The AI Democratization Paradox:

Large-scale Evidence from a Decentralized Knowledge Platform

Abstract

Does AI democratize knowledge production or amplify existing disparities? We investigate this tension by studying the deployment of neural machine translation across more than 100 Wikipedia language communities. Leveraging rich, fine-grained data and exogenous variation from a natural experiment, we uncover the "AI Democratization Paradox," where the technology simultaneously drives democratizing and concentrating forces. AI lowered barriers, leading to a substantial increase in content creation across diverse target languages without sacrificing quality or readership. However, the benefits were concentrated: well-resourced communities captured disproportionate gains—3-4 times larger than mid-tier editions. While editors actively leveraged AI to address representation gaps, translating female biographies at twice the expected rate, structural constraints limited the impact in high-need areas. We conclude that technological solutions alone cannot overcome structural inequalities; AI's distributional impact is contingent on the interplay between technological capabilities and existing social structures.

Keywords: artificial intelligence, machine translation, digital inequality, decentralized platforms, distributional effects, knowledge production, socio-technical systems

1 Introduction

Artificial intelligence (AI) is rapidly reshaping how knowledge is created, curated, and distributed globally (Goldfarb and Tucker, 2019; Agrawal et al., 2019, 2022). This transformation is particularly pronounced within decentralized knowledge platforms, ranging from Wikipedia's multilingual encyclopedias to GitHub's collaborative coding repositories. These platforms operate as instances of commons-based peer production (Benkler, 2002) within two-sided market structures (Rochet and Tirole, 2003; Parker et al., 2016), relying on distributed contributors—often volunteers—to build vast information resources that constitute a modern digital commons (Ostrom, 1990).

Within these unique ecosystems, AI tools—such as machine translation and large language models—promise to lower barriers to participation. However, their deployment raises a critical question that sits at the heart of contemporary debates on technology and inequality (Capraro et al., 2024): Does AI democratize knowledge production by leveling the playing field, or does it amplify existing inequalities by disproportionately benefiting already advantaged actors?

This debate centers on two competing perspectives. The *technological optimism* view posits that AI acts as a democratizing force by reducing technical barriers—such as the need for multilingual proficiency—thereby enabling broader participation (Jensen, 2007; Brynjolfsson et al., 2019, 2025). Conversely, the *structural reproduction* view argues that technology often reinforces existing disparities, as participants with greater complementary resources are better positioned to capture the benefits of new tools (Cohen and Levinthal, 1990; Forman et al., 2012). Understanding which force dominates, and under what conditions, is crucial for designing equitable digital platforms and anticipating AI's societal impact.

We investigate this tension by examining the impact of AI-powered machine translation (MT) on translation-mediated content creation across Wikipedia's global platform. Our study leverages a unique natural experiment: the integration of Google Translate's advanced neural MT system into Wikipedia's Content Translation tool in January 2019. This setting

provides an ideal laboratory. With over 100 language editions affected simultaneously, we can observe how the same technological shock interacts with diverse community structures, ranging from the massive English edition to smaller, resource-constrained languages, each governed by its own set of community-enforced norms and practices.

Employing a difference-in-differences strategy that exploits variation in the timing and coverage of Google Translate's rollout, we uncover what we term the *AI Democratization Paradox*. We find that advanced machine translation simultaneously unleashes both democratizing and concentrating forces.

On the one hand, the technology substantially democratizes content creation. We observe a remarkable 139% increase in translation-mediated content creation, generating an estimated 12.3 million additional pageviews over two years. Crucially, these gains are distributed across a diverse array of target languages, extending knowledge access to speakers of previously underserved linguistic communities worldwide. Beyond this quantitative expansion, editors actively leverage the efficient gains to address systemic representation gaps—female biographies, for instance, are translated at twice the rate expected based on existing content distributions.

On the other hand, concentrating forces significantly counteract these democratizing gains. The benefits flow disproportionately to well-resourced language communities, with the largest editions experiencing gains 3-4 times larger than mid-tier communities and over 10 times larger than the smallest editions. English Wikipedia's dominance as a source language intensifies from 68% to 81%, while regions with the greatest need—such as Sub-Saharan Africa—see minimal impact.

Our study reconciles these paradoxical outcomes by demonstrating that AI's impact is critically shaped by the *interaction between technological capabilities and existing social structures*. Machine translation operates through an "expertise unbundling" mechanism, decoupling content knowledge from the need for specific language skills. This lowers the barrier to entry. However, realizing the benefits of this unbundling requires complementary

resources. Our distributional analyses reveal that communities with greater editorial capacity, larger existing knowledge bases, and established reader networks are better equipped to leverage the technology, leading to the concentration of benefits. Technological capability alone cannot overcome structural inequalities.

This research makes two primary contributions. First, we quantify and contextualize the socio-economic significance of AI-powered translation in a major commons-based peer production environment. By dramatically accelerating translation efforts—compressing what would have taken decades of manual labor into just two years—AI tools facilitated an estimated 12.3 million additional pageviews during our study period, a substantial expansion in access to knowledge. Second, we move beyond "AI increases productivity" narratives by providing novel, large-scale empirical evidence of how socio-technical systems shape AI's distributional impact. By foregrounding the AI Democratization Paradox, we offer a conceptual framework that reconciles competing views on technology's role in digital inequality.

Finally, our findings offer actionable insights for platform governance and AI deployment strategy. We demonstrate that providing equal access to AI tools is necessary but insufficient to ensure equitable outcomes. To realize AI's democratizing potential, platforms must implement complementary interventions—such as targeted editorial support, capacity building, and incentive structures—that address the resource constraints of disadvantaged communities. Organizations deploying AI in decentralized systems should adopt a distributional lens, anticipating not just aggregate productivity gains but also where technology will alleviate versus amplify existing disparities.

The remainder of this paper is organized as follows: Section 2 reviews the relevant literature and develops our theoretical framework. Section 3 describes the empirical setting, and Section 4 details the data used in our study. Section 5 outlines our empirical strategy. Section 6 presents the aggregate impact of AI on content production and consumption, while Section 7 analyzes the distributional effects across language communities and content representation, revealing the AI Democratization Paradox. Finally, Section 8 concludes with a

discussion of the theoretical and managerial implications.

2 Theoretical Framework

This study addresses the intersection of AI adoption, digital platforms, and inequality. To articulate the central theoretical tension, we first define the economic and organizational context of our study—a decentralized knowledge platform—then examine competing perspectives on how AI deployment should affect such systems, and finally clarify our contribution.

2.1 Economic Organization of Decentralized Knowledge Platforms

The rapid advancement of AI is reshaping the digital economy (Agrawal et al., 2019, 2022). While much research focuses on AI adoption within firms, decentralized knowledge platforms present a distinct and theoretically important context. These platforms—including Wikipedia, GitHub, and Reddit—represent a third mode of economic organization, separate from traditional firms and markets, known as commons-based peer production (Benkler, 2002). In these systems, large numbers of geographically dispersed individuals, driven by a diverse set of intrinsic and social motivations rather than direct financial compensation, work cooperatively to generate vast repositories of public information goods.

From a structural perspective, these platforms function as two-sided markets (or multi-sided markets), creating value by facilitating interactions between distinct groups of users—in Wikipedia's case, content contributors (editors) and content consumers (readers) (Rochet and Tirole, 2003; Parker et al., 2016). The platform's operator (the Wikimedia Foundation) does not produce the content but provides the infrastructure and governance framework that enables the community to thrive (Teh, 2022).

The governance of these platforms is also unique. Rather than relying on managerial hierarchy (as in a firm) or explicit contracts (as in a market), they are often self-governed by

complex systems of community-developed rules, norms, and enforcement mechanisms. This model of collective action to manage a shared resource closely mirrors the principles for successful commons governance identified by Ostrom (1990). This socio-technical environment, characterized by volunteer labor, non-market incentives, and community-based governance, creates complex and unpredictable dynamics (Zhang and Zhu, 2011; Faraj et al., 2016; Zhu et al., 2020). The deployment of AI in such a context raises a fundamental question: Will it empower the collective and level the playing field, or will it exacerbate the structural inequalities already present in the system?

2.2 Democratizing Potential of AI: A Technological Optimism Perspective

The technological optimism perspective suggests that AI can act as a democratizing force. By reducing technical barriers and transaction costs—such as the need for multilingual proficiency—AI enables broader participation. This view, articulated in the theory of democratizing innovation, posits that as the costs of design and production tools fall, the locus of innovation can shift from centralized producers to individual users, who can now create what they need for themselves (Von Hippel, 2006).

In multilingual environments like Wikipedia, language barriers represent a significant impediment to knowledge sharing, creating what scholars term the "digital language divide". Machine translation, powered by recent advances in neural networks (Wu et al., 2016; Vaswani et al., 2017), directly addresses this constraint by lowering the transaction costs of cross-linguistic communication. Prior research provides strong causal evidence for this effect, showing that the introduction of high-quality machine translation on an e-commerce platform significantly increased international trade by making it easier for buyers and sellers to understand each other (Brynjolfsson et al., 2019).

In the context of knowledge production, this leads to what we term an expertise unbundling mechanism. Traditionally, creating a Wikipedia article in a new language required

both domain expertise (knowledge of the topic) and specific production skills (fluency in the target language). Advanced machine translation decouples these requirements. It allows a contributor with deep knowledge of a topic in one language to create a high-quality draft in another language without being fluent in it, leaving the final refinement to a native-speaking editor. Theoretically, this unbundling should flatten hierarchies and enable wider participation, particularly from under-resourced language communities seeking to import knowledge from larger ones. This process should, in theory, accelerate the growth of smaller encyclopedias by tapping into the cumulative growth dynamics observed on the platform, where a larger base of existing content attracts even more contributions (Aaltonen and Seiler, 2016).

2.3 Concentrating Effects of AI: A Structural Reproduction Perspective

Conversely, a substantial body of literature on socio-technical systems and digital inequality cautions that technology adoption often reinforces, rather than reduces, existing disparities. This structural reproduction view argues that the benefits of new technologies are not uniformly distributed but are instead mediated by existing social structures and the availability of complementary resources (Wessels, 2013). The features of a technology do not operate in a vacuum; their impact depends on how users can leverage them within specific social contexts.

This perspective suggests that communities already possessing advantages—such as a large base of active editors, a substantial existing knowledge repository to translate from, and a large, engaged readership—are better positioned to adopt and exploit new AI tools. For example, a language community with many active editors can more effectively review, correct, and integrate the surge of machine-translated drafts, while a community with few editors may be overwhelmed. This differential capacity can lead to a concentration of benefits among the already advantaged, a dynamic that parallels findings from labor economics suggesting AI's gains will flow disproportionately to high-skill workers who can complement the technology

(Cazzaniga et al., 2024).

Indeed, extensive research has documented persistent structural inequalities within Wikipedia, despite its open and collaborative mission. Significant gaps exist in content representation across gender, geography, and language dimensions (Graham et al., 2014; Wagner et al., 2016; Miquel-Ribé and Laniado, 2020). From a structural reproduction viewpoint, without carefully designed interventions, AI deployment risks exacerbating these gaps by primarily accelerating content creation in the very communities that are already dominant. This aligns with a broader historical perspective which posits that power relations and institutional capacity, not technology alone, are the ultimate determinants of how the gains from innovation are distributed across society (Cohen and Levinthal, 1990; Scheidel, 2024).

