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

Tin Cheuk Leung†
Wake Forest University

Koleman Strumpf‡
Wake Forest University

January 2026

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 by examining the short-run impact of the introduction of Grokipedia, an AI-generated online encyclopedia operated by xAI, which provides automated entries that could either substitute for human editing or draw in new contributors. We develop a simple theoretical framework in which AI entries can redirect user attention and stimulate human editing through novel framing, yielding ambiguous effects on UGC ex ante. Using a new panel dataset covering 1.4 million Wikipedia pages of notable individuals, we exploit the fact that only a subset have comparable Grokipedia entries to estimate the causal effect of AI on subsequent human contributions, 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 entries act as an attention amplifier rather than a pure substitute for Wikipedia content. Exploiting variation in the semantic similarity between Grokipedia entries and their corresponding Wikipedia articles, we further show that pages with *lower* similarity experience *larger* increases in editing after Grokipedia's launch, consistent with the model's predictions. Taken together, our findings indicate 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 Joshua Gans, Marit Hinnosaar, Shane Greenstein, Imke Reimers, Peter Thiel and Russell Tsz-Nga Wong for their comments. We also 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

Generative artificial intelligence (GAI) systems rely fundamentally on vast amounts of user-generated content (UGC). Large language models are trained primarily on text drawn from the open internet, including encyclopedias, discussion forums, question-and-answer platforms, blogs, and news articles—sources that are overwhelmingly produced and curated by unpaid human contributors (Dodge et al., 2021). Wikipedia alone comprises tens of millions of volunteer-authored articles and is widely recognized as a core high-quality training source for modern language models. This reliance creates an inherent feedback loop: AI systems trained on UGC increasingly generate outputs that may substitute for the very platforms that supply their training data. 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 entries for notable individuals and topics using proprietary and publicly available data, offering users an alternative source of information outside the Wikipedia ecosystem.[2]

Theoretically, the effect of Grokipedia on Wikipedia's UGC is ambiguous. AI may act as a substitute: if users rely on Grokipedia entries 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

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

[2]One concern is that Grokipedia has limited visibility. However, as soon as it was released the site had significant web traffic (Section 3). There is also a revealed preference argument for its importance: we find it has economically significant impacts on Wikipedia activity. Finally, several Wikipedia Talk pages, where editors discuss potential improvements, explicitly mention Grokipedia as an inspiration for changes (for example see "Turkey" or "Runescape").

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 entry 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 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 entry.[3] This episode provides anecdotal evidence that AI-generated entries on an external platform may inform or prompt human editing behavior on Wikipedia.

To study this, we first develop a simple theoretical framework that clarifies how generative AI can simultaneously affect user attention and human editing incentives. In the model, AI entries may redirect traffic toward or away from Wikipedia while also introducing novel framing that lowers the cognitive cost of identifying gaps or inconsistencies in existing content. These forces generate two testable predictions: (i) AI can increase editing activity on average when attention expansion or novelty dominates substitution, and (ii) conditional on AI being present, pages for which AI entries differ more from existing content should experience larger editorial responses.

A key challenge in the literature examining AI's empirical impact is to find well-defined and exogenous variation in exposure. In our application we have explicit treatments and controls, and we use free public data which allows other researchers to verify and extend our results.

We assemble a new panel dataset of approximately 1.4 million Wikipedia pages of notable individuals across five occupational domains—Academia, Culture, Leaders, Politics, and Sports. Only about 170,000 of these pages (roughly 12%) are covered by Grokipedia

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

at launch, while the remaining 88% do not receive an AI-generated entry. This sparse treatment assignment is central to our identification strategy, as it provides a large and heterogeneous pool of untreated but otherwise comparable Wikipedia pages that can serve as credible counterfactuals for pages exposed to Grokipedia.[4]

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 identifying Grokipedia coverage. Our empirical design follows the matching framework developed in Section 6. Because Grokipedia coverage is concentrated among highly visible pages and baseline visibility varies systematically across occupations, we construct matched samples using Mahalanobis-distance nearest-neighbor matching within occupational fields. Treated pages—those with a Grokipedia entry—are paired with the closest untreated pages based on pre-treatment views, editing histories, long-run readership, and page characteristics, approximating the counterfactual trajectory each treated page would have followed absent Grokipedia. We then estimate treatment effects using a Difference-in-Differences framework, which compares changes in views and edits before and after Grokipedia's launch between treated pages and their matched controls.

Across both approaches, our results reveal a surprising and robust pattern: the availability of Grokipedia entries leads to an increase, rather than a decrease, in human-generated edits. Wikipedia pages with an AI entry experience significantly more editing activity relative to matched control pages. We also document a substantial rise in Wikipedia readership following Grokipedia's introduction, suggesting that AI-generated entries 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 semantic similarity measure between Grokipedia entries and their corresponding Wikipedia pages. The distribution is skewed toward high overlap—about 20% of entries are nearly

[4]Our analysis focuses on the short-run impact of Grokipedia's introduction. We show that Grokipedia attracted an immediate surge in user attention when it was introduced. Concentrating on this initial post-launch period is a natural way to study how the sudden availability of AI-generated entries interacts with ongoing human editing activity on Wikipedia, before longer-run adjustments in user behavior, platform integration, and editorial norms take place. Characterizing these longer-run dynamics is an important topic for future research but lies beyond the scope of this paper.

identical to Wikipedia text—yet exhibits meaningful variation, with most scores between 0.5 and 1 and a mode around 0.8. We then use this variation to test a mechanism suggested by early anecdotal evidence: AI-generated entries may introduce new frames or details that prompt Wikipedia 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 develop a stylized model of generative AI and user-generated content supply in Section 4. We describe the data in Section 5. We then discuss our empirical strategy and results in Section 6. The last section concludes.

2 Literature

Our study contributes to three distinct streams of literature: the economic impact of generative AI on knowledge production, the dynamics of platform competition, and the incentives underlying the private provision of public goods.

First, we extend the rapidly expanding literature on how generative AI shapes user content creation. A central debate in this field is whether AI functions as a substitute that crowds out human effort or a complement that enhances it. Recent empirical work

has largely documented a substitution effect, finding that the availability of tools like ChatGPT is associated with declines in human-authored posts and participation in online communities (Lyu et al., 2025; Burtch, Lee, and Chen, 2024; Glickman and Zhang, 2024; del Rio-Chanona, Laurentsyeva, and Wachs, 2024).[5] These studies typically model AI as a productivity tool adopted by the agent to lower the marginal cost of content generation (Brynjolfsson, Li, and Raymond, 2025).

We depart from this framework by studying AI not as a tool used by the worker, but as an external competitive force (Grokipedia) entering the knowledge market. In our setting, the AI does not lower the cost of editing Wikipedia directly; rather, it produces a rival product that competes for user attention. We find that this competitive entry stimulates, rather than displaces, human contribution. This aligns with recent work on "human-in-the-loop" systems where AI serves as a "recombinant" tool, lowering the search costs for identifying knowledge gaps or framing new ideas (Agrawal, Gans, and Goldfarb, 2018). Our results suggest that AI-generated content serves as a reference point that highlights missing information or alternative perspectives. In other words, AI activates an "inspiration" channel: its novel framing and distinct informational content make gaps in the existing record salient, prompting editors to engage in expansionary and integrative work to close the distance between the platforms.[6]

Second, our findings connect to the literature on digital platform competition and strategic differentiation. Standard models of platform entry suggest that when a low-cost, automated entrant (Grokipedia) appears, the incumbent platform (Wikipedia) faces pressure to differentiate (Eisenmann, Parker, and Van Alstyne, 2011; Zhu and Iansiti, 2012). While we do not observe the Wikipedia platform owner making strategic changes,

[5]To this extent, our results also 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 consider 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.

[6]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.

we observe the community of editors endogenously responding to the entrant.

Our theoretical mechanism—where AI redirects attention and introduces novel framing—differs from supply-side models such as Gans (2024). Gans models AI as reducing the cost of low-complexity tasks, which shifts human effort toward higher-difficulty problems through a free-rider threshold effect. In contrast, our framework emphasizes demand-side channels—attention, visibility, and semantic disagreement—through which AI outputs can actively stimulate human editing even when task costs are unchanged. As a result, similarity between AI and human output plays opposite roles in the two models: redundancy dampens contribution in Gans (2024), whereas in our setting novelty and framing differences can directly trigger corrective and expansionary human effort.

Third, we contribute to the literature on the private provision of public goods, particularly in online knowledge commons.[7] Prior work establishes that contributors to platforms like Wikipedia are motivated by peer recognition, social impact, and audience size (Zhang and Zhu, 2011; Gallus, 2017; Hinnosaar et al., 2022). A key concern in this literature is whether automated content generation will erode these intrinsic motivations by making human effort feel redundant.

Our results alleviate this concern by highlighting the "guardianship" and "expansion" motives. We show that editing activity is highest precisely where the AI's output differs most from the human consensus. This suggests that the presence of an AI competitor activates an identity motive among editors to not only correct the record but to expand it—ensuring their platform remains the superior source of truth. This dynamic mirrors findings in organizational science where professional communities respond to automation by reasserting jurisdiction over complex or nuanced tasks (Autor, Levy, and Murnane, 2003).

In summary, we bridge the gap between labor-substitution models and platform strategy. We show that in a non-monetary, public-goods setting, the entry of an AI competitor generates a "flight to quality," where human contributors increase their investment to maintain the differentiation and accuracy of their platform.

[7]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.

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 attention was also visible within the Wikimedia community itself. In a public blog post published on November 10, 2025—less than two weeks after Grokipedia's launch—the Wikimedia Foundation explicitly addressed the rise of generative AI encyclopedias, emphasizing Wikipedia's continued value and implicitly positioning AI-native projects such as Grokipedia as salient points of comparison (Wikimedia Foundation, 2025).

