Zhu, 2012). Some have argued that Grokipedia reflects the partisan views of its owner Elon Musk. It will be interesting to see how Wikipedia responds to Grokipedia articles with a partisan slant: will they seek to maintain neutrality or to offset the bias with an opposing slant? Answering this question is important as evidence accumulates that such on-line knowledge-based impact and shape off-line behavior (Hinnosaar et al., 2023).

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics* 40 (S1):S293–S340.

Brooks, Creston, Samuel Eggert, and Denis Peskoff. 2024. "The Rise of AI-Generated Content in Wikipedia." *arXiv preprint* URL `https://arxiv.org/abs/2410.08044`.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics* 140 (2):889–942.

Burtch, Gordon, Dokyun Lee, and Zhichen Chen. 2024. "The Consequences of Generative AI for Online Knowledge Communities." *Scientific Reports* 14 (1):10413.

del Rio-Chanona, Maria, Nadzeya Laurentsyeva, and Johannes Wachs. 2023. "Are Large Language Models a Threat to Digital Public Goods? Evidence From Activity on Stack Overflow." *arXiv preprint* URL `https://arxiv.org/abs/2307.07367`.

Gallus, Jana. 2017. "Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia." *Management Science* 63 (12):3999–4015.

Glickman, Mark and Yi Zhang. 2024. "AI and Generative AI for Research Discovery and Summarization." *arXiv preprint* URL `https://arxiv.org/abs/2401.06795`.

Greenstein, Shane and Feng Zhu. 2012. "Is Wikipedia Biased?" *American Economic Review: Papers Proceedings* 102 (3):343–348.

Hale, James. 2025. "Wikipedia goes to war against AI slop with new deletion policy." *PCWorld* Accessed: 2025-11-23.

Heater, Brian. 2025. "Wikipedia editors push back against surge of AI-generated content." *The Verge* Accessed: 2025-11-23.

Hinnosaar, Marit, Toomas Hinnosaar, Michael Kummer, and Olga Slivko. 2023. "Wikipedia matters." *Journal of Economics & Management Strategy* 32 (3):657–669.

Laouenan, Morgane, Palaash Bhargava, Jean-Benoît Eyméoud, Olivier Gergaud, Guillaume Plique, and Etienne Wasmer. 2022. "A cross-verified database of notable people, 3500BC–2018AD." *Scientific Data* 9:290. URL https://doi.org/10.1038/s41597-022-01369-4.

Lyu, Liang, James Siderius, Hannah Li, Daron Acemoglu, Daniel Huttenlocher, and Asuman Ozdaglar. 2025. "Wikipedia Contributions in the Wake of ChatGPT." *arXiv preprint* URL https://arxiv.org/abs/2503.00757.

Miller, Marshall. 2025. "New User Trends on Wikipedia." *Diff, a Wikimedia community blog* Accessed: 2025-11-25.

Owen, Laura Hazard. 2025. "Wikipedia announces new AI strategy to support human editors." *Nieman Lab* Accessed: 2025-11-23.

Ren, Yuqing, Jilin Chen, and John Riedl. 2016. "The Impact and Evolution of Group Diversity in Online Open Collaboration." *Management Science* 62 (6):1668–1686.

Triedman, Harold and Alexios Mantzarlis. 2025. "What did Elon change? A comprehensive analysis of Grokipedia." *arXiv preprint* URL https://arxiv.org/abs/2511.09685.

Wikipedia contributors. 2024. "Wikipedia: Artificial intelligence." https://en.wikipedia.org/wiki/Wikipedia:Artificial_intelligence. Accessed: 2025-11-23.

Yasseri, Taha. 2025. "How Similar Are Grokipedia and Wikipedia? A Multi-Dimensional Textual and Structural Comparison." *arXiv preprint* URL https://arxiv.org/abs/2510.26899.

Zhang, Xiaoquan (Michael) and Feng Zhu. 2011. "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia." *American Economic Review* 101 (4):1601–1615.

Zhu, Kai, Dylan Walker, and Lev Muchnik. 2020. "Content Growth and Attention Contagion in Information Networks: Addressing Information Poverty on Wikipedia." *Information Systems Research* 31 (2):491–509.