2.4 AI Democratization Paradox and Our Contribution

These competing perspectives—technological optimism and structural reproduction—create a critical theoretical tension regarding AI's impact on decentralized knowledge production. Our contribution is the articulation and empirical validation of the AI Democratization Paradox: the hypothesis that a single technological intervention can simultaneously drive democratization (by unbundling expertise and reducing barriers for all) and concentration (by amplifying the advantages conferred by complementary resources).

This framework helps explain why simply providing equal access to a powerful tool may be insufficient to address digital inequalities. It highlights the critical interplay between technological capabilities and pre-existing social structures in shaping the future of digital knowledge ecosystems. Our research, therefore, moves beyond asking if AI increases productivity to ask for whom it increases productivity and under what structural conditions.

3 Setting

Wikipedia serves as an ideal empirical setting to examine the interaction between AI deployment and decentralized knowledge production. As the largest collaborative encyclopedia in the world, Wikipedia operates through a volunteer-driven model across more than 300 language editions, with each language community possessing vastly different levels of resources and organizational capacity. This heterogeneity enables us to study how a uniform technological shock—AI-powered translation—interacts with diverse existing social structures.

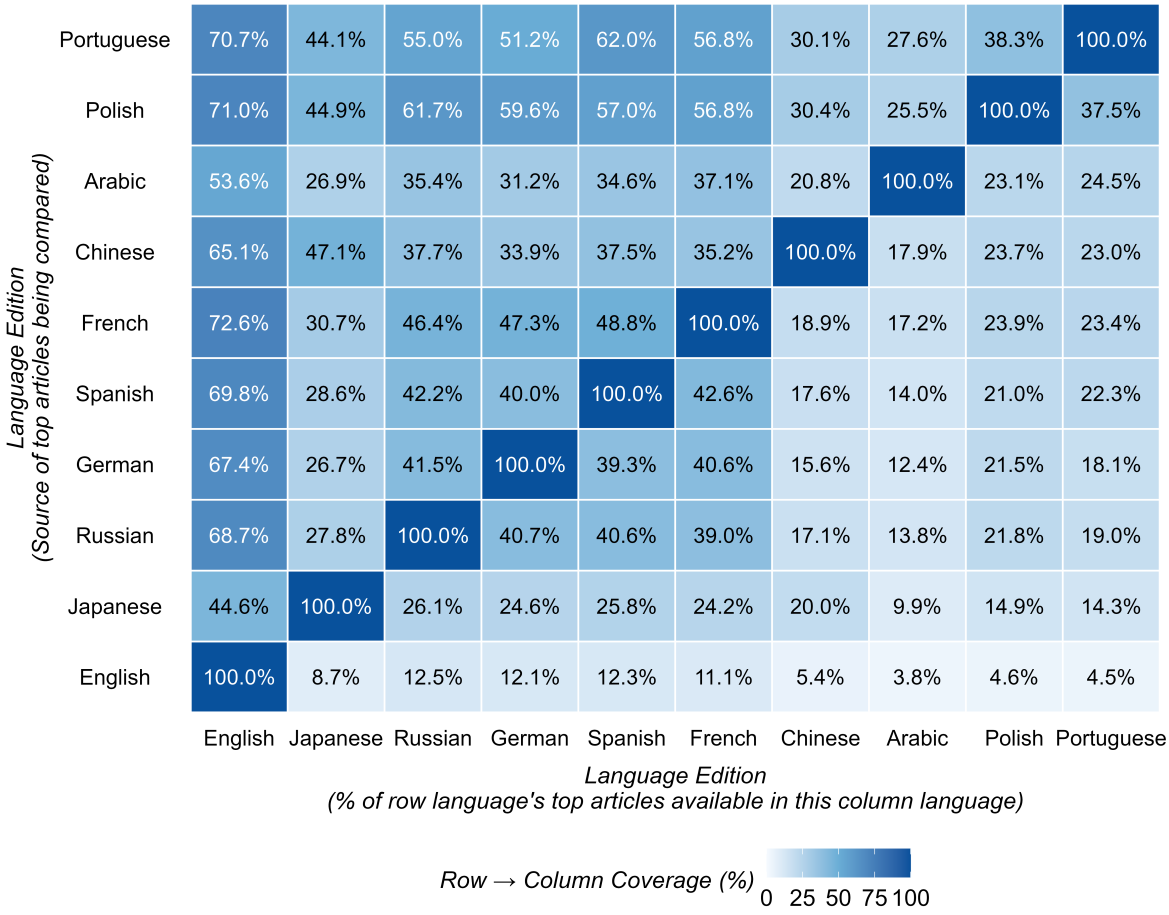
3.1 The Challenge: Cross-Linguistic Knowledge Gaps

Despite Wikipedia's global reach, knowledge availability varies sharply across languages. Content overlap between language editions is limited, meaning that readers of one language often lack access to information that are widely consumed in another.

Figure 1 quantifies these differences using the top 5% most-viewed articles in 2024 for the ten largest language editions. Each cell reports *directional coverage*: the share of high-traffic articles in the row language that also exist in the column language. Two patterns stand out. First, coverage is generally low—pairwise overlap rarely exceeds 70% even among major languages. Second, coverage is asymmetric: for example, more than 60–70% of Portuguese or Spanish high-traffic articles appear in English, whereas only about 5–13% of English high-traffic articles appear in most other editions (see the bottom row). These patterns indicate substantial cross-linguistic gaps and suggest that much of Wikipedia's most demanded knowledge remains siloed by language.

Historically, creating an article in a new language required writing from scratch—a time-intensive process demanding both domain expertise and proficiency in the target language. These high fixed costs advantaged well-resourced editions and left smaller ones behind. To alleviate this constraint, the Wikimedia Foundation introduced the Content Translation system in 2015, which enables *translation-mediated content creation* and lowers the barriers

Figure 1: Large and Asymmetric Cross-Language Gaps in High-Traffic Wikipedia content.



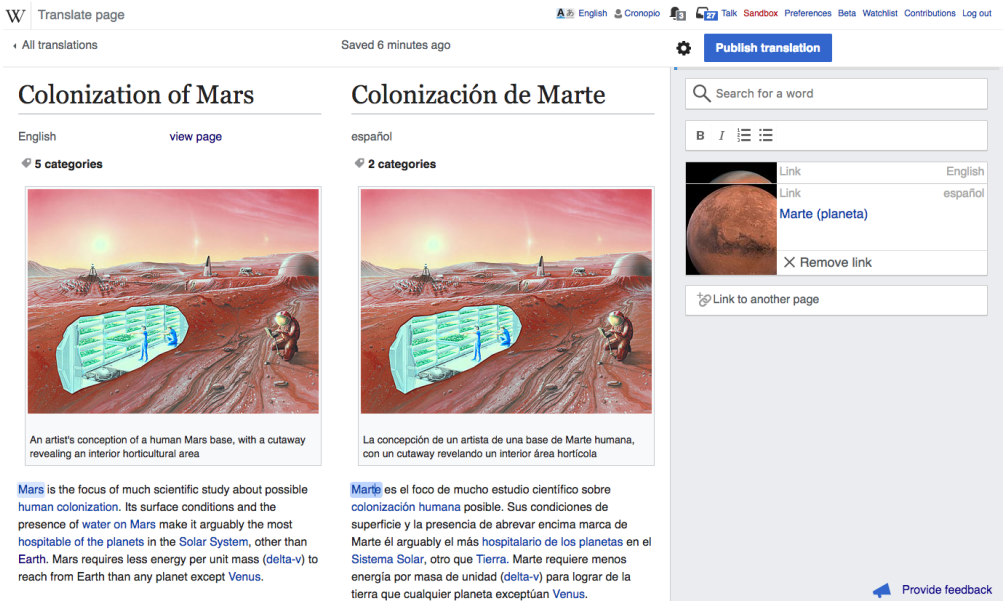
Notes: Each cell shows the share (%) of the row language's top 5% most-viewed 2024 articles that are also available in the column language (*row* → *column* coverage). Diagonal cells equal 100% by construction. Pairwise coverage is generally low and asymmetric: coverage seldom exceeds 70%, and English in particular exhibits low outward coverage (roughly 5-13% into other editions) while covering a large share of other language's high-traffic content.

to systematic knowledge transfer across language editions.

3.2 The Human-AI Collaboration Workflow

The Content Translation system implements a human-in-the-loop workflow that combines machine translation capabilities with human editorial oversight. When editors select an article for translation, the system presents a dual-panel interface (Figure 2): the source article appears on one side, while a workspace for creating the translated version appears on the other. Importantly, the system automates technical aspects—preserving formatting, maintaining link structures, and transferring citations—allowing editors to focus on ensuring linguistic accuracy and cultural adaptation.

Figure 2: The Human-AI Collaboration Workflow in Wikipedia's Content Translation Tool



Note: Screenshot of the Content Translation tool interface. The system employs a human-in-the-loop workflow, presenting the source article (left panel) alongside a workspace for the translated version (right panel). It generates machine translations as initial drafts, which editors then review, revise, and refine before publication. The system also automates technical aspects like formatting and link preservation.

The system generates machine translations as initial drafts, which editors then review, revise, and refine before publication. This structured human-AI collaboration is essential for maintaining Wikipedia's quality standards while leveraging the efficiency gains from automation. The human-in-the-loop design distinguishes this setting from fully automated

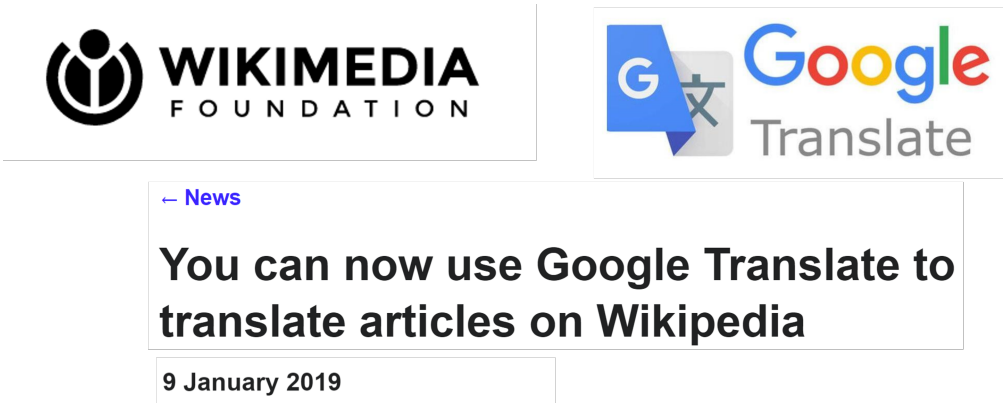
translation applications.

By late 2024, this system had facilitated the creation of over 2 million articles—a volume equivalent to creating three medium-sized Wikipedia language editions or roughly one-third the size of English Wikipedia. This substantial output demonstrates the system's transformative impact on cross-linguistic knowledge transfer.

3.3 The Technological Shock: Google Translate Integration

The Content Translation system initially relied on basic machine translation engines like Apertium and Yandex, which offered limited accuracy and language coverage. The critical intervention we study occurred in January 2019, when the Wikimedia Foundation integrated Google Translate's advanced neural machine translation system (Figure 3).[1]

Figure 3: The Technological Shock: Integration of Google Translate's Neural Machine Translation System



Note: Screenshot of the Wikimedia Foundation's announcement on January 9, 2019, regarding the integration of Google Translate into the Content Translation tool. This integration introduced advanced Neural Machine Translation (NMT) capabilities, representing a substantial, exogenous technological shock by significantly improving translation quality and lowering the cognitive burden for editors.