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.[8] Launched in Oct 27th, 2025, its stated purpose is to provide a faster-moving, AI-supported alternative to Wikipedia.

[8]Grokipedia's announcement was largely a surprise. Elon Musk announced he would create an AI-based alternative to Wikipedia only three weeks before the roll-out, and the first public mention of this project occurred a month before that. Google Trends data for "Grokipedia" was at zero until the week of Elon Musk's announcement.

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.

From an industrial-organization perspective, competition between Wikipedia and Grokipedia is highly non-standard. Both platforms supply privately produced services with public-good characteristics and compete primarily for user attention rather than price, while allowing unrestricted content spillovers across platforms. Unlike conventional product markets, neither side can exclude the other from inspecting, adapting, or contesting its output. The two platforms also face starkly asymmetric production technologies and governance structures: Wikipedia relies on a resource-constrained volunteer labor force supplemented by automation, whereas Grokipedia operates as a centralized, capital-intensive AI system with minimal marginal costs but idiosyncratic ownership control. These differences imply that competitive pressure need not generate displacement. Instead, rivalry may reallocate effort across human and machine inputs, potentially stimulating selective editorial responses rather than crowding out human contribution.[9]

Crucially for our analysis, Grokipedia's launch coincided with a sharp and highly concentrated surge in user attention. Traffic estimates from SimilarWeb indicate that prior to launch, Grokipedia attracted only a few thousand global visits per day. Beginning on October 27, 2025—the day the platform went live—traffic spiked abruptly, reaching over 500,000 daily visits worldwide during its first week, with more than 100,000 daily visits per day originating from the United States alone. Peak attention occurred immediately after launch, exceeding two million global visits on October 28, before declining in the following days.

[9]Wikipedia content is released under a Creative Commons Attribution–ShareAlike license, which permits extensive reuse and adaptation, implying that Grokipedia's close textual overlap does not raise standard copyright or fair-use constraints despite exceeding what would typically be permissible in proprietary markets. More importantly for imitation incentives, Wikipedia enforces community-governed sourcing and neutrality norms, including lists of non-authoritative sources that editors generally refuse to cite. Grokipedia does not publicly articulate comparable constraints. As a result, similarity reflects heterogeneous informational quality: overlap based on reputable sources may facilitate imitation, whereas overlap driven by partisan or non-authoritative outlets need not translate into editorial response.

This sharp but short-lived surge in attention motivates our focus on short-run outcomes. Rather than attempting to characterize long-run platform substitution or equilibrium adoption—processes that unfold gradually as user habits, platform integration, and editorial norms evolve—we study the immediate post-launch period, when Grokipedia was most visible and salient to Wikipedia editors. During this window, editors were suddenly exposed to a competing AI-generated entry of the same underlying subjects, often framed differently and updated rapidly. Focusing on this short-run response allows us to directly assess how the introduction of AI-generated encyclopedic content interacts with ongoing human editorial activity: whether it primarily displaces human effort or instead acts as a complementary upstream input that triggers additional editing, fact-checking, and article expansion on Wikipedia.

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 A Stylized Model of Generative AI and UGC

We present a stylized model to clarify when generative AI acts as a substitute for, or a complement to, user-generated content (UGC), and to derive predictions consistent with our empirical findings.

Consider a Wikipedia page indexed by i. A unit mass of users consume information about topic i, and a fraction $\lambda \in (0, 1)$ of users are potential editors who may contribute to the Wikipedia page. Editors are intrinsically motivated and receive no monetary compensation.

An external AI platform (Grokipedia) may generate a entry for page i. Let $G_i \in \{0, 1\}$

indicate whether a Grokipedia entry exists. When $G_i = 1$, the AI entry may overlap with the existing Wikipedia content. Let $s_i \in [0, 1]$ denote the semantic similarity between the AI entry and the pre-existing Wikipedia article, where higher values correspond to greater overlap.

4.1 User Attention

User traffic to Wikipedia depends on baseline informational value and on the presence of AI entries. Specifically, page views are given by

$$\text{Views}_i = \bar{V}_i + G_i(\theta - \delta s_i),$$

where $\bar{V}_i$ is baseline attention, $\theta \geq 0$ captures an attention-amplification effect from AI exposure, and $\delta \geq 0$ reflects substitution when the AI entry closely replicates Wikipedia content.

Lemma 1. ***(Average Effect of AI on Page Views)*** *Suppose $G_i = 1$ for treated pages and let $s_i \in [0, 1]$ denote the semantic similarity between the AI entry and the pre-existing Wikipedia article. Then the introduction of AI entries increases Wikipedia page views on average if*

$$\theta > \delta\, \mathbb{E}[s_i].$$

Moreover, the condition $\theta \geq \delta$ is a sufficient (but stronger) condition that guarantees page views weakly increase for all treated pages.

Lemma 1 highlights the distinction between average and pointwise effects of AI on user attention. The parameter θ captures the extent to which AI entries act as an attention gateway—by increasing awareness, lowering search costs, or redirecting users toward Wikipedia—while δs_i captures substitution, which is stronger when the AI entry closely replicates existing content. When similarity varies across pages, substitution affects some pages more than others, but attention expansion applies uniformly. As a result, page views can increase on average even if AI substitutes for Wikipedia on highly similar pages. The sufficient condition $\theta \geq \delta$ rules out substitution dominating attention for any treated page, while the weaker condition $\theta > \delta\mathbb{E}[s_i]$ allows for heterogeneous effects across pages.

4.2 Editing Incentives

Each potential editor e faces an idiosyncratic editing cost c_{ie} drawn from a continuous distribution with CDF $F(\cdot)$ and density $f(\cdot) > 0$. Editors choose whether to contribute a discrete unit of effort.

The benefit from editing page i is

$$b_i = \kappa \cdot \text{Views}_i + G_i \cdot \rho(1 - s_i),$$

where $\kappa > 0$ captures pro-social benefits that increase with audience size (Zhang and Zhu, 2011), and $\rho > 0$ reflects an "inspiration" channel: when AI introduces novel framing or information, editors are more likely to respond.

Editor e contributes if $b_i \geq c_{ie}$. Aggregate editing activity for page i is therefore

$$E_i = F(b_i).$$

4.3 Implications and Testable Predictions

The model delivers three testable predictions that guide our empirical analysis.

Proposition 1. ***(Average Complementarity)*** *Suppose $\rho > 0$ and $\theta \geq \delta$. Then the introduction of a Grokipedia entry ($G_i = 1$) weakly increases Wikipedia editing activity for every treated page, and strictly increases expected editing activity on average relative to pages without a Grokipedia entry ($G_i = 0$).*

Proposition 1 highlights that generative AI can complement user-generated content even when it partially substitutes for information consumption. The key distinction is between the incentives of readers and editors. While high-similarity AI entries may reduce the marginal value of visiting Wikipedia for some users, the model assumes that Grokipedia does not reduce overall attention to treated pages and instead acts as an attention gateway. At the same time, by introducing new framing or emphasizing different aspects of a topic, AI entries create salient contrasts that motivate human editors to respond. As long as Grokipedia expands attention and provides some novelty, these forces

dominate substitution effects, leading to a net increase in Wikipedia editing activity on average.

Proposition 2. ***(Similarity Gradient Among Treated Pages)*** *Among Wikipedia pages that receive a Grokipedia entry ($G_i = 1$), aggregate editing activity is strictly decreasing in the semantic similarity between the Grokipedia entry and the pre-existing Wikipedia article:*

$$\frac{\partial E_i}{\partial s_i} < 0.$$

Proposition 2 reflects the idea that the editorial response to generative AI depends on how informative the AI output is relative to existing content. When a Grokipedia entry closely mirrors the Wikipedia article, it reveals little that editors did not already know, generating minimal incentive to revise or expand the page. By contrast, when the AI entry departs from Wikipedia—by introducing new facts, alternative framing, or emphasizing neglected aspects—it makes gaps and discrepancies salient, prompting editors to respond. In the model, higher similarity weakens both the direct "inspiration" effect of AI and the attention substitution effect operating through page views, leading to a strictly negative relationship between similarity and subsequent editing activity.

Proposition 3. ***(Irrelevance of Similarity for Control Pages)*** *For Wikipedia pages without a Grokipedia entry ($G_i = 0$), editing activity is independent of semantic similarity s_i.*

Together, these predictions imply that similarity matters only within the treated sample, a feature that aligns directly with our empirical design. Formal proofs of all propositions are provided in Appendix A.[10]

[10] Appendix B extends the baseline model to allow Wikipedia editing to endogenously affect semantic similarity with Grokipedia and to capture editors' incentives to deliberately differentiate Wikipedia content in order to mitigate substitution. This extension provides a theoretical rationale for why the relationship between similarity and page views may be weak in the data, as strategic differentiation can offset substitution effects. Importantly, allowing similarity to be endogenous does not alter any of the model's main predictions: Propositions 1, 2, and 3 continue to hold, and are formally shown to be invariant to this extension.

5 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 will 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. Applying this procedure to the 1.4 million notable individuals in our sample, we identify approximately 170,000 individuals with a Grokipedia entry, implying that about 12% of notable people with an English Wikipedia page in our focal fields were covered by Grokipedia during our data-collection window.[11] This coverage, however, varies substantially across occupational domains: Grokipedia presence is highest in Culture (22.2%), followed by Leaders (12.1%), Politics (8.2%), Sports (8.2%), and Academia (6.5%). Combining all elements of the scraped data yields a unified person-level dataset that includes biographical iden-

[11] 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.

tifiers, detailed Wikipedia traffic and edit histories, pre- and post-treatment Wikipedia content, and Grokipedia presence and text where available.