This integration represented a substantial technological shock to the Wikipedia ecosystem. Google Translate employs Neural Machine Translation (NMT) architecture, which

[1]See the official announcement: https://wikimediafoundation.org/news/2019/01/09/you-can-now-use-google-translate-to-translate-articles-on-wikipedia/

processes entire sentences holistically rather than word-by-word, capturing context and linguistic nuance far more effectively than previous statistical methods. Published evaluations demonstrated that this technology reduced translation errors by 55-85% across major language pairs (Wu et al., 2016), substantially lowering the cognitive burden for editors engaged in translation-mediated content creation.

This empirical setting offers several distinct advantages for our research. First, the integration provides a clear, plausibly exogenous technological shock affecting over 100 language communities simultaneously. Second, Wikipedia's open architecture provides comprehensive data on content creation patterns, quality metrics, and user engagement. Finally, the platform's decentralized, volunteer-driven nature enables us to study how AI capabilities interact with existing community structures, offering key insights into the conditions under which technology democratizes or concentrates knowledge production.

4 Data

4.1 Data Sources and Sample Construction

Our empirical analysis draws on rich data from Wikipedia's decentralized ecosystem, enabling a detailed examination of both the aggregate and distributional effects of AI deployment on translation-mediated content creation.

Primary Data: Translation Records. Our core dataset comprises complete translation records from Wikipedia's Content Translation System, which logs every article translation facilitated through the tool. Each record contains detailed metadata including source and target languages, article titles, publication timestamps, and the specific machine translation engine used (see Section A.1 for technical details). Our sample spans January 2017 to December 2020, providing two years of observations before and after the primary integration of Google Translate in January 2019.

Complementary Data Sources. We enrich our primary dataset with multidimensional

metrics at the article, language, and editor levels to capture the full scope of the intervention's impact.

Article-Level Metrics: To measure reader engagement, we use pageview statistics. To assess content quality, we employ two complementary signals: the Wikimedia Foundation's language-agnostic structural quality scores and community deletion logs that track articles removed for failing to meet standards. To analyze representation patterns, we leverage Wikidata's structured metadata providing demographic information (gender for biographical articles, geographic location for place-based content) and employ language-agnostic topic classifications (Johnson et al., 2021) to systematically examine content gaps across domains.

Language-Level Characteristics: To understand how existing social structures mediate AI's impact, we compile metrics capturing each language community's complementary resources. These include: (1) knowledge base size, measured by total existing articles; (2) editorial capacity, captured through active editor counts; (3) readership network size, indicated by total in the language edition.

Sample Construction We applied several filtering criteria to construct our analysis sample. We focus our analysis on target languages, including only those with at least 20 translations during the sample period. We exclude editions exhibiting abnormal translation spikes suggestive of bot activity or coordinated campaigns,[2] and remove English Wikipedia as a translation destination due to its unique content policies restricting translations from other languages.

This process yields a substantial dataset of 540,488 translated articles—a volume equivalent to the entire content of a small to medium-sized Wikipedia language edition. For panel analyses, we aggregated these articles to the language-month level, creating a balanced panel dataset of 132 languages observed over 48 months from January 2017 to December 2020.

[2]Defined as monthly volumes exceeding 20 times a language's historical average.

4.2 Pre-Intervention Patterns and Structural Asymmetries

The baseline distribution of translation activity reveals structural asymmetries that existed before the AI intervention, providing key context for interpreting our subsequent findings. These pre-existing patterns highlight a centralized knowledge-sharing network characterized by structural asymmetries.

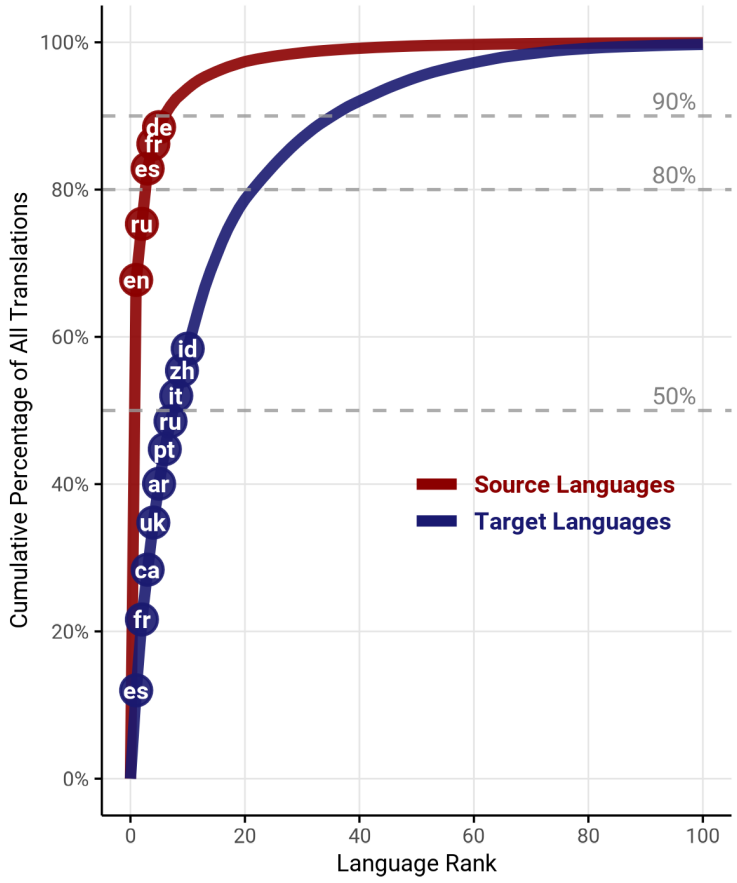
Figure 4 illustrates the stark concentration of translation activity. English Wikipedia dominates as the information source for 68% of all translations, establishing its role as the primary global knowledge hub. The top five source languages—English, Russian, Spanish, French, and German—collectively account for nearly 90% of translated content. While target languages show greater diversity, concentration remains significant: the top 10 recipient languages receive 58% of all translations, rising to 95% for the top 50 languages. These patterns reveal that, while translation extends knowledge beyond its original sources, the benefits of knowledge transfer are heavily concentrated in a relatively small subset of Wikipedia's 300+ language editions.

This concentration masks a *dual inequality* in how different communities depend on translation, as revealed in Figure 5. Panel (a) shows that large, well-resourced communities dominate in absolute volume—Spanish leads with approximately 45,000 translated articles, followed by French (35,000) and Catalan (30,000). These communities possess the editorial capacity to absorb and integrate substantial translated content.

However, Panel (b) reveals a strikingly different pattern when examining relative dependency. Smaller, resource-constrained communities rely most heavily on translation as a proportion of their total content. In language editions such as Lao, Punjabi, and Hakka Chinese, translated articles constitute 15–20% of the entire encyclopedia—suggesting that, for these smaller editions, translation functions not merely as a supplement but as a critical engine of growth.

These baseline patterns establish the asymmetric structure that predated AI intervention: large language editions dominate absolute translation volumes while smaller editions depend

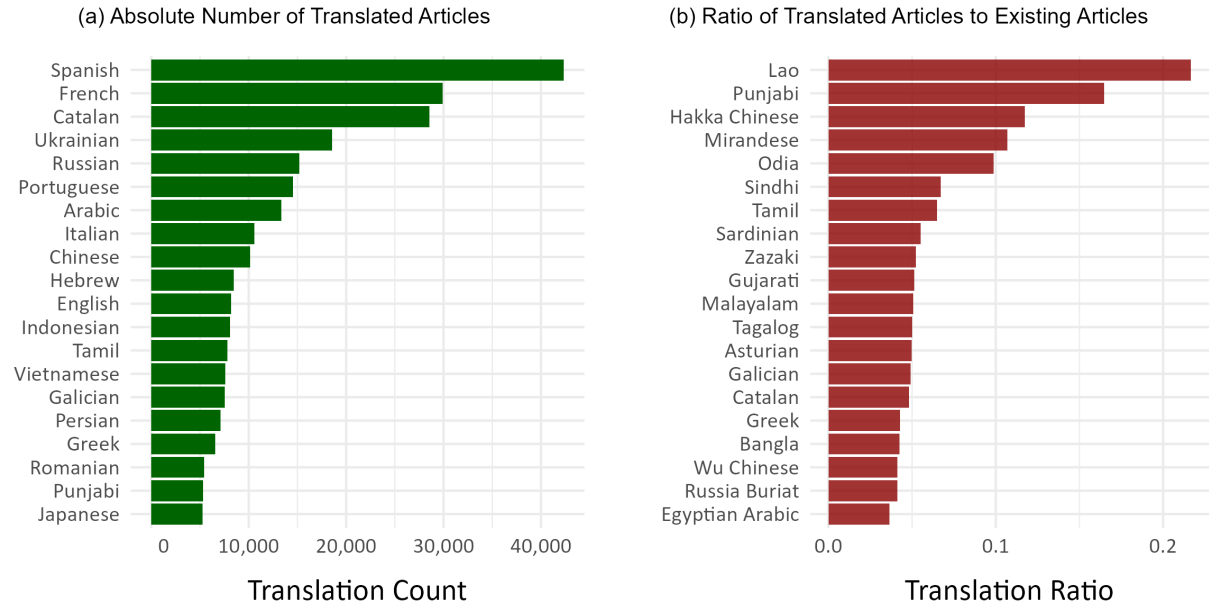
Figure 4: Pre-Intervention Centralization of the Global Knowledge-Sharing Network (2017-2018)



Note: Cumulative distribution of translation activity across source and target languages before the Google Translate intervention. The source languages (red) are highly concentrated: English dominates as the source for 68% of all translations, and the top five languages account for nearly 90%. Target languages (blue) show more diversity, yet the top 10 recipient languages still receive 58% of all translations.

more critically on translation for their development. This dual inequality provides essential context for understanding how a technological shock like Google Translate deployment interacts with existing disparities to produce the democratization paradox we identify.

Figure 5: The Dual Inequality in Pre-Intervention Translation Dependency (2017-2018)



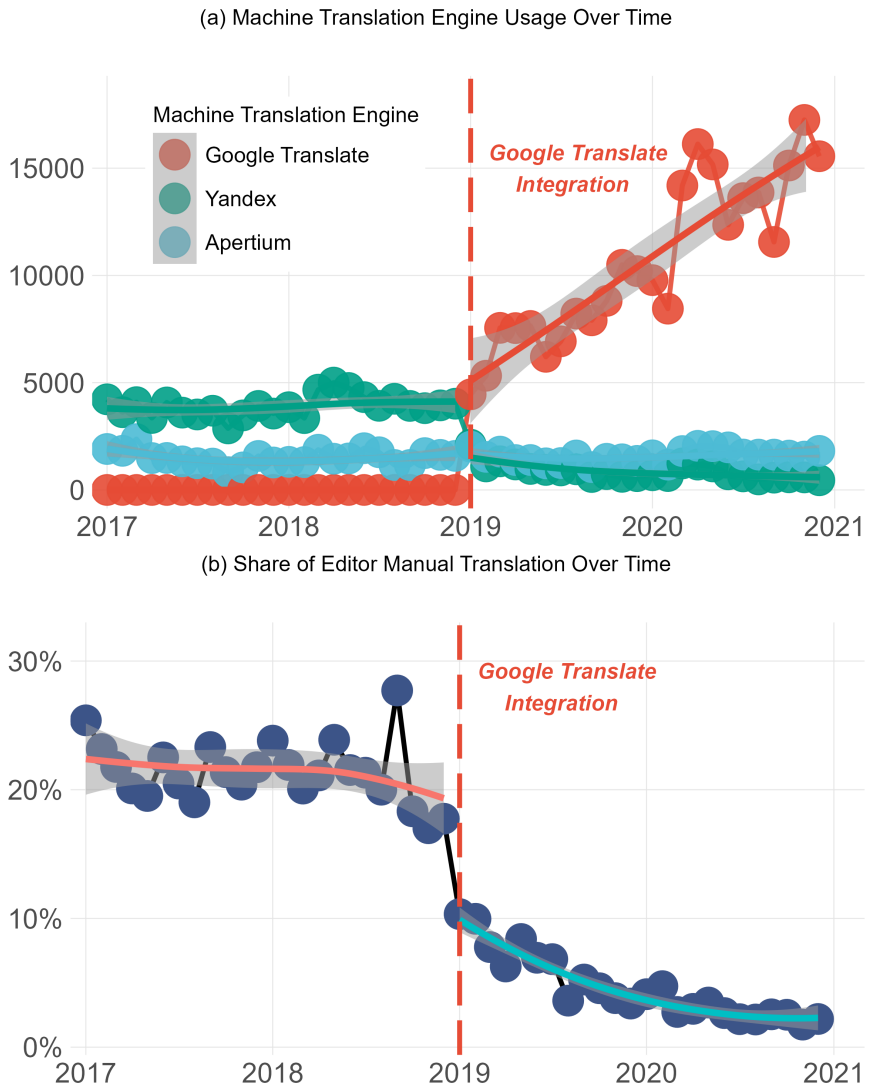
Note: This figure illustrates the structural asymmetries in translation reliance across Wikipedia language editions before the Google Translate intervention. Panel (a) shows the absolute volume of translated articles by target language, indicating that large, well-resourced communities (e.g., Spanish, French) dominate in total volume. Panel (b) displays the ratio of translated articles to the total existing articles for each language, revealing that smaller, resource-constrained communities (e.g., Lao, Punjabi) rely much more heavily on translation (15–20% of total content) for their overall growth.

4.3 Descriptive Evidence: Technology Adoption and Workflow Shift

Initial descriptive analysis reveals dramatic behavioral shifts following Google Translate's integration, providing preliminary evidence of the technology's transformative impact. Figure 6 documents these changes from two complementary perspectives.

Panel (a) tracks the usage patterns of different machine translation engines over time. Before January 2019 (marked by the vertical line), Yandex and Apertium maintained stable but modest usage levels. Following integration, Google Translate exhibits rapid adoption, simultaneously displacing existing engines and dramatically expanding overall machine-assisted translation volume. This pattern suggests that Google Translate did not merely substitute for existing tools but fundamentally expanded the scope of AI-assisted content creation.

Figure 6: Rapid Adoption of Advanced Machine Translation and the Transformation of Translation Workflows (2017-2020)



Note: Behavioral shifts following the integration of Google Translate (marked by the vertical red line in January 2019). Panel (a) tracks the monthly usage of different machine translation engines, showing the rapid adoption of Google Translate, which both displaced existing engines (Yandex, Apertium) and expanded overall volume. Panel (b) displays the proportion of articles translated manually (without machine assistance), which dropped sharply from approximately 20% to single digits, indicating a near-complete transition to AI-assisted workflows.

Panel (b) reveals the corresponding transformation in editor workflows. The proportion of manual translations—articles created without any machine assistance—dropped precipitously from approximately 20% before integration to single digits afterward. This near-

complete transition to AI-assisted workflows indicates that the improved technology became the de facto standard for cross-linguistic content creation.

Together, these patterns provide suggestive evidence of a fundamental transformation in content creation processes, motivating our formal causal analysis of the technology's aggregate and distributional effects.

5 Empirical Strategy

We exploit Google Translate's integration into Wikipedia's Content Translation system as a quasi-experimental shock to identify the causal impact of advanced AI on knowledge production. Our identification strategy leverages variation in both the timing and coverage of this technological deployment.

5.1 Identification Strategy and Treatment Definition

The integration of Google Translate provides a plausibly exogenous source of variation for causal identification. The deployment resulted from a top-down partnership between the Wikimedia Foundation and Google, making the timing independent of individual language communities' activities or characteristics. Moreover, treatment assignment was determined by Google's pre-existing language support rather than Wikipedia-specific factors, creating quasi-random variation in AI availability.

Within our sample of 132 Wikipedia language editions, the deployment occurred in waves (Table 1). The primary wave in January 2019 simultaneously treated 78 languages, followed by staggered additions of 12 languages through August 2020. The remaining 42 languages, unsupported by Google Translate during our sample period, serve as controls. This staggered structure, while creating some complexity, provides rich variation for identification, with the vast majority of treated languages (78 of 90) receiving simultaneous access.

This implementation pattern enables a difference-in-differences research design with our

language-month panel data. Our identification relies on the parallel trends assumption: absent the intervention, treated and control languages would have followed similar trajectories in translation activity. We validate this assumption through multiple pre-trend tests and placebo analyses presented below.

Table 1: Staggered Rollout of Google Translate Across Wikipedia Language Editions

Implementation Status	**Deployment Period**	**Languages**
Completed Implementation (Treated)		
Initial Major Rollout	January 2019	78
Staged Rollout	June 2019 – August 2020	12
Pending Implementation (Control)	—	42
Total in Sample		**132**

Note: Summary of the Google Translate deployment timeline across the 132 languages in our analysis sample. The deployment occurred in waves, providing variation in timing and coverage used for the difference-in-differences identification strategy. The 42 languages unsupported during the sample period serve as controls.

5.2 Primary Estimation: Counterfactual Imputation Approach

Given the staggered rollout and potential for heterogeneous treatment effects, we employ contemporary difference-in-differences methods that address recent econometric concerns about traditional two-way fixed effects (TWFE) estimators. Our primary approach uses counterfactual imputation, with robustness checks using multiple alternative estimators including Stacked DiD (Cengiz et al., 2019), Interaction Weighted estimation (Sun and Abraham, 2021), and Doubly Robust methods (Callaway and Sant'Anna, 2021). All approaches yield consistent results (detailed in Appendix C).

We adopt counterfactual imputation as our primary method for three reasons. First,

causal inference fundamentally addresses a missing data problem (Holland, 1986)—we observe treated units' outcomes under treatment but not their counterfactual outcomes without treatment. Directly imputing these counterfactual outcomes provides a natural and transparent framework for estimating treatment effects. Second, this approach accommodates both staggered rollout timing and treatment effect heterogeneity across units and time periods. Third, it avoids the negative weighting problems that have been highlighted in recent applied econometrics literature regarding traditional TWFE models (De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021).

Specifically, we implement the matrix completion estimator within the Counterfactual Imputation framework (Athey et al., 2021).[3] The matrix completion method constructs a panel data matrix with language editions as rows, time periods as columns, and translation volumes as entries. The method leverages patterns across both dimensions—using information from control units and pre-treatment periods—to impute counterfactual outcomes for treated observations. This approach is particularly well-suited to our setting as it flexibly accommodates the heterogeneous responses we expect across Wikipedia's diverse language communities.

5.3 Alternative Identification and Estimation

To strengthen causal inference and address potential concerns, we implement two complementary identification strategies that rely on different sources of variation and identifying assumptions.

Time-Shift Difference-in-Differences Design. As an alternative approach, we implement a time-shift difference-in-differences design focusing exclusively on the 78 languages treated in January 2019. We compare their translation activity in July 2018–June 2019 (post-treatment) against July 2017–June 2018 (pre-treatment), using the same languages'

[3]We verify robustness across multiple imputation methods provided the *fect* R package, including fixed effects counterfactual, interactive fixed effects, and matrix completion approaches, all yielding consistent results.

earlier activities as the counterfactual.

This within-language comparison alleviates concerns about unobserved heterogeneity affecting growth trajectories differently between different language groups. The limitation of this approach is the necessarily shorter time window, as we must reserve the earlier period as a control.

Blocked Difference-in-Differences Design. For our heterogeneity analyses examining how effects vary across language characteristics and content types, we employ a simplified blocked difference-in-differences design. This approach focuses on the 78 languages treated simultaneously in January 2019, excluding those with staggered timing. This creates a clean setting with single treatment timing, enabling transparent estimation using classical two-way fixed effects:

$$Y_{it} = \alpha \times GT_{it} + \beta \times GT_{it} \times X_i + \gamma_i + \delta_t + \epsilon_{it}$$

where Y_{it} represents translation volume for language i in month t, GT_{it} indicates Google Translate availability (1 for treated languages post-January 2019, 0 otherwise), X_i captures time-invariant language characteristics (e.g., editorial capacity, knowledge base size), and γ_i and δ_t are language and time fixed effects. The coefficient vector β identifies heterogeneous treatment effects—our primary interest in understanding the distributional impact of AI deployment.

This multi-pronged empirical strategy, combining different estimators and identification approaches, ensures our findings are robust to various econometric consideration and provides convinciing evidence on both the aggregate and distributional effects of AI-powered translation.

6 Aggregate Impact: The Democratizing Potential of AI

We begin by examining the aggregate impact of Google Translate integration on translation-mediated content creation. This analysis tests the technological optimism perspective: whether advanced AI successfully lowers barriers to knowledge production through expertise unbundling, and whether this increased output maintains quality standards and reader value. We employ our counterfactual imputation framework for language-level analysis and comparative analysis over time for article-level quality metrics.

6.1 The Productivity Surge

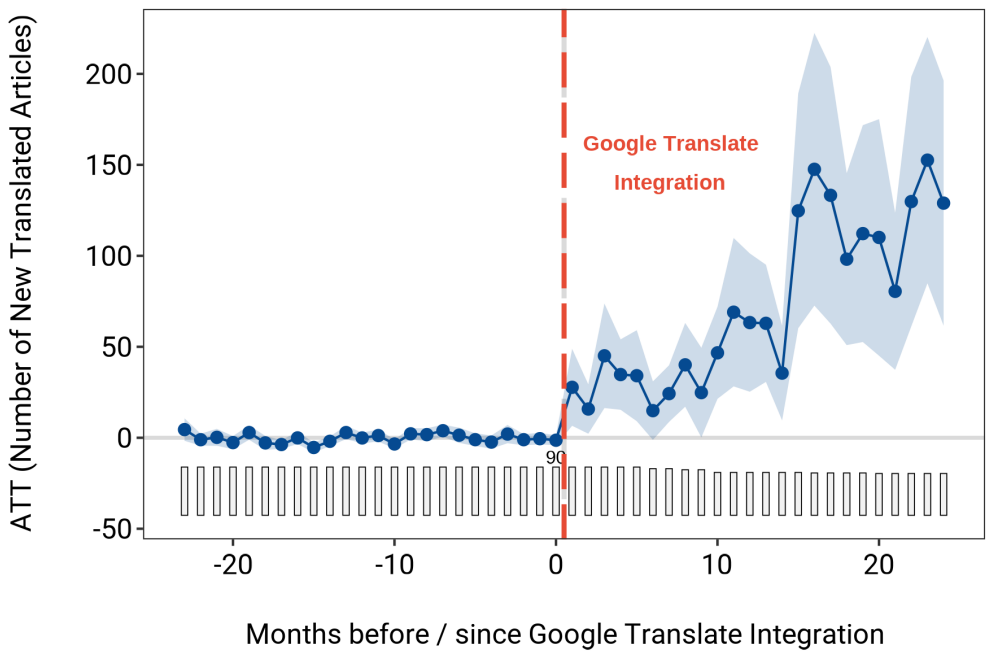
We first analyze the impact of Google Translate integration on the volume of translation-mediated content creation. This analysis tests whether advanced AI effectively decouples language skills from content knowledge, thereby increasing output.