5.1 Who Were Selected to Have a Grokipedia Page?

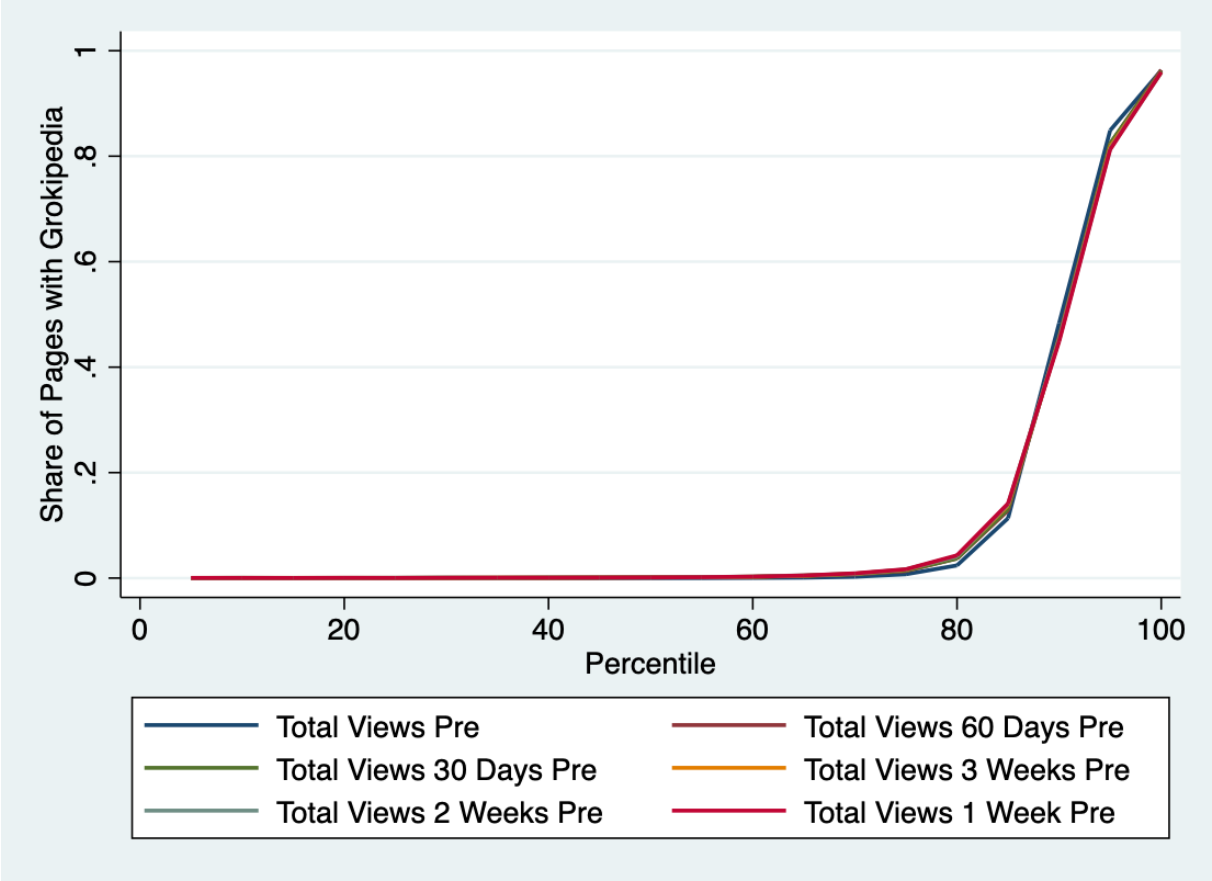
Although the precise criteria Grokipedia uses to decide which Wikipedia pages receive a Grokipedia entry are not publicly documented, it is not unreasonable to assume that the platform prioritizes pages that already attract substantial attention on Wikipedia. This motivates our approach of examining 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. We 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 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 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.[12]

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.

[12] As will be discussed in the next Section, the estimates require a Wikipedia page with a Grokipedia entry (treatment) to have a comparable page without one (control). Pages in 95th percentile and above lack a control (since almost all are treated), so they are not part of the estimation sample.

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



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 visibility of the page immediately before treatment. We progressively enrich the specification by adding prior editing activity, 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.

Table 1: OLS Regression on Grokipedia Presence

	Dependent variable: Grok indicator			
	(1)	(2)	(3)	(4)
ln(Views 60 Pre)	0.151***	0.147***	0.147***	0.145***
	(0.000133)	(0.000148)	(0.000152)	(0.000181)
ln(Edits 60 Pre +1)		0.0280***	0.0261***	0.0192***
		(0.000412)	(0.000412)	(0.000430)
Field: Culture			0.0665***	0.0667***
			(0.000743)	(0.000799)
Field: Leaders			0.0402***	0.0455***
			(0.000822)	(0.000896)
Field: Politics			0.0338***	0.0360***
			(0.000821)	(0.000880)
Field: Sports			0.0669***	0.0639***
			(0.000710)	(0.000764)
Wikipedia Word Count (10,000)				0.0309***
				(0.000325)
Gender: Male				-0.00540***
				(0.000570)
Gender: Other				0.0893***
				(0.0117)
Age				-0.0000341***
				(0.00000142)
Constant	-0.791***	-0.771***	-0.823***	-0.816***
	(0.000826)	(0.000872)	(0.00108)	(0.00132)
Observations	1422093	1422093	1422093	1291333
R^2	0.477	0.479	0.483	0.503

Standard errors in parentheses

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

Column (2) augments this baseline specification by adding the logarithm of the number of Wikipedia edits in the same pre-launch window, ln(Edits 60 Pre + 1), capturing pre-existing editorial activity beyond page views alone. The results are very similar to those in column (1). The coefficient on ln(Views 60 Pre) remains stable at around 0.15, while ln(Edits 60 Pre + 1) enters positively and is statistically significant, with a coefficient of approximately 0.03. This suggests that pages with more active editing histories were somewhat more likely to receive Grokipedia coverage, but that this effect is economically

much smaller than that of page views. Consistent with this interpretation, the explanatory power of the model increases only marginally, with the R^2 rising from 0.477 to 0.479.

Column (3) 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 0.479 in column (2) to 0.483 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 (4) 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 the R^2 from 0.483 in column (3) to 0.503 in column (4). 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.

5.2 Similarity Between Grokipedia and Wikipedia Text

To assess how closely Grokipedia entries resemble their corresponding Wikipedia pages, we compute a semantic similarity score for the set of *treated* Wikipedia pages for which both the Grokipedia text and the relevant historical version of the corresponding Wikipedia article are available (Grokipedia uses the same slugs to identify pages as Wikipedia).[13] Throughout the paper, we refer to a Wikipedia page as *treated* if it is linked to a dedicated Grokipedia entry, and as a *control* if no such Grokipedia page exists. These treated and control pages are paired in our empirical design to exhibit similar pre-treatment trends in Wikipedia activity—such as page views and edit histories—so that differences observed after Grokipedia's introduction can be more plausibly attributed to the presence of AI-generated entries rather than to pre-existing differences. While we introduce our formal matching procedure and causal framework in Section 6, this terminology is useful here to distinguish pages that are directly exposed to AI-generated entries from those that are not.

Our similarity measure is based on a hybrid approach that combines two conceptually distinct ways of comparing text: semantic similarity and lexical similarity. Each captures a different notion of how similar two documents are, and neither is sufficient on its own in our setting.

The first component relies on a transformer-based sentence embedding model (`all-MiniLM-L6-v2`). This class of models represents an entire document as a point in a high-dimensional semantic space, where documents that convey similar meanings are located close to one another even if they use different words, sentence structures, or organizational patterns. For example, two biographies that describe the same career trajectory but differ in phrasing, level of detail, or ordering of events will typically be mapped to nearby vectors. We quantify this notion of similarity using cosine similarity between the two embedding vectors, yielding a score between 0 and 1. A higher value indicates greater semantic alignment. Importantly, this measure is designed to be robust to paraphrasing and reorganization: it captures whether two texts "say the same thing," not whether they say it in the same way.

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

The second component uses a traditional TF–IDF–based cosine similarity measure. TF–IDF (term frequency–inverse document frequency) represents each document as a weighted vector of words, placing greater weight on terms that are informative within the corpus while downweighting common function words. Cosine similarity between these vectors captures lexical overlap—that is, the extent to which two texts use the same vocabulary and phrases. This measure is particularly sensitive to near-verbatim copying or light editing, but it treats paraphrases that use different wording as substantially less similar.

The motivation for combining these two measures is that they respond differently to the types of variation we observe between Grokipedia and Wikipedia. To illustrate, consider the entry for *Rodney (wrestler)*. The Grokipedia and Wikipedia texts cover essentially the same biographical content—including early life, professional wrestling career, championships, and post-retirement activities—but differ in presentation. The Grokipedia version compresses some paragraphs, removes explicit section headers, and occasionally reorders information relative to Wikipedia. As a result, the TF–IDF similarity between the two texts is relatively high (0.877), reflecting substantial overlap in names, dates, organizations, and wrestling-related terminology, while the embedding-based semantic similarity is lower (0.772), reflecting differences in structure and narrative flow rather than content. Both measures are informative: the TF–IDF score indicates close textual overlap, while the embedding score indicates that the texts are not literal copies but remain strongly aligned in meaning.

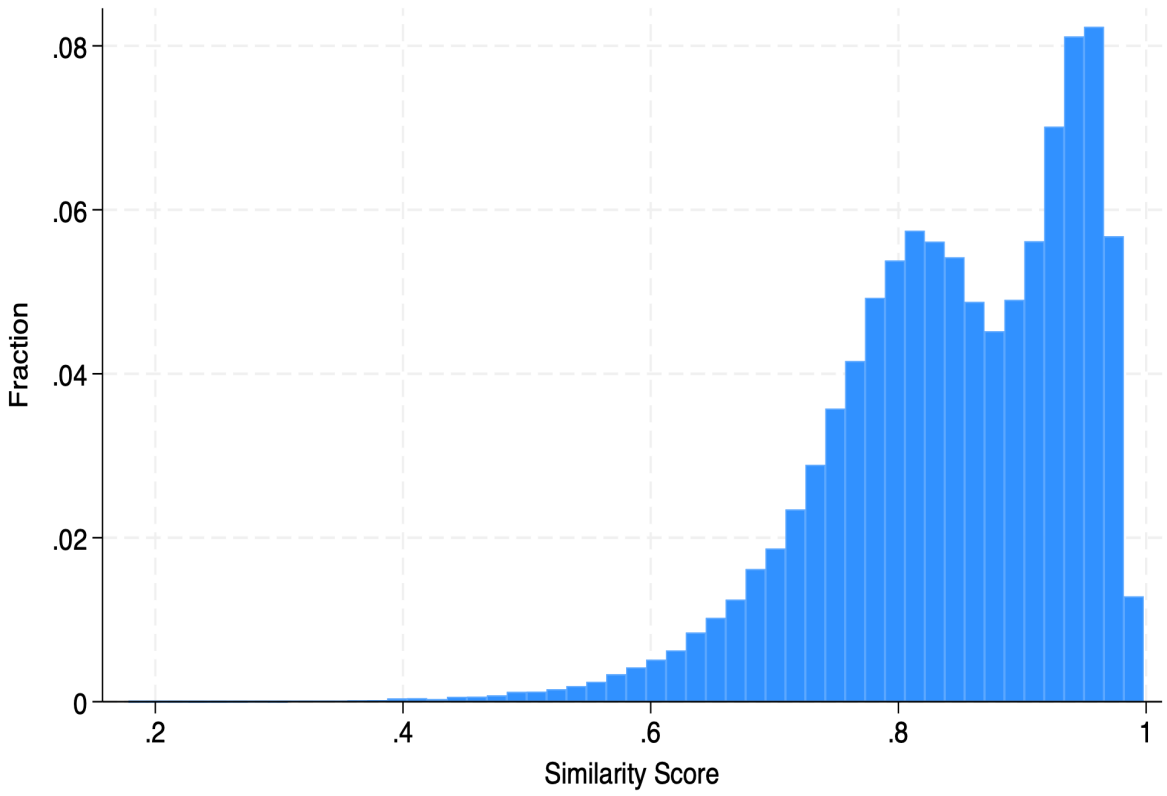
Formally, our final similarity score is a weighted average of the embedding-based similarity and the TF–IDF similarity, with greater weight placed on the semantic embedding component.[14] This weighting reflects our emphasis on meaning rather than exact wording, while still allowing the measure to identify cases in which Grokipedia closely reproduces Wikipedia text with minimal modification. The resulting hybrid metric therefore distinguishes between entries that are nearly identical, those that closely track Wikipedia

[14]We place greater weight on the embedding-based similarity (0.6 versus 0.4 for TF–IDF) because our primary interest is in substantive overlap in content rather than purely mechanical copying. Embedding-based measures are more robust to paraphrasing and reorganization, which are common differences between Grokipedia and Wikipedia entries, while TF–IDF similarity is more sensitive to near-verbatim reproduction. Our results are qualitatively similar under alternative weighting schemes and when using either component on its own.

while introducing modest stylistic or organizational changes, and those that diverge more substantially in emphasis or detail.

Figure 2 plots the distribution of similarity scores across the treated sample. Three features stand out. First, a substantial mass of observations is concentrated at very high similarity levels: close to 20% of pages have a similarity score above 0.95, indicating that the Grokipedia entry is nearly identical to its Wikipedia counterpart. These cases are consistent with direct replication or only trivial modification of the Wikipedia text. Second, among the remaining pages, similarity scores are broadly distributed between roughly 0.5 and 1, with a clear mode around 0.8. This pattern suggests that many Grokipedia entries closely track Wikipedia while introducing modest but non-trivial changes, such as rephrasing, reordering of material, or selective elaboration. Finally, very low similarity scores are rare, indicating that even when Grokipedia departs from Wikipedia, it typically retains substantial semantic overlap with the original article rather than introducing entirely unrelated content.

Figure 2: Similarity Scores between Grokipedia and Matched Wikipedia Pages



6 Matching 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.

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.[15] 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. Applying this procedure to the full set of 174,771 treated Wikipedia pages linked to a Grokipedia entry, we are able to identify an acceptable matched control for 138,989 pages; the remaining treated pages are excluded from the matched sample

[15] Matching is implemented in Stata using the `psmatch2` command.

due to the absence of a suitable control satisfying the caliper and field-specific matching criteria.[16]

After constructing the matched sample, we estimate the Difference-in-Differences (DiD) estimates by comparing changes in outcomes for treated pages relative to matched control pages before and after the launch of Grokipedia on October 27, 2025. 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 DiD estimates isolate differences in Wikipedia activity that can be attributed to Grokipedia exposure rather than to pre-existing heterogeneity in visibility or page characteristics.

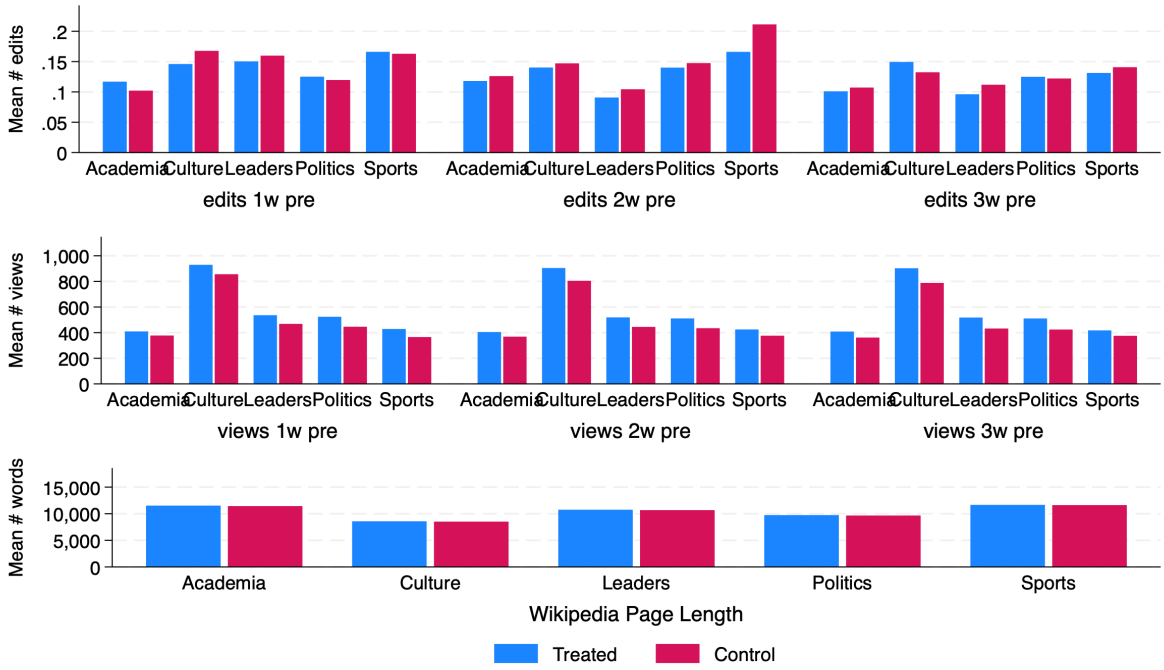
Figure 3 provides a visual comparison of pre-treatment characteristics between the treated pages (those that eventually receive a Grokipedia entry) and their matched controls. The top panel reports the mean number of Wikipedia edits in the one-, two-, and three-week windows prior to Grokipedia's launch; the middle panel shows mean pageviews over the same pre-treatment horizons; and the bottom panel displays baseline Wikipedia page length measured in number of words.

Across fields, treated and control pages are highly similar along most dimensions. In particular, the top panel shows no meaningful differences in pre-treatment editing activity: for all three pre-periods and across all categories, the mean number of edits for treated and control pages is nearly identical. Likewise, the bottom panel indicates that Wikipedia page length is very closely balanced between the two groups, suggesting that treated and control pages are comparable in terms of baseline article scope and content depth.

The only systematic difference visible in the figure arises in the middle panel, which plots pre-treatment pageviews. Treated pages exhibit slightly higher Wikipedia pageviews than their matched controls across fields and time windows. However, these differences are modest in magnitude and small relative to overall variation in views, especially when

[16] In Appendix C, we implement a substantially finer matching strategy as a robustness check. Matching is restricted to individuals within the same occupational field *and* country of citizenship, and we additionally incorporate birth year into the Mahalanobis distance metric. These stricter requirements sharply reduce the number of matched treated observations but yield even larger and highly statistically significant treatment effects on both views and edits, reinforcing the main conclusions of this section.

Figure 3: Matched Pre-Treatment Summary Statistics: Treated vs Control



compared to the post-treatment effects documented in Section 6.1.

Overall, the matching procedure produces a credible comparison set for our empirical analysis of how Grokipedia's introduction affects subsequent Wikipedia activity. Appendix D provides a more detailed discussion of the corresponding summary statistics, including formal treated–control differences and their statistical significance.

6.1 Results

Table 2 reports Difference-in-Differences (DiD) estimates obtained from the matched sample, i.e. we compare changes in Wikipedia activity for treated pages relative to matched control pages before and after the launch of Grokipedia on October 27, 2025.[17] 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), relative to the corresponding three-week pre-period.