Main Results. The event study results, estimated using the Counterfactual Imputation method, are presented in Figure 7. We observe parallel trends prior to the intervention, followed by an immediate and sustained increase in translation activity post-integration.

Our difference-in-differences estimates show that Google Translate's integration increased translation-mediated article creation by 71.2 articles per language edition per month (ATT; $p<0.01$). In the steady-state phase—the final six months of the post-integration period—the monthly increase rises to 119 translated articles, a 139% gain over the pre-integration average of 85.4 articles per month. Together, these magnitudes highlight the transformative "expertise unbundling" mechanism: by sharply lowering the linguistic skill required for translation, machine translation democratized cross-linguistic content creation.

Diagnostic Tests. In line with best practices in difference-in-differences analysis, we conduct two diagnostic tests (Figure 8) to strengthen our causal inference. First, an equivalence test confirms the absence of differential pre-trends ($p=0.863$), validating our parallel trends assumption. Second, a placebo test artificially shifting the intervention six months

Figure 7: The Technology-Driven Productivity Surge: A Significant Increase in Translation-Mediated Content Creation

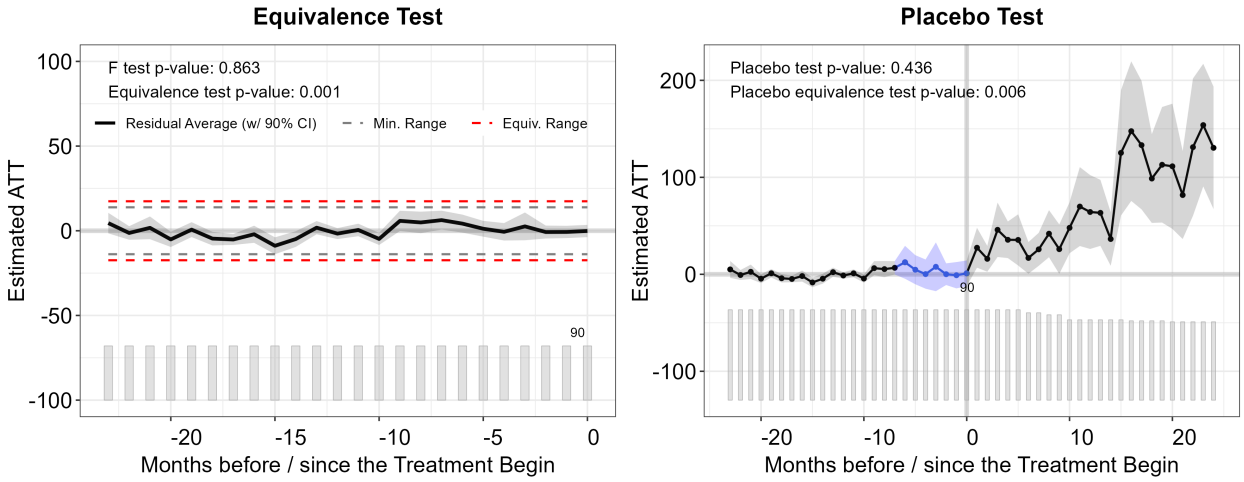


Note: Event study plot showing the impact of Google Translate integration (t=0) on monthly translation volume, estimated using the Counterfactual Imputation method. The plot shows parallel trends prior to the intervention, followed by an immediate and sustained increase in activity. The estimated Average Treatment Effect on the Treated (ATT) is a increase of 71.2 additional articles per language/month from pre-integration levels. The blue line represents point estimates, and the shaded area indicates 95% confidence intervals.

earlier to August 2018 yields no significant effect, confirming that our model correctly identifies null effects when no true intervention has occurred. Together, these diagnostics provide strong support for our causal interpretation of the observed increases in translation activity post-integration.

Further Robustness. Our findings are robust across a range of contemporary difference-in-differences estimators, including Callaway and Sant'Anna's Doubly Robust estimator, Sun and Abraham's Interaction-Weighted approach, Cengiz et al.'s Stacked DID design, and traditional two-way fixed effects models. As an additional validation, our within-language time-shift design—which using the same languages' earlier activities as the controls—produces highly consistent results (ATT = 79.1 articles/month). Together, this robustness across

Figure 8: Validation of the Causal Identification Strategy: Pre-trend and Placebo Tests



Note: Diagnostic tests validating the difference-in-differences strategy. The left panel shows the equivalence test, which confirms the absence of differential pre-trends (p=0.863), supporting the parallel trends assumption. The right panel presents the placebo test, which artificially shifts the intervention six months earlier (August 2018; blue region) and correctly identifies no significant effect.

methodologies with varying identifying assumptions, weighting schemes, and control group definitions strengthens confidence in the causal validity of our estimates. Detailed results are reported in Appendix C.

6.2 Maintaining Quality at Scale

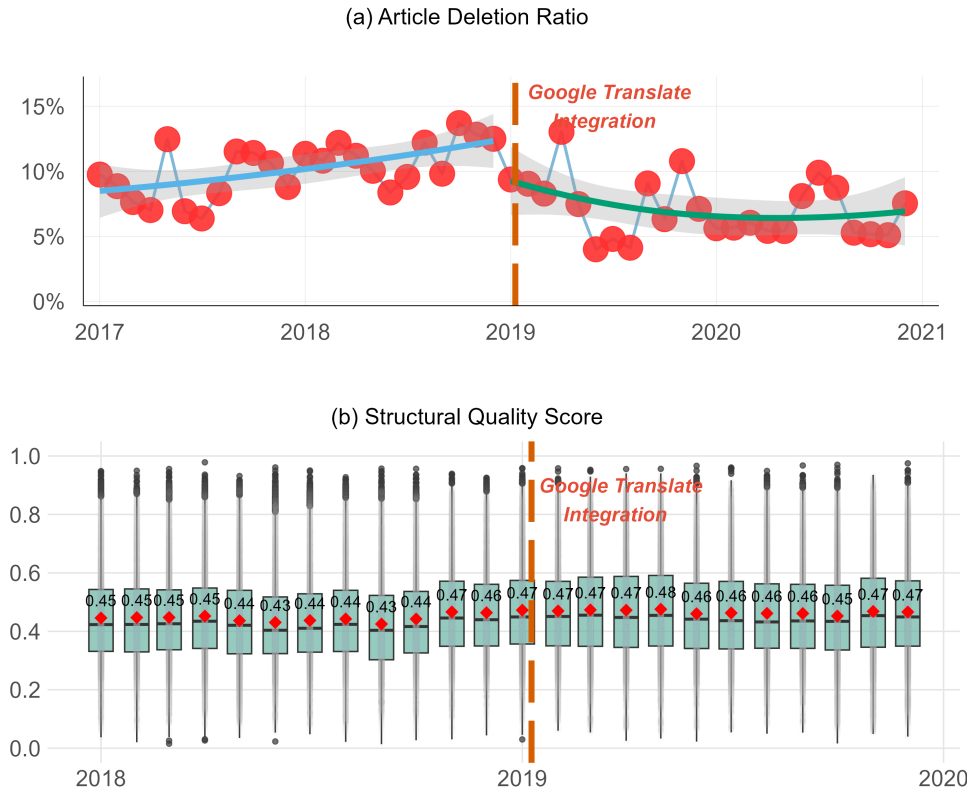
A critical concern with AI-assisted content creation is whether rapid scaling compromises quality. We address this through two complementary quality assessments: community validation through deletion processes and algorithmic structural quality evaluation.

Community Quality Standards. Wikipedia employs a collaborative deletion process to remove articles that fail to meet quality standards. The deletion ratio serves as a key indicator of community validation and represents the primary quality signal monitored by the Wikimedia Foundation in managing the Content Translation system.[4]

[4]We calculate the monthly deletion ratio by aggregating across all languages: (Number of articles translated in month M that were eventually deleted) / (Total number of articles translated in month M). Ag-

Panel (a) of Figure 9 reveals a surprising finding: deletion rates actually *declined* following Google Translate integration. Rather than degrading quality through automation, the improved translation technology, operating within Wikipedia's human-in-the-loop framework, actually enhanced editors' ability to create content that meets community quality standards.

Figure 9: Content Quality Was Maintained Amidst the Technology-Driven Productivity Surge



Note: Analysis of content quality before and after Google Translate integration (January 2019). Panel (a) tracks the community validation metric (deletion ratio: proportion of translated articles subsequently deleted). The LOESS curves show that deletion rates declined post-integration. Panel (b) presents the distribution of algorithmic structural quality scores (based on features like reference density and section organization). The improvement in median scores indicates that structural quality was maintained or slightly enhanced following Google Translate deployment (t=0).

Objective Quality Assessment. We complement community-based validation with the Wikimedia Foundation's language-agnostic quality prediction model, which evaluates articles based on structural features such as reference density, section organization, article

gregating results across all languages provides a stable monthly metric, as smaller language editions display high variability due to low monthly translation volumes.

length, and multimedia elements. This algorithmic assessment generates a structural quality score for each Wikipedia article across different languages.

Panel (b) of Figure 9 shows the distribution of structural quality scores for articles created before and after AI integration. We observe a modest but statistically significant improvement in structural quality following integration, with an average increase of 0.02 units ($p < 0.01$; detailed results in Section D.1). This suggests that AI assistance contributed to the creation of better-structured and documented content.

Together, these complementary quality measures demonstrate that the dramatic surge in AI-assisted content creation did not come at the expense of content quality. Both community validation through deletion processes and objective structural assessments confirm that advanced machine translation, when integrated within collaborative human-AI workflows, maintained content standards while enabling production at substantially greater scale.

6.3 Sustained Reader Engagement

Beyond production metrics, we next examine the downstream impact on knowledge consumption to assess whether the surge in AI-assisted content creation translated into meaningful reader engagement and knowledge access.

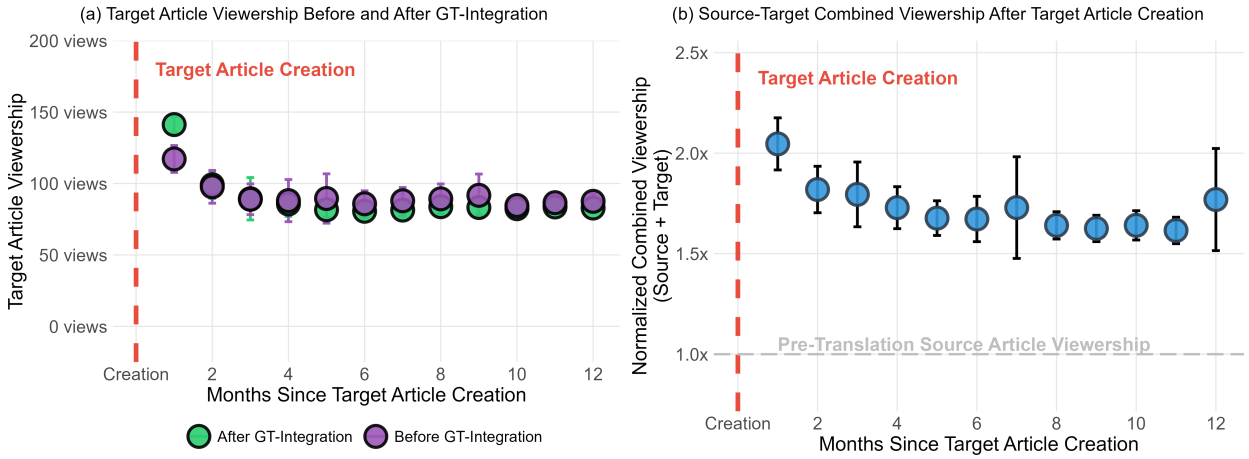
Content Utility. We analyze article pageviews to assess the utility of translated content from the readers' perspective. Figure 10 Panel (a) tracks pageview trajectories for articles created before and after Google Translate integration. Both cohorts exhibit remarkably consistent patterns: an initial spike of 140-150 views in the creation month, stabilizing at 80-90 monthly views thereafter.

Critically, we observe *no significant difference* in readership patterns between articles created before and after Google Translate integration. This stability represents a key finding: despite the remarkable surge in translation volume, the per-article engagement value remained constant. This pattern contradicts concerns that AI-assisted content might be less engaging or useful to readers, instead validating the effectiveness of the human-AI collabo-

rative framework in maintaining content utility at scale.

Expansion vs. Substitution. An interesting question is whether translated content expands overall knowledge consumption or substitutes readership from source languages. Either outcome would be valuable: even pure substitution would enable readers to access content in their preferred language. To explore this, we analyze the combined readership of source-target article pairs, normalizing against the pre-translation source readership baseline (set to 1.0).

Figure 10: Sustained Reader Engagement and the Expansion of Knowledge Access



Note: Analysis of reader engagement (pageviews) with AI-assisted content. Panel (a) compares the monthly pageview trajectories for articles created before (purple) and after (green) Google Translate integration (t=0 is the creation month). Both cohorts show consistent engagement, stabilizing at 80-90 monthly views, indicating maintained content utility despite the volume surge. Panel (b) analyzes the combined readership of source-target article pairs, normalized against the pre-translation source readership (baseline=1.0). The combined readership increases to approximately 1.6 times the baseline, demonstrating genuine consumption expansion rather than substitution.

Panel (b) of Figure 10 reveals that following translation, combined readership increases to approximately 1.6 times the baseline and persists at this elevated level. This demonstrates genuine consumption expansion: the availability of content in target languages attracts new readership rather than diverting existing readers from source articles. This finding indicates that AI-assisted translation creates net value by making knowledge accessible to previously underserved linguistic communities.[5]

[5]Here we examine substitution vs expansion at individual article level. For an analysis at the macro

These readership patterns reveal that AI-assisted translation creates measurable value: it maintains content utility while expanding knowledge access for underserved linguistic communities.

6.4 Economic and Social Significance

While 80-90 monthly views per article might appear modest in isolation, the aggregate impact is substantial and economically significant in the context of global knowledge access.

The scale of aggregate knowledge consumption is remarkable. Our estimates imply approximately *12.3 million additional pageviews* over our two-year study period—knowledge consumption that would not have occurred without AI assistance.[6] This equals the total traffic of a mid-sized Wikipedia edition like Hebrew or Czech, created through AI augmentation rather than decades of manual effort.

These viewing patterns reveal that translated articles represent high-value additions to Wikipedia's knowledge base. In Wikipedia's characteristic long-tail distribution, where the median article receives only 10-20 monthly views, translated articles consistently achieving 80-90 views, placing them in the 75th-80th percentile of engagement. This indicates that AI-facilitated content addresses genuine knowledge gaps rather than creating low-quality filler.

For smaller language communities, the acceleration is transformative. Consider Bengali Wikipedia, serving almost 300 million speakers but having accumulated only 80,000 articles by 2019. The AI-enabled acceleration compresses what would require decades of manual effort into just two years, fundamentally altering the trajectory of knowledge development for underserved language communities.

The significance extends beyond immediate metrics through positive externalities. Expanded Wikipedia access influences scientific research (Thompson and Hanley, 2018), judicial

language level, see Appendix E

[6]Calculation: 71.2 articles/month/language × 90 treated languages × 24 months × 80 views/month = 12.3 million views.

decisions (Thompson et al., 2024), and economic outcomes (Hinnosaar et al., 2023). These multiplier effects amplify the social value of expanded knowledge access beyond the direct readership numbers.

In summary, these aggregate results provide compelling evidence for AI's democratizing potential. The deployment of advanced machine translation drove massive productivity gains while maintaining quality and significantly expanding global knowledge access. However, this aggregate view masks critical disparities in how these benefits are distributed, which we examine next.

7 Distributional Benefits: AI Democratization Paradox

The previous section established a significant aggregate increase in AI-assisted content creation, supporting the technological optimism view. However, aggregate effects can mask critical disparities. We now examine the distributional impact of this AI deployment, investigating the tension between democratizing forces and structural reproduction.

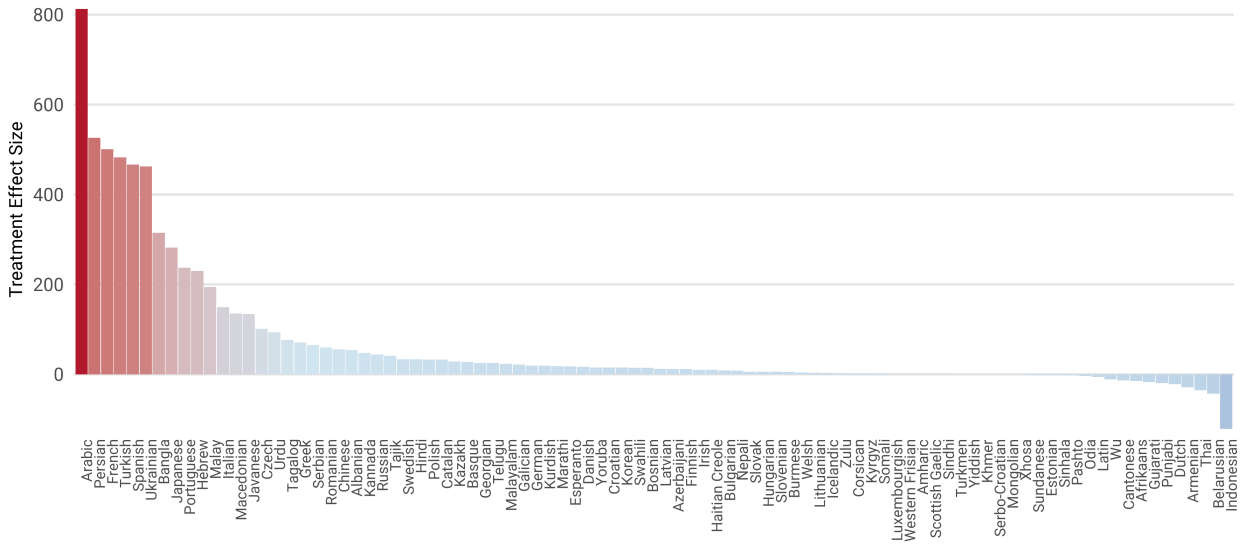
7.1 Heterogeneous Impacts Across Languages

We next examine how the benefits of AI deployment are distributed across Wikipedia's diverse language landscape. This section tests the *structural reproduction* hypothesis: even when communities gain equal technical access to better translation technology, do those with greater complementary resources capture disproportionate gains?

Language-Level Treatment Effects. Our counterfactual imputation approach yields language-specific treatment effects by directly imputing post-integration outcomes in the absence of the tool. Figure 11 plots the estimated effect for each of the 90 treated languages, providing a distributional view of how translation-mediated content creation changed across lanuages. Most languages experience positive gains, but the distribution is highly skewed: a

small set of languages see very large improvements, while a majority realize modest effects. This concentration motivates a systematic analysis of which language characteristics are associated with larger gains.

Figure 11: Language-level treatment effects from AI-enabled translation.



Notes: Each bar reports the estimated treatment effect (post-integration outcome minus the imputed counterfactual) for one of the 90 treated language editions, sorted from largest to smallest. Values reflect changes in translation-mediated article creation during the post period; positive bars indicate increases relative to the counterfactual without the tool.

Unequal Returns to AI: A Matthew Effect. To probe the determinants of heterogeneity, we analyze treatment effects along three dimensions of community resources: (i) editorial capacity (active editors), (ii) readership networks (pageviews), and (iii) knowledge base size (existing articles). We partition language editions into ten deciles for each dimension and estimate a heterogeneous difference-in-differences model that interacts treatment with decile indicators.

Figure 12 reveals a pronounced gradient. While most communities benefit from Google Translate, well-resourced language editions gain far more. Languages in the top decile realize

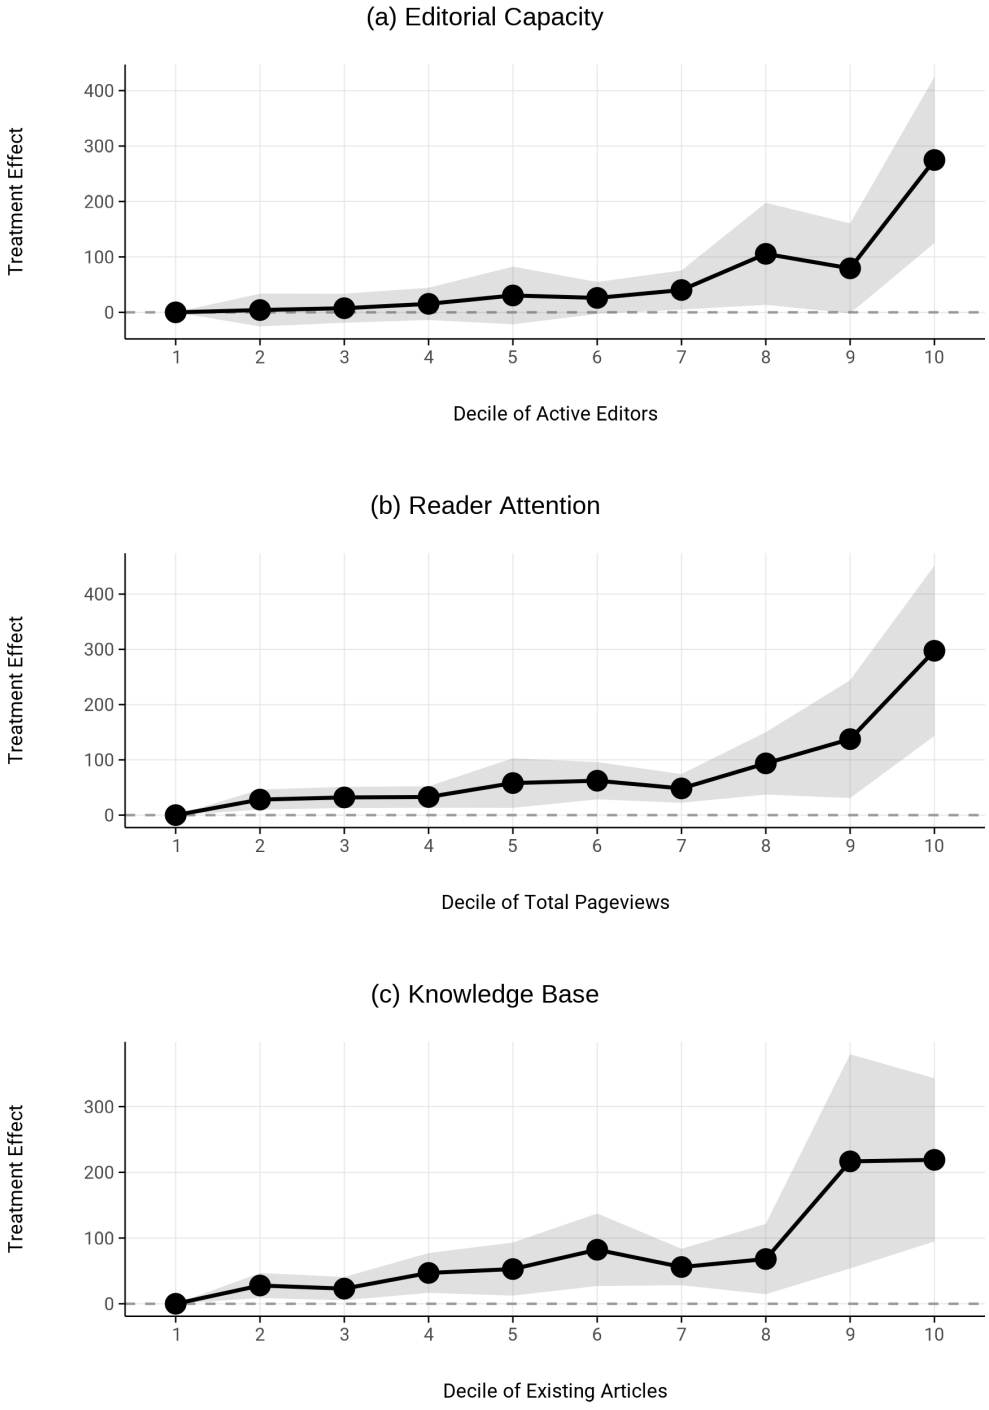
effects that are roughly three to four times those of mid-tier communities and more than an order of magnitude larger than the smallest editions. The pattern is consistent across all three resource dimensions.