Formally, let $j(i)$ denote the matched control page for treated page i, constructed

[17] Appendix E reports the average treatment effects on the treated (ATT) which are broadly consistent with the DiD results presented in this section.

using Mahalanobis-distance matching on pre-treatment page activity and page characteristics within each occupational field. The matched-pair Difference-in-Differences (DiD) estimand is given by

$$\widehat{\text{DiD}} \;=\; \frac{1}{N_T} \sum_{i:\,\text{Treated}_i=1} \Big[\big(Y_i^{\text{post}} - Y_{j(i)}^{\text{post}}\big) - \big(Y_i^{\text{pre}} - Y_{j(i)}^{\text{pre}}\big) \Big],$$

where Y_i^{post} and Y_i^{pre} denote the post-launch and pre-launch outcomes for page i, respectively, and N_T is the number of treated pages in the matched sample. This estimator differences out both permanent differences between treated and control pages and common time shocks affecting all pages.

Table 2: DiD Estimates of Grokipedia Impact

	Academia		Culture		Leaders		Politics		Sports	
	Views	Edits	Views	Edits	Views	Edits	Views	Edits	Views	Edits
DiD	119.63***	0.09***	85.34***	0.03**	87.46***	0.12***	104.53***	0.16***	81.12***	0.10***
	(17.99)	(0.03)	(17.76)	(0.01)	(17.27)	(0.02)	(22.43)	(0.03)	(16.57)	(0.02)
N (# Treated)	6239	6239	69048	69048	18738	18738	12303	12303	32661	32661

Standard errors in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$

Across all occupational fields, the DiD estimates reveal a clear and economically meaningful increase in Wikipedia activity following inclusion in Grokipedia. To gauge magnitudes, we benchmark the DiD estimates against the average pre-treatment outcomes for treated pages.

In Academia, treated pages experience an increase of approximately 120 post-period views relative to matched controls. Compared to a pre-treatment average of roughly 407 weekly views for treated academic pages, this corresponds to an increase of about 30 percent over the three-week post-launch window. The number of edits rises by about 0.09 over the same period, which represents an increase of roughly 80 percent relative to the treated group's average pre-period editing rate of approximately 0.11 edits per week.

Similar patterns emerge in Culture and Leaders, though with more moderate relative magnitudes. Cultural figures see an increase of roughly 85 views, corresponding to about a 9 percent increase relative to a pre-treatment average of approximately 912 weekly views. Edits increase by about 0.03, or roughly 20 percent relative to the pre-period editing

mean. For Leaders, the estimated increase of approximately 87 views corresponds to a 15–20 percent rise relative to a pre-treatment average of about 524 weekly views, while the 0.12 increase in edits represents an increase of roughly 100 percent relative to pre-period editing activity.

The largest relative effects are observed among Political figures and Sports figures. Political pages experience an increase of approximately 105 views, corresponding to a 20 percent increase relative to a pre-treatment average of about 514 weekly views, alongside an increase of 0.16 edits, which amounts to an increase of more than 100 percent relative to pre-period editing levels. Sports figures see an increase of roughly 81 views, or about 20 percent relative to a pre-treatment mean of approximately 423 weekly views, and an increase of 0.10 edits, corresponding to a 60–70 percent rise relative to baseline editing activity.

Taken together, these magnitudes indicate that the effects of Grokipedia are not only statistically significant but also economically substantial. While view increases tend to range from roughly 10 to 30 percent across fields, the relative impact on edits is markedly larger, often exceeding 50 percent and in several cases approaching or exceeding a doubling of pre-treatment editing activity. These patterns underscore the importance of the attention-amplification channel emphasized in the model and highlight the responsiveness of human editorial effort to relatively modest expansions in readership.

These findings speak directly to the theoretical framework developed in Section 4. The positive DiD effects on page views across all fields are consistent with Lemma 1, which characterizes the conditions under which AI-generated entries expand user attention to Wikipedia by acting as complements rather than substitutes for the underlying articles. The fact that treated pages experience larger increases than matched controls after the launch indicates that the attention-expansion channel dominates any substitution arising from content overlap.

At the same time, the positive and statistically significant DiD effects on edits across every field provide direct empirical support for Proposition 1. In the model, average complementarity arises when increased attention and novel framing jointly raise the marginal return to human editorial effort. The results show that, relative to matched controls,

treated pages receive more edits after Grokipedia's launch, implying that generative AI exposure increases rather than crowds out voluntary human contributions.

Importantly, the matched-pair DiD design sharpens the interpretation of these results by differencing out both permanent page-level heterogeneity and common time shocks. While Lemma 1 governs how AI exposure affects user attention, Proposition 1 governs how that attention translates into human effort. The consistent positive effects for both views and edits across all fields indicate that the attention-amplification channel emphasized in the model is active and that, on net, this expansion in audience stimulates editorial activity.

Taken together, these findings indicate that the launch of Grokipedia meaningfully redirects both user attention and editor activity toward treated Wikipedia pages. In line with Lemma 1 and Proposition 1, generative AI operates not as a crowding-out force but as a catalyst that expands readership and encourages complementary human contributions across a wide range of occupational domains.

6.2 Dynamic Treatment Effects

Figure 4 presents the dynamic average treatment effects on the treated (ATT) of Grokipedia's introduction on Wikipedia activity in the three weeks before and after the platform's launch on October 27, 2025. Formally, for each outcome Y_{it} in week t, we estimate:[18]

$$\widehat{\text{ATT}}_t \;=\; \frac{1}{N_T} \sum_{i:grok_i=1} \Big(Y_{it} - Y_{j(i)t} \Big),$$

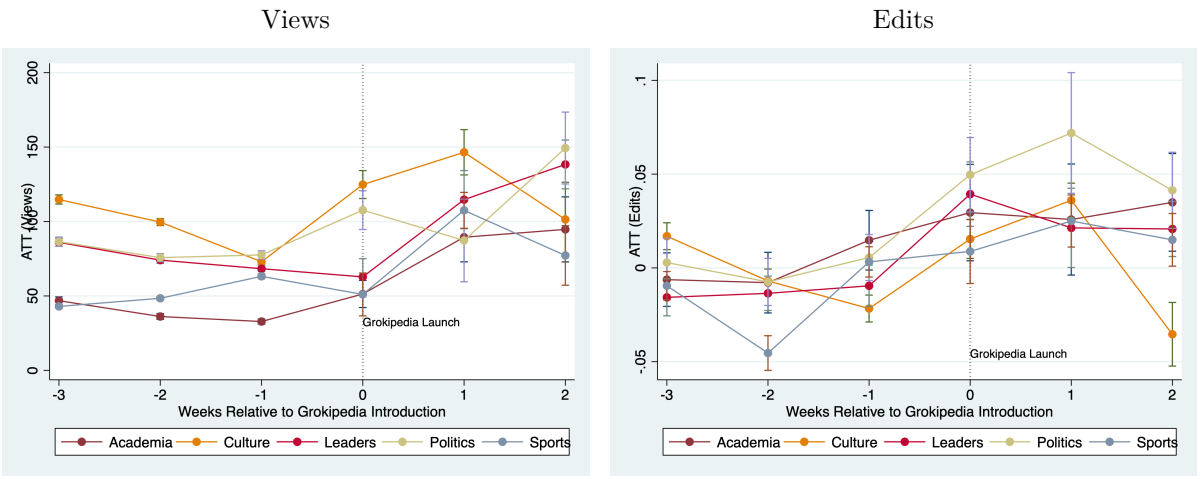
where $j(i)$ denotes the nearest matched control for treated page i.

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 pre- or post-treatment week, together with its 95% confidence interval.

Several patterns emerge clearly. First, the weekly ATTs are economically meaningful and quantitatively consistent with the main results reported in Table 2. In particular,

[18] This section considers ATT rather than DiD as in the previous section. This is done to consider values for the pre-treatment period. The post-treatment patterns discussed below continue to hold if we look at DiD (net out the pre-treatment differences).

Figure 4: ATT Effects Over Time Before and After Grokipedia



the differences between the post-treatment and pre-treatment week-specific estimates are each roughly one-third of the cumulative three-week effect reported in the main table, which is exactly what one would expect if the treatment effect is approximately constant over the three-week post-launch 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, for both views and edits, the estimated ATTs are relatively stable in the weeks prior to Grokipedia's introduction, with no systematic upward or downward pattern, providing little evidence of differential pre-trends between treated and control pages. At the time of Grokipedia's launch (week 0) or right after the launch (week 1), however, we observe a clear and discrete increase in the ATT across most fields, indicating an immediate response to the introduction of AI-generated entries.

Following the launch, the treatment effect remains remarkably stable over time. For both outcomes, the post-treatment week-specific ATTs do not exhibit meaningful attenuation or decay as the weeks progress. The estimates in weeks 0, 1, and 2 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. This pattern contrasts with many novelty-driven shocks, which tend to fade quickly, and instead suggests that Grokipedia's introduction generates a persistent increase in Wikipedia engagement over at least the immediate post-launch horizon. The main exception is pages in the Culture category,

where the ATT becomes negative and statistically significant by the third post-launch week, indicating a short-lived or potentially crowd-out effect in that domain.

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.

6.3 Which Type of Wikipedia Pages Receive More Edits?

As discussed in the Introduction, the Wikipedia entry for Donald Trump had extensive edits that seemed to mirror new text from the Grokipedia entry. The anecdote illustrates how AI-generated content can introduce new framing or highlight omissions that prompt human editors to revisit and update Wikipedia entries.

Proposition 2 formalizes this intuition. In the model, once a Grokipedia entry exists, higher semantic similarity between the AI entry and the Wikipedia article weakens editing incentives by revealing little new information, whereas lower similarity makes gaps and discrepancies more salient. As a result, aggregate editing activity among treated pages is predicted to be strictly decreasing in similarity.[19]

Guided by this prediction, we examine whether Wikipedia pages for which Grokipedia provides more novel information—operationalized as a lower textual similarity between the Grokipedia entry and the pre-launch Wikipedia article—experience larger increases in editing activity following Grokipedia's debut. To do so, we exploit the matched treated–control design and focus on DiD estimates in outcomes within each matched pair. Let i index a treated Wikipedia page and let $j(i)$ denote its matched control. Our outcome of interest is the DiD estimates in total edits,

$$\Delta Y_i^{\text{DiD}} \equiv (Y_i^{\text{post}} - Y_{j(i)}^{\text{post}}) - (Y_i^{\text{pre}} - Y_{j(i)}^{\text{pre}}).$$

[19] In our theoretical framework, similarity can also affect Wikipedia page views through a substitution channel, captured by the parameter δ, whereby higher similarity reduces click-through to Wikipedia. Empirically, however, we find no statistically significant relationship between similarity and the post-treatment difference in page views (Appendix F), suggesting that δ is close to zero or that offsetting mechanisms cancel out in the data. Importantly, this does not affect the theoretical results in Section 4: Proposition 1 requires only that Grokipedia does not reduce attention on average, and Proposition 2 relies on the direct "inspiration" channel linking similarity to editing incentives, which does not depend on similarity affecting page views.

All the conclusions discussed below continue to hold if instead we use as the outcome variable an ATT, $\Delta Y_i^{\text{ATT}} \equiv Y_i^{\text{post}} - Y_{j(i)}^{\text{post}}$.

We estimate the following OLS specification:

$$\Delta Y_i^{\text{DiD}} = \alpha + \beta \, \text{Similarity}_i + \mathbf{X}_i' \boldsymbol{\gamma} + \varepsilon_i,$$

where Similarity_i is the semantic similarity score between the Grokipedia entry and the pre-launch Wikipedia text for treated page i, and $\mathbf{X}_i$ denotes a vector of controls. A lower similarity score indicates that Grokipedia introduces information or framing not already present in the Wikipedia article. Consistent with Proposition 2, the model predicts $\beta < 0$.

We begin with a specification that includes only the similarity score. We then add field indicators—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. Table 3 presents the results.

Across all three specifications, the coefficient on the similarity score is negative and statistically significant. In Column (1), the estimate is -0.354 and highly significant. Incorporating field fixed effects in Column (2) reduces the magnitude to -0.321, 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 (-0.311) and statistically significant. While the overall explanatory power of the regression is modest—editorial activity is inherently noisy—the stability of the similarity coefficient across specifications provides strong support for the similarity gradient predicted by Proposition 2.[20]

Beyond statistical significance, the estimates also imply economically meaningful effects. Focusing on Column (3), a one–standard deviation increase in the similarity score (approximately 0.1) is associated with a decline of about 0.03 edits in pre-treatment edit-

[20]The interpretation of the similarity coefficient depends on the nature of the information introduced by Grokipedia. When an AI entry incorporates new facts from reputable sources, imitation by Wikipedia editors is expected, as such content lowers the cost of identifying verifiable omissions. In contrast, overlap driven by non-authoritative sources (e.g., partisan outlets) might not translate into imitation. To test this, we construct an indicator for whether a Grokipedia page cites any source classified by Wikipedia as banned or generally unreliable and interact it with the similarity score. The coefficient on similarity remains negative and significant (about –0.27), while the interaction term is positive (around 0.06) but statistically insignificant, consistent with imitation operating primarily through credible sources.

Table 3: OLS Regression of Δ in Edits between Treated and Control Pages

	(1)	(2)	(3)
Similarity Score	-0.354***	-0.321***	-0.311***
	(0.0843)	(0.0861)	(0.0900)
Field: Culture		-0.0568	-0.0650
		(0.0416)	(0.0438)
Field: Leaders		0.0281	0.0253
		(0.0460)	(0.0486)
Field: Politics		0.0615	0.0633
		(0.0490)	(0.0511)
Field: Sports		0.0442	0.0383
		(0.0438)	(0.0459)
Wikipedia Word Count (10,000)			0.0431***
			(0.0106)
Gender: Male			-0.0543*
			(0.0218)
Gender: Other			-0.115
			(0.277)
Age			-0.0000301
			(0.0000475)
Constant	0.379***	0.361***	0.364***
	(0.0716)	(0.0821)	(0.0885)
Observations	133241	133241	125598
R^2	0.000	0.000	0.001

Standard errors in parentheses

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

ing activity. While this absolute change may appear small, it corresponds to roughly 2.5 percent of a standard deviation in pre-treatment edits among treated pages (with a standard deviation of about 1.2). Given the inherently low and highly skewed baseline level of individual-page editing activity, this magnitude is non-trivial.

Taken together, these results show that the complementary effect of generative AI on Wikipedia editing is not uniform across treated pages. Instead, it depends critically on how much new information or framing the AI entry introduces relative to existing content. When Grokipedia closely mirrors Wikipedia, the editorial response is muted; when it departs from Wikipedia, editors are more likely to engage. This pattern reconciles the positive average effect documented earlier with substantial within-treated heterogeneity and underscores the role of AI as a catalyst for human contribution rather than a simple substitute.

7 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 of 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).

A separate and promising avenue for future work concerns the longer-run and dynamic interactions between AI-generated knowledge and human editorial processes. Beyond the short-run attention effects studied in this paper, longer horizons may reveal different patterns as user attention stabilizes, platform integration evolves, and norms around AI use mature. Relatedly, subsequent research could examine how revisions to AI-generated content—rather than its initial introduction—affect human editing behavior, drawing closer parallels to the iterative editing process that characterizes Wikipedia itself. Another important direction is the interaction between generative AI and informational neutrality, a longstanding concern in the Wikipedia literature (Greenstein and Zhu, 2012). To the extent that AI-generated encyclopedic content reflects the perspectives or incentives of its developers, future work could study how human editors respond—whether by reinforcing neutrality norms or by counteracting perceived bias—and the broader implications for how online knowledge platforms shape offline beliefs and 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.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.

Autor, David H, Frank Levy, and Richard J Murnane. 2003. "The skill content of recent technological change: An empirical exploration." *The Quarterly Journal of Economics* 118 (4):1279–1333.

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, R Maria, Nadzeya Laurentsyeva, and Johannes Wachs. 2024. "Large language models reduce public knowledge sharing on online Q&A platforms." *PNAS Nexus* 3 (9):pgae400.

Dodge, Jesse, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, and Matt Gardner. 2021. "Documenting the English Colossal Clean Crawled Corpus." *ArXiv* abs/2104.08758. URL `https://api.semanticscholar.org/CorpusID:233296858`.

Eisenmann, Thomas, Geoffrey Parker, and Marshall Van Alstyne. 2011. "Platform envelopment." *Strategic Management Journal* 32 (12):1270–1285.

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

Gans, Joshua S. 2024. "Will user-contributed AI training data eat its own tail?" *Economics Letters* 242:111868.

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 E. Kummer, and Olga Slivko. 2022. "Externalities in knowledge production: evidence from a randomized field experiment." *Experimental Economics* 25 (2):706–733.

———. 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.

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

Wikimedia Foundation. 2025. "In the AI era, Wikipedia has never been more valuable." URL https://wikimediafoundation.org/news/2025/11/10/in-the-ai-era-wikipedia-has-never-been-more-valuable/. Wikimedia Foundation blog post.

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, Feng and Marco Iansiti. 2012. "Entry into platform-based markets." *Strategic Management Journal* 33 (1):88–106.

Appendix A Proofs for Section 4

We provide formal proofs for the three propositions stated in Section 4.

Appendix A.1 Setup

Aggregate editing activity on page i is given by

$$E_i = F(b_i),$$

where $F(\cdot)$ is a continuous cumulative distribution function with strictly positive density $f(\cdot)$, and

$$b_i = \kappa \cdot \text{Views}_i + G_i \cdot \rho(1 - s_i).$$

Page views satisfy

$$\text{Views}_i = \bar{V}_i + G_i(\theta - \delta s_i),$$

with parameters $\kappa, \rho, \theta, \delta > 0$.

Appendix A.2 Proof of Proposition 1 (Average Complementarity)

Aggregate editing activity on page i is given by $E_i = F(b_i)$, where $F(\cdot)$ is a strictly increasing cumulative distribution function and

$$b_i = \kappa \cdot \text{Views}_i + G_i \cdot \rho(1 - s_i).$$

Recall that page views satisfy

$$\text{Views}_i = \bar{V}_i + G_i(\theta - \delta s_i).$$

For a treated page ($G_i = 1$), aggregate editing activity is therefore

$$E_i(1) = F\Big(\kappa(\bar{V}_i + \theta - \delta s_i) + \rho(1 - s_i)\Big),$$

while for an untreated page ($G_i = 0$),

$$E_i(0) = F(\kappa \bar{V}_i).$$

The difference in editing activity between treated and untreated pages is

$$E_i(1) - E_i(0) = F\Big(\kappa(\bar{V}_i + \theta - \delta s_i) + \rho(1 - s_i)\Big) - F(\kappa \bar{V}_i).$$

Because $F(\cdot)$ is strictly increasing, a sufficient condition for $E_i(1) > E_i(0)$ is that

$$\kappa(\theta - \delta s_i) + \rho(1 - s_i) \geq 0.$$

Since $s_i \in [0, 1]$, this condition holds for all treated pages whenever

$$\kappa(\theta - \delta) \geq 0 \quad \text{and} \quad \rho > 0.$$

Under these parameter restrictions, the introduction of a Grokipedia entry weakly increases editing incentives for every treated page, with strict inequality for all pages with $s_i < 1$.