These results indicate that technological benefits systematically concentrate among already advantaged communities (e.g., large editions such as French, Spanish, Italian, and Arabic). Those editions were already central hubs of translation activity prior to the intervention (see Figure 4 and Figure 5).

We interpret these findings as robust conditional correlations demonstrating systematic variation in community responses to AI deployment, rather than strict causal mechanisms of community size per se. The key insight is that technological capabilities interact with existing social structures to produce differential outcomes.

Together, the findings support the structural reproduction hypothesis. Although AI lowers language-related barriers, communities with greater complementary resources are better positioned to translate, review, and integrate content at scale. As a result, the shock can widen, rather than narrow, gaps between resource-rich and resource-poor communities.

Figure 12: The Matthew Effect in AI deployment: larger gains for well-resourced communities.



Notes: Heterogeneous difference-in-differences estimates by deciles of complementary resources: knowledge base size (existing articles), editorial capacity (active editors), and readership networks (pageviews). Points/lines show estimated effects relative to the pre-integration baseline; shaded bands denote 95% confidence intervals. Top-decile languages experience gains 3–4× larger than mid-tier communities and over 10× larger than the smallest editions.

Centralized Sources, Diversified Access. We then ask how these heterogeneous gains reshape cross-linguistic knowledge flows. Figure 13 shows two simultaneous trends after integration.

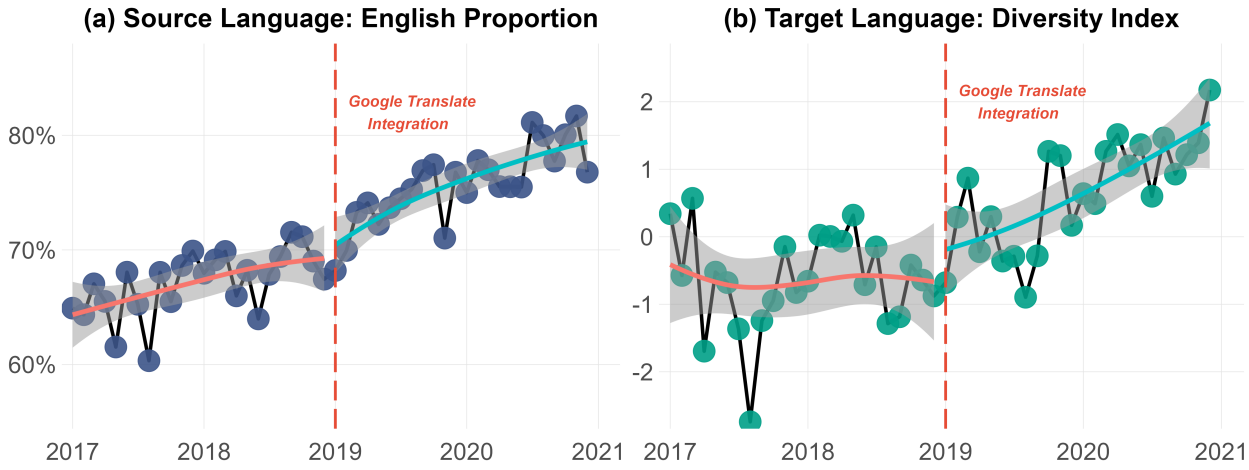


Figure 13: The paradox of global knowledge flows: more centralized sources, more diversified access.

Notes: Evolution around the Google Translate integration (vertical red line). Panel (a): source centralization—English Wikipedia's share as the origin of translation rises from 68% to 81%. Panel (b): access diversification—the Shannon entropy of target languages increases, indicating that translated knowledge reaches a broader array of recipient languages.

Panel (a) documents increasing centralization of sources, with English's dominance intensifying post-integration. Panel (b), however, shows concurrent democratization of access: the diversity of target languages rises markedly, as measured by Shannon entropy. In short, AI produces a more centralized *origin* of knowledge flows while simultaneously broadening *access* across recipient languages—a systemic paradox of concentration and democratization.

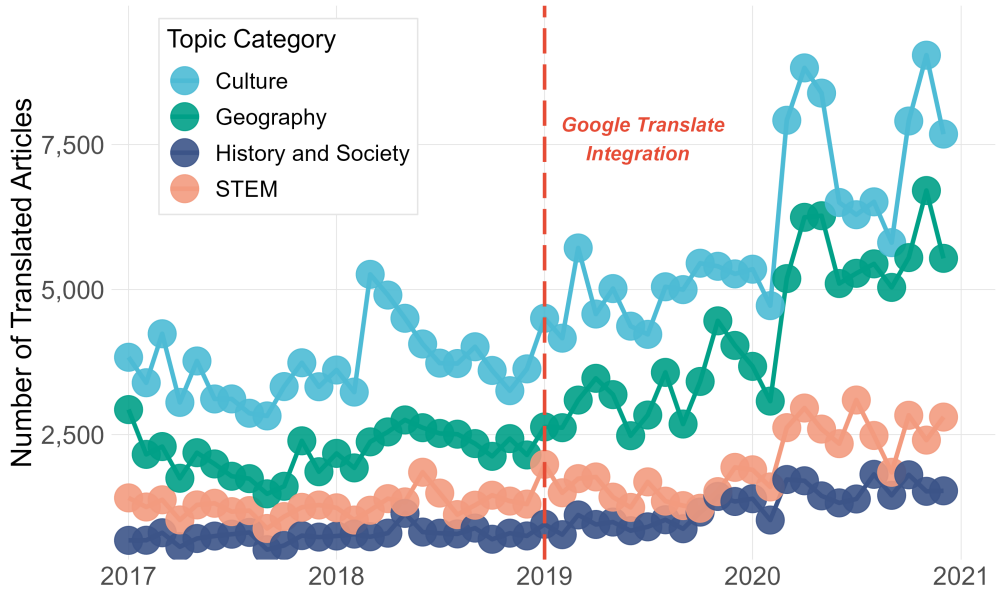
7.2 Content Representation

Beyond volume effects, we now examine *what* content gets translated, investigating whether efficiency gains address systemic representation gaps (democratization through agency) or merely reproduce existing biases (structural reproduction).

Figure 14 shows heterogeneous growth across knowledge domains. Cultural and geo-

graphical articles show the most pronounced increases and collectively represent 75% of translations. Given that these are also critical domains for representation, we focus our subsequent analysis on these areas.

Figure 14: Heterogeneous Growth Across Knowledge Domains, Dominated by Cultural and Geographical Content



Note: Monthly translation volume across major content categories before and after Google Translate integration (red line). The growth is heterogeneous across domains, with Cultural and Geographical articles showing the most pronounced increases. These two categories collectively represent 75% of all translations post-integration.

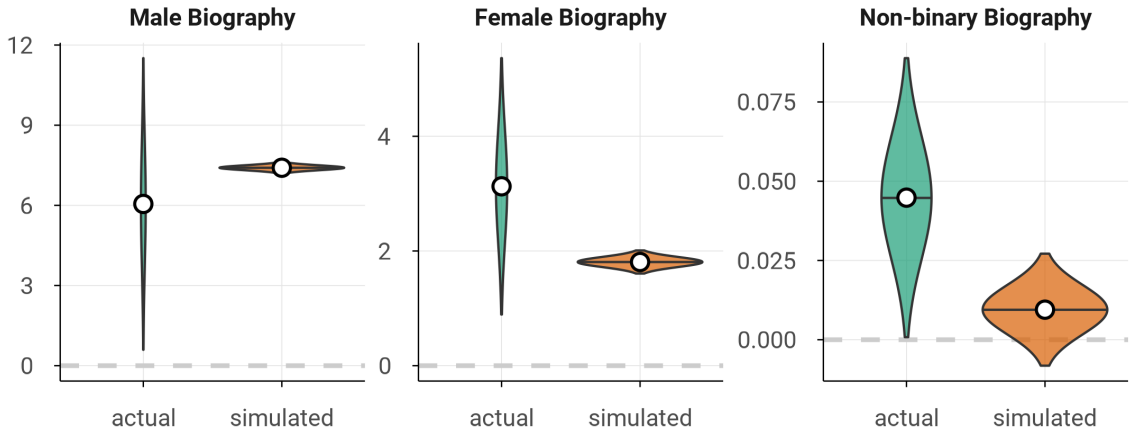
Detecting Human Agency via Simulation Benchmarks. To rigorously assess agency, we employ a simulation benchmark analysis. We establish a counterfactual baseline by simulating the translation patterns that would emerge if editors selected articles according to the distribution of the existing resource pool. This null model represents the expected outcome if translation patterns simply reproduced existing structural imbalances without intentional intervention. By comparing the actual observed translation patterns (estimated via DiD) against this simulated benchmark, we can detect active efforts to address representation gaps.

Gender Representation: Democratization with Limitation. We first examine

gender representation in biographical articles, where women and non-binary individuals face well-documented underrepresentation. English Wikipedia's baseline distribution starkly illustrates this imbalance: 82.2% male, 17.7% female, 0.1% non-binary biographies.

If editors passively reproduced these patterns, translated content would mirror these proportions. Instead, actual post-integration translations show substantial rebalancing: 62.4% male, 37.2% female, 0.4% non-binary. Figure 15 compares our DiD estimates against simulated benchmarks based on random selection from the existing source pool.[7]

Figure 15: Evidence of Human Agency: Leveraging AI to Address Gender Representation Gaps



Note: Comparison of observed post-integration translation patterns (DID estimates with 95% CIs) against simulated benchmarks (violin plots representing 500 replications). Benchmarks represent expected outcomes based on random selection from the existing, biased source pool (e.g., English Wikipedia: 17.7% female). Female biographies are translated at approximately twice the expected rate (actual observed rate: 37.2%), and non-binary translations also significantly exceed their benchmarks. Male biographies fall significantly below expectations. This suggests editors actively used AI efficiency gains to prioritize underrepresented content.