It follows immediately that

$$\mathbb{E}[E_i(1)] > \mathbb{E}[E_i(0)],$$

establishing that the introduction of Grokipedia increases expected Wikipedia editing activity on average. □

Appendix A.3 Proof of Proposition 2 (Similarity Gradient Among Treated Pages)

For treated pages ($G_i = 1$), aggregate editing activity is

$$E_i(s_i) = F(\kappa \cdot \text{Views}_i + \rho(1 - s_i)).$$

Differentiating with respect to s_i yields

$$\frac{\partial E_i}{\partial s_i} = f(\kappa \cdot \text{Views}_i + \rho(1 - s_i)) \cdot (-\kappa\delta - \rho)\,.$$

Because $\kappa, \rho > 0$, $\delta \geq 0$, and $f(\cdot) > 0$, it follows that

$$\frac{\partial E_i}{\partial s_i} < 0.$$

Hence, among treated pages, editing activity is strictly decreasing in similarity. □

Appendix A.4 Proof of Proposition 3 (Irrelevance of Similarity for Control Pages)

For control pages ($G_i = 0$),

$$b_i = \kappa \cdot \bar{V}_i, \quad E_i = F(\kappa \cdot \bar{V}_i),$$

which does not depend on s_i. Therefore,

$$\frac{\partial E_i}{\partial s_i} = 0,$$

and similarity has no effect on editing activity for untreated pages. □

Appendix B Model Extension: Endogenous Similarity and Strategic Differentiation

This section extends the baseline model by allowing Wikipedia editing activity to endogenously affect the semantic similarity between Wikipedia and Grokipedia, and by introducing incentives for editors to deliberately differentiate Wikipedia content from AI-generated entries. We then show formally that this extension does not alter the main theoretical predictions.

Appendix B.1 Similarity Dynamics

Conditional on editing, an editor chooses a level of differentiation effort $d_i \geq 0$. Similarity evolves according to

$$s_i' = (1 - \lambda)s_i + \lambda \bar{s}_i - \eta d_i,$$

where $\bar{s}_i$ denotes the similarity that would arise absent strategic differentiation, $\lambda \in (0, 1)$ governs persistence, and $\eta > 0$ measures the effectiveness of differentiation effort. In other words, $\bar{s}_i$ denotes the similarity that would arise if editors exert no differentiation effort ($d_i = 0$), reflecting unavoidable overlap due to shared facts and common source material.

Appendix B.2 Editing Benefits with Differentiation

With endogenous similarity, the gross benefit from editing page i is

$$B_i(d_i) = \kappa \cdot \text{Views}_i + \rho\Big(1 - s_i\Big) + \psi\delta\Big(1 - s_i'\Big),$$

where $\rho > 0$ captures an "inspiration" channel triggered by the pre-edit novelty of the AI entry relative to the existing Wikipedia text, and $\psi \geq 0$ captures a strategic differentiation motive: editors value reducing post-edit similarity in order to mitigate substitution when δ is large.

Page views are given by

$$\text{Views}_i = \bar{V}_i + G_i(\theta - \delta s_i'),$$

and editing costs are quadratic in differentiation effort.

Appendix B.3 Optimal Differentiation Choice

Let the cost of differentiation be

$$C(d_i) = \frac{c_d}{2} d_i^2, \qquad c_d > 0.$$

Under the extension, inspiration depends on the *pre-edit* novelty of the AI entry relative to Wikipedia, while substitution and strategic differentiation depend on the *post-edit* similarity s_i'. Specifically, conditional on editing, the editor chooses $d_i \geq 0$ to maximize

$$\max_{d_i \geq 0} \left\{ \kappa\Big(\bar{V}_i + G_i(\theta - \delta s_i')\Big) + G_i \rho(1 - s_i) + G_i \psi \delta (1 - s_i') - \frac{c_d}{2} d_i^2 \right\},$$

where similarity evolves as

$$s_i' = (1 - \lambda) s_i + \lambda \bar{s}_i - \eta d_i.$$

Lemma 2. ***(Endogenous Differentiation)*** *For treated pages ($G_i = 1$), the optimal differentiation effort is*

$$d_i^* = \frac{\eta}{c_d}(\kappa + \psi)\delta,$$

which is (weakly) increasing in the substitution parameter δ and strictly positive whenever $\delta > 0$. For untreated pages ($G_i = 0$), $d_i^* = 0$.

Proof. Substitute $s_i' = (1 - \lambda) s_i + \lambda \bar{s}_i - \eta d_i$ into the objective. The terms involving d_i are

$$\kappa\Big(- G_i \delta s_i' \Big) + G_i \psi \delta (1 - s_i') - \frac{c_d}{2} d_i^2 = \text{const} + G_i \eta d_i (\kappa + \psi)\delta - \frac{c_d}{2} d_i^2,$$

where $\rho(1 - s_i)$ does not depend on d_i because inspiration is triggered by the pre-edit similarity. For $G_i = 1$, the objective is strictly concave in d_i and the first-order condition is

$$\eta(\kappa + \psi)\delta - c_d d_i = 0,$$

which yields $d_i^* = \frac{\eta}{c_d}(\kappa + \psi)\delta$. For $G_i = 0$, d_i only increases costs and provides no benefits, so $d_i^* = 0$. □

Appendix B.4 Reduced-Form Net Benefit

Define the optimized net benefit

$$\tilde{b}_i \equiv \max_{d_i \geq 0} b_i(d_i),$$

where $b_i(d_i)$ is the editor's payoff inside the braces above.

Lemma 3. *For treated pages ($G_i = 1$), the optimized net benefit is*

$$\tilde{b}_i = \kappa\Big(\bar{V}_i + \theta - \delta\Big[(1-\lambda)s_i + \lambda\bar{s}_i\Big]\Big) + \rho(1-s_i) + \psi\delta\Big(1 - \Big[(1-\lambda)s_i + \lambda\bar{s}_i\Big]\Big) + \frac{\eta^2}{2c_d}\Big((\kappa+\psi)\delta\Big)^2.$$

For untreated pages ($G_i = 0$), $\tilde{b}_i = \kappa\bar{V}_i$.

Proof. Let $s_i^0 \equiv (1-\lambda)s_i + \lambda\bar{s}_i$ denote post-edit similarity absent differentiation ($d_i = 0$). Substituting d_i^* from Lemma 2 implies $s_i' = s_i^0 - \eta d_i^*$. Plugging d_i^* into the objective yields the stated expression, where the last term is the standard quadratic maximization gain:

$$\max_{d_i \geq 0}\left\{\eta d_i(\kappa+\psi)\delta - \frac{c_d}{2}d_i^2\right\} = \frac{\eta^2}{2c_d}\Big((\kappa+\psi)\delta\Big)^2.$$

For $G_i = 0$, the objective does not contain s_i' and d_i, so $\tilde{b}_i = \kappa\bar{V}_i$. □

Aggregate editing activity remains

$$E_i = F(\tilde{b}_i),$$

where $F(\cdot)$ is the CDF of editing costs.

Appendix B.5 Invariance of the Main Propositions

We now show that Propositions 1–3 continue to hold.

Proposition 1. From Lemma 3, for treated pages,

$$\tilde{b}_i(1) - \tilde{b}_i(0) = \kappa\Big(\theta - \delta s_i^0\Big) + \rho(1-s_i) + \psi\delta(1-s_i^0) + \frac{\eta^2}{2c_d}\Big((\kappa+\psi)\delta\Big)^2, \quad s_i^0 = (1-\lambda)s_i + \lambda\bar{s}_i.$$

Since $s_i^0 \in [0,1]$, the condition $\theta \geq \delta$ implies $\kappa(\theta - \delta s_i^0) \geq \kappa(\theta - \delta) \geq 0$. Moreover, $\rho(1 - s_i) \geq 0$ with strict inequality whenever $s_i < 1$, and the remaining terms are nonnegative. Therefore $\tilde{b}_i(1) \geq \tilde{b}_i(0)$ pointwise, with strict inequality on a set of positive measure. Because $F(\cdot)$ is strictly increasing, it follows that $E_i(1) \geq E_i(0)$ and Proposition 1 holds.

Proposition 2. For treated pages, differentiating Lemma 3 with respect to s_i yields

$$\frac{\partial \tilde{b}_i}{\partial s_i} = -\rho - (1-\lambda)\delta(\kappa + \psi) < 0,$$

since $\rho > 0$, $\kappa > 0$, $\delta \geq 0$, $\psi \geq 0$, and $\lambda \in (0,1)$. Since $f(\cdot) > 0$,

$$\frac{\partial E_i}{\partial s_i} = f(\tilde{b}_i)\frac{\partial \tilde{b}_i}{\partial s_i} < 0,$$

establishing Proposition 2.

Proposition 3. For untreated pages, $\tilde{b}_i = \kappa \bar{V}_i$ does not depend on s_i, implying

$$\frac{\partial E_i}{\partial s_i} = 0.$$

Thus Proposition 3 continues to hold. □

Appendix C Robustness on Mahalanobis Matching

This appendix presents a robustness check that strengthens the matching design used in Section 6 by imposing substantially finer restrictions on the construction of matched control pages. The goal is to address the concern that residual heterogeneity—particularly along geographic and cohort dimensions—may influence both the likelihood of receiving a Grokipedia entry and subsequent Wikipedia activity.

In the baseline analysis, treated and control Wikipedia pages are matched within occupational fields using Mahalanobis-distance nearest-neighbor matching on pre-treatment measures of page activity and article characteristics. While this approach ensures comparability in visibility and editing dynamics, individuals within the same field may still differ systematically across countries in ways that affect public attention, language usage, and editorial norms.

To further tighten the counterfactual comparison, we refine the matching procedure in two ways. First, we restrict matching to occur within the same country of citizenship, as identified by the individual's primary citizenship variable. This ensures that treated and control pages are drawn from the same national context, thereby absorbing country-level differences in Wikipedia usage, media attention, and contributor behavior. Second, we augment the Mahalanobis distance metric to include the individual's birth year, which helps align treated and control pages along cohort dimensions that may proxy for historical prominence, generational relevance, or differential exposure to contemporary media environments.

Formally, matching is now performed within the intersection of occupational field and country, using Mahalanobis distance on the same set of pre-treatment activity variables as in the baseline specification, with birth year added as an additional covariate. All other aspects of the matching procedure—including one-to-one nearest-neighbor matching with replacement and the use of a caliper—remain unchanged.

As expected, imposing these finer restrictions substantially reduces the number of treated observations for which an acceptable match can be found. Many treated pages—particularly those from smaller countries or rare cohorts—do not have a sufficiently close untreated counterpart within the same field–country cell and are therefore excluded.

The resulting matched sample sizes are considerably smaller than in the baseline analysis. Specifically, the number of treated observations retained under the finer matching procedure is 2,536 for Academia, 40,788 for Culture, 7,563 for Leaders, 4,457 for Politics, and 18,015 for Sports. These reductions reflect the more stringent identification strategy rather than data loss or changes in outcome measurement.

Table A1 reports the difference-in-difference estimates (DiD) using this finer matching design. Despite the sharp reduction in sample size, the main findings of Table 2 not only persist but become stronger. Across all five occupational fields, the estimated effects of Grokipedia exposure on both post-treatment page views and edits remain positive and statistically significant at conventional levels.

Table A1: DiD Estimates of Grokipedia Impact (Finer Matching)

	Academia		Culture		Leaders		Politics		Sports	
	Views	Edits	Views	Edits	Views	Edits	Views	Edits	Views	Edits
DiD	69.97***	0.10***	176.90***	0.14***	98.67***	0.02	191.41***	0.18***	103.60***	0.13***
	(17.21)	(0.03)	(14.67)	(0.01)	(20.48)	(0.03)	(41.51)	(0.04)	(14.33)	(0.01)
N (# Treated)	2536	2536	40788	40788	7563	7563	4457	4457	18015	18015

Standard errors in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$

Quantitatively, the magnitudes of most of the DiD estimates are larger than those reported in the baseline specification. For example, in the Culture field, treated pages experience an increase of approximately 177 post-treatment views and 0.14 edits, compared to substantially smaller effects (85 for views and 0.03 for edits) in the baseline matching. Similar amplification is observed across Politics, and Sports. The precision of the estimates remains high despite the smaller samples (with the exception of number of edits among Leaders which is still positive but statistically insignificant), indicating that the effects are not driven by a small number of influential observations.

Taken together, these results provide strong evidence that the positive impact of Grokipedia on subsequent Wikipedia activity is not an artifact of coarse matching across heterogeneous populations. Even when comparisons are restricted to individuals from the same country and similar birth cohorts—thereby absorbing a wide range of unobserved geographic and generational factors—the introduction of AI-generated entries continues to increase both readership and human editing activity.

Appendix D Detailed Summary Statistics on Matching

This appendix provides detailed summary statistics for the matched treated and control samples used in the empirical analysis. The discussion here complements the exposition in Section 6, where the matching procedure is introduced, and Figure 3, which presents a compact visual comparison of treated–control differences without focusing on specific numerical values. The purpose of this appendix is to document the underlying magnitudes and statistical properties of these pre-treatment differences in greater detail.

Table A2 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.

Table A2: Comparing Pre-trends: Summary Statistics After Matching (T=Treated, C=Control)

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.

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 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 pre-treatment editing activity, treated and control pages are highly comparable across all fields. For each of the three pre-period windows—one week, two weeks, and three weeks prior to Grokipedia's launch—the mean number of Wikipedia edits differs only marginally between the two groups. These differences are small in magnitude and, as shown formally in Table A3, are generally statistically insignificant. This similarity in baseline editing intensity supports the validity of the matched control group as a counterfactual for treated pages.

A similar conclusion applies to Wikipedia page length. Across all fields, treated and control pages have nearly identical baseline lengths. Mean differences are small relative to overall page size and are not statistically distinguishable from zero. The close alignment in page length further reinforces that the matching procedure successfully balances core article characteristics prior to treatment.

The only dimension along which systematic differences emerge is pageviews. As documented in Table A3, treated pages exhibit statistically significant differences in pre-treatment views across all fields and time windows. However, these differences are economically modest in size. Even in high-visibility categories such as Culture, Politics, and Sports, the absolute gaps are small relative to overall variation in pageviews and are consistent with treated individuals being only slightly more publicly visible prior to receiving a Grokipedia entry. Importantly, these pre-treatment view differences are substantially smaller than the post-treatment divergences documented in Section 6.1, indicating that they are unlikely to drive the main findings.

Overall, the evidence in this appendix confirms that the treated and control samples are closely balanced along key pre-treatment dimensions. Differences in edits and page length are negligible and statistically insignificant, while differences in views—though

Table A3: Difference between Treated and Control

Variable	Academia	Culture	Leaders	Politics	Sports
# of Edits	0.01	-0.02**	-0.01	0.01	0.00
(1 week pre)	(1.18)	(-2.25)	(-0.87)	(0.54)	(0.42)
# of Edits	-0.01	-0.01	-0.01	-0.01	-0.05***
(2 week pre)	(-0.59)	(-0.70)	(-1.39)	(-0.63)	(-4.73)
# of Edits	-0.01	0.02*	-0.02	0.00	-0.01
(3 week pre)	(-0.51)	(1.92)	(-1.58)	(0.26)	(-1.35)
# of Views	32.83***	72.98***	68.32***	77.56***	63.19***
(1 week pre)	(4.90)	(3.29)	(7.14)	(8.39)	(10.41)
# of Views	36.20***	99.61***	74.07***	75.75***	48.43***
(2 week pre)	(5.45)	(4.61)	(8.08)	(8.46)	(8.29)
# of Views	46.84***	114.80***	86.01***	86.44***	42.98***
(3 week pre)	(7.12)	(5.22)	(9.26)	(9.56)	(7.09)
Wikipedia Page	72.99	52.29	68.14	86.30	45.13
Length	(0.29)	(0.40)	(0.30)	(0.39)	(0.29)

Diff is treated minus control within matched sample (weighted).
Robust t-stat in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

statistically detectable—are small in magnitude and economically limited. This pattern supports the interpretation that the post-launch changes documented in the main analysis reflect the impact of Grokipedia's introduction rather than pre-existing trends or imbalances.

Appendix E Average Treatment Effects on the Treated

Table A4 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.

Table A4: 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)	6239	6239	69048	69048	18738	18738	12303	12303	32661	32661

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

Across all occupations, the results show a clear and economically meaningful increase in Wikipedia activity following inclusion in Grokipedia, consistent with the difference-in-difference results in Table 2. Among academics and sportsmen, total post-period views rise by approximately 236, which represents an increase of roughly 20–25 percent relative to the typical pre-period weekly volume observed in these fields. The number of edits also increases, by about 0.05-0.09 edits over three weeks, and the estimate is statistically significant at the 1 percent level.

The effects are even larger for cultural, leadership, and political figures. For instance, 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.

Just like our difference-in-difference results in Section 6.1, these results speak directly to the theoretical framework developed in Section 4.

Appendix F Similarity and Wikipedia Page Views

In this appendix, we examine whether the semantic similarity between Grokipedia entries and pre-existing Wikipedia articles predicts changes in Wikipedia page views. This analysis parallels the mechanism regression for editing activity in Section 6.3, but replaces edits with page views as the outcome of interest.

Specifically, for each treated Wikipedia page i, we define the DiD estimates in views relative to its matched control page $j(i)$ as

$$\Delta\text{Views}_i^{\text{DiD}} \equiv (\text{Views}_i^{\text{post}} - \text{Views}_{j(i)}^{\text{post}}) - (\text{Views}_i^{\text{pre}} - \text{Views}_{j(i)}^{\text{pre}}).$$

We then estimate regressions of $\Delta\text{Views}_i^{\text{DiD}}$ on the similarity score between the Grokipedia entry and the pre-treatment Wikipedia article, progressively adding field indicators and page-level controls.

Table A5 reports the results. Across all specifications, the estimated coefficient on the similarity score is small in magnitude and statistically indistinguishable from zero. This finding contrasts with the corresponding regressions for editing activity, where similarity exhibits a robust and negative relationship with post-treatment edits.

The absence of a systematic relationship between similarity and page views suggests that, conditional on receiving a Grokipedia entry, Wikipedia traffic does not vary meaningfully with how closely the AI-generated entry resembles the original article. In the context of the model presented in Section 4, this result is consistent with the possibility that the substitution effect of AI entries on Wikipedia attention—captured by the parameter δ—is close to zero, or that opposing mechanisms affecting click-through behavior largely offset one another.

Importantly, this empirical finding does not affect the theoretical predictions in Section 4. Proposition 1 (Average Complementarity) requires only that the introduction of Grokipedia does not reduce attention to treated pages on average and that AI entries provide some degree of novel framing. Proposition 2 (Similarity Gradient Among Treated Pages) relies on the direct "inspiration" channel—whereby lower similarity increases editorial incentives—and does not depend on similarity affecting page views. Consequently, even if $\delta \approx 0$, the model continues to predict a positive average effect of Grokipedia on

Table A5: OLS Regression of Δ in Views between Treated and Control Pages

	(1)	(2)	(3)
Similarity Score	-13.59	-2.082	17.14
	(104.2)	(106.4)	(111.8)
Field: Culture		-35.66	-59.02
		(51.43)	(54.41)
Field: Leaders		-32.40	-34.17
		(56.89)	(60.33)
Field: Politics		-16.10	-31.44
		(60.58)	(63.52)
Field: Sports		-36.45	-49.71
		(54.21)	(56.96)
Wikipedia Word Count (10,000)			-13.65
			(13.23)
Gender: Male			-23.38
			(27.05)
Gender: Other			-84.30
			(344.0)
Age			-0.0337
			(0.0589)
Constant	102.2	124.5	157.8
	(88.56)	(101.4)	(109.9)
Observations	133241	133241	125598
R^2	0.000	0.000	0.000

Standard errors in parentheses

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

editing activity and a negative relationship between similarity and subsequent edits, both of which are borne out in the data.