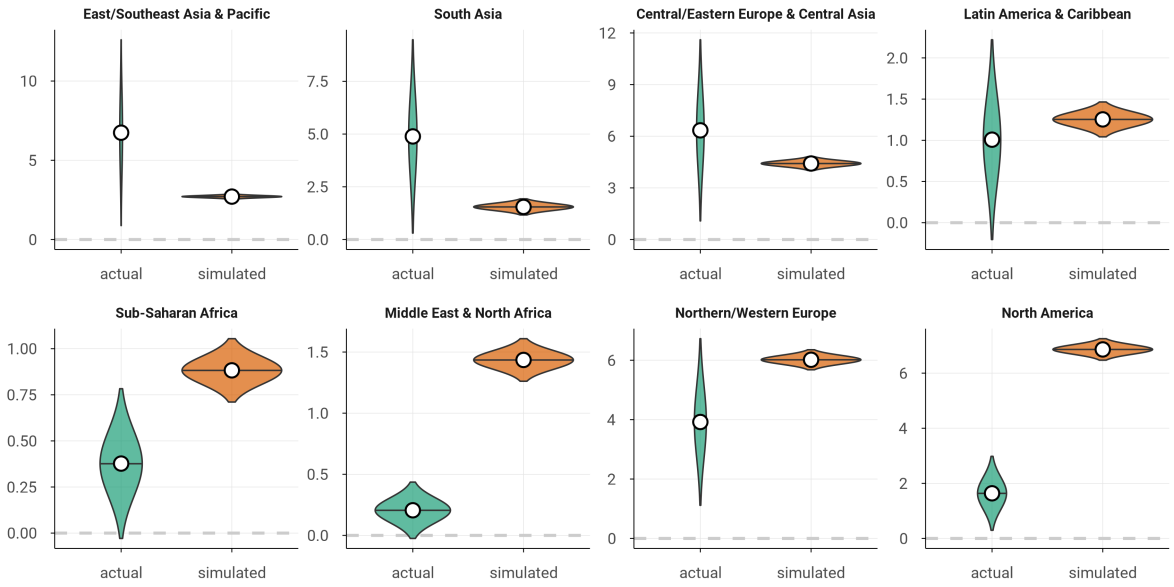
The results reveal significant deviations from the null model. Female biography translations occur at approximately twice the expected rate, while non-binary translations also substantially exceed their benchmark. Male biographies fall significantly below simulated expectations.

[7]See Table A.6 in Section D.4 for detailed regression results.

These patterns provide suggestive evidence of human agency in leveraging AI tools. Editors are actively using the efficiency gains from advanced machine translation to prioritize underrepresented content. However, the persistence of male biographical dominance in absolute terms suggests that technological solutions alone, while beneficial, may require complementary interventions to fully address deeply embedded structural inequalities.

Geographic Representation: Structural Constraints. Applying the same methodology to geographical content across eight global regions (defined by Wikimedia Foundation; see Section D.6) reveals a more complex picture. Figure 16 shows mixed evidence when comparing between actual and simulated translation patterns across regions.

Figure 16: The Limits of Agency: Structural Constraints Hinder Progress in Geographic Representation



Note: Comparison of observed translation patterns (DID estimates) against simulated benchmarks for geographical content across global regions. While some rebalancing occurs (East/Southeast Asia and Central/Eastern Europe exceed benchmarks; North America and Western Europe fall below), regions with the greatest need (e.g., Sub-Saharan Africa, Middle East & North Africa) show minimal absolute gains and do not significantly deviate from random selection. This highlights how structural constraints (lack of source material and local editors) limit the impact of AI.

Some progress emerges: East/Southeast Asia and Central/Eastern Europe significantly

exceed simulated baselines, while North America and Western Europe fall below expectations, suggesting intentional rebalancing. However, regions with the greatest needs—Sub-Saharan Africa and Middle East & North Africa—show minimal absolute gains (0.4 and 0.3 articles/month) that don't significantly deviate from random selection.

This reveals the boundaries of human agency. Where source content exists and some editorial capacity remains, editors can leverage AI for bias correction. But in regions lacking both source material and local editors, even advanced AI combined with good intentions cannot overcome structural constraints.

7.3 Summary of the Paradox

Our distributional and content analyses reveal what we term the *AI Democratization Paradox*: the simultaneous operation of democratizing and concentrating forces following AI deployment. This paradox manifests across three interconnected key dimensions:

The Mechanism-Outcome Contradiction. Google Translate operates through an inherently democratizing mechanism—expertise unbundling that universally lowers language barriers. Yet this barrier-reducing technology produces concentrating outcomes, with well-resourced communities capturing gains at least 3-4 times larger than their disadvantaged counterparts. The paradox lies in how a universally accessible tool amplifies rather than alleviates structural inequalities.

Coexistence of Opposing Forces. The same technological intervention generates contradictory dynamics simultaneously. Democratizing forces are evident: knowledge access in more diverse languages, and female biographies are translated at twice the expected rate. Yet concentrating forces operate in parallel: English Wikipedia's dominance increases from 68% to 81%, and regions with highest need see minimal benefit. AI functions simultaneously as both equalizer and amplifier of inequality.

The Equal Access Fallacy. Perhaps most striking, our findings challenge the intuitive belief that equal access to democratizing technology should narrow gaps between

communities. Instead, identical access produces sharply divergent outcomes. This reveals a deeper truth: the binding constraint on democratization isn't technological capability but complementary resources—editorial capacity, existing knowledge, organizational infrastructure—that remain fundamentally unequal.

This paradox demonstrates that AI's distributional impact emerges not from its technical features but from the complex interaction between technological affordances and social structures. While AI can lower barriers and enable human agency to address inequalities, these democratizing effects are systematically mediated—and often overwhelmed—by the unequal distribution of complementary resources. Understanding this dynamic is essential for designing AI deployments that achieve rather than undermine equitable outcomes.

8 Conclusion and Implications

This study provides the first large-scale empirical evidence on how AI deployment reshapes knowledge production in decentralized platforms. By leveraging Wikipedia's integration of Google Translate as a natural experiment across more than 100 diverse language communities, we examined the impact of advanced AI on cross-linguistic content creation. Our findings reveal a central tension in AI deployment within collaborative systems, which we term the *AI Democratization Paradox*.

Our empirical analysis demonstrates that AI deployment simultaneously unleashes both democratizing and concentrating forces. The democratizing forces are substantial: we found a more than double increase in aggregate translation-mediated content creation. This generates an estimated 12.3 million additional pageviews over two years, critically distributed across a diverse array of target languages, making knowledge accessible to speakers of previously underserved linguistic communities worldwide. Moreover, we observed clear evidence of human agency leveraging AI to address systemic inequalities, with female biographies translated at twice the expected rate.

However, these gains are profoundly mediated by strong concentrating forces. Benefits disproportionately flow to well-resourced language communities with greater editorial capacity, larger existing knowledge bases, and established reader networks. These communities experienced gains 3-4 times larger than smaller editions, while regions with the highest need, such as Sub-Saharan Africa, saw minimal impact. This paradox highlights that AI's impact critically depends on the interaction between technological capabilities and existing social structures.

Our study advances understanding of AI adoption, digital platforms, and inequality by moving beyond "AI increases productivity" narratives to demonstrate how socio-technical systems shape AI's distributional impact. We provide a conceptual framework reconciling technological optimism and structural reproduction views, showing both mechanisms operate simultaneously. We empirically validate that technological capabilities alone cannot overcome structural inequalities, as the capacity to leverage AI depends on existing human capital and organizational infrastructure. Finally, we highlight human agency's role in steering AI toward equitable outcomes while identifying constraints on that agency.

Our findings offer actionable insights for organizations deploying AI in collaborative systems. Providing equal access to AI tools is necessary but insufficient for equitable outcomes. Platforms must recognize that "one-size-fits-all" AI deployment risks amplifying existing disparities. To realize AI's democratizing potential, technological solutions must be coupled with complementary interventions supporting under-resourced communities through targeted training, capacity-building, and infrastructure support. Organizations should evaluate AI impact through a distributional lens rather than focusing solely on aggregate productivity metrics. Policymakers addressing the digital divide must recognize that equitable access to information requires fostering both technological access and the human and organizational capacity to leverage that technology effectively.

Avenues for Future Research. Our findings open several promising directions that extend both theoretical understanding and practical applications of AI in collaborative systems. First, could platform operators mitigate AI's concentrating effects through targeted interventions? Future research should examine how governance choices—such as providing editorial support, capacity-building programs, or differential incentives to under-resourced communities—can shape AI's distributional impact. This connects our framework to the growing literature on platform governance, offering actionable insights for organizations deploying AI in decentralized systems.

Second, as AI systems become increasingly autonomous, will the democratization paradox intensify or evolve? Advanced generative models might prove more democratizing by reducing dependence on complementary human resources, or they may create new inequality mechanisms through AI skill gaps and sophisticated content verification challenges. Understanding how our paradox transforms with increasingly capable AI represents a promising research frontier.

Third, does AI-assisted content creation, which predominantly flows from dominant languages like English, promote cultural homogenization or perpetuate embedded biases? Future work should examine whether efficiency gains from AI come at the cost of cultural authenticity and diverse knowledge traditions. Exploring these nuanced dimensions of value in global knowledge commons offers rich opportunities for interdisciplinary collaboration.

In conclusion, realizing AI's democratizing potential requires sustained investigation of the complex interplay between technological capabilities, social structures, and governance mechanisms. Recognizing and addressing this challenge—by coupling technological innovation with targeted structural support—is essential for ensuring that the AI revolution contributes to a more equitable digital future.

References

Aaltonen, A. and Seiler, S. (2016). Cumulative growth in user-generated content production: evidence from wikipedia. *Management Science*, 62(7):2054–2069.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). *The economics of artificial intelligence: An agenda*. University of Chicago Press.

Agrawal, A., Gans, J., and Goldfarb, A. (2022). *Prediction Machines, Updated and Expanded: The Simple Economics of Artificial Intelligence*. Harvard Business Press.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G., and Khosravi, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730.

Benkler, Y. (2002). Coase's penguin, or, linux and" the nature of the firm". *Yale law journal*, pages 369–446.

Brynjolfsson, E., Hui, X., and Liu, M. (2019). Does machine translation affect international trade? evidence from a large digital platform. *Management Science*, 65(12):5449–5460.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Capraro, V., Lentsch, A., Acemoglu, D., Akgun, S., Akhmedova, A., Bilancini, E., Bonnefon, J.-F., Brañas-Garza, P., Butera, L., Douglas, K. M., et al. (2024). The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS nexus*, 3(6):pgae191.

Cazzaniga, M., Jaumotte, M. F., Li, L., Melina, M. G., Panton, A. J., Pizzinelli, C., Rockall, E. J., and Tavares, M. M. M. (2024). *Gen-AI: Artificial intelligence and the future of work*. International Monetary Fund.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American economic review*, 110(9):2964–2996.

Faraj, S., von Krogh, G., Monteiro, E., and Lakhani, K. R. (2016). Special section introduction—online community as space for knowledge flows. *Information systems research*, 27(4):668–684.

Forman, C., Goldfarb, A., and Greenstein, S. (2012). The internet and local wages: A puzzle. *American Economic Review*, 102(1):556–575.

Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Eonomic Literature*, 57(1):3–43.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of econometrics*, 225(2):254–277.

Graham, M., Hogan, B., Straumann, R. K., and Medhat, A. (2014). Uneven geographies of user-generated information: Patterns of increasing informational poverty. *Annals of the Association of American Geographers*, 104(4):746–764.

Hinnosaar, M., Hinnosaar, T., Kummer, M., and Slivko, O. (2023). Wikipedia matters. *Journal of Economics & Management Strategy*, 32(3):657–669.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960.

Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The Quarterly Journal of Economics*, 122(3):879–924.

Johnson, I., Gerlach, M., and Sáez-Trumper, D. (2021). Language-agnostic topic classification for wikipedia. In *Companion Proceedings of the Web Conference 2021*, pages 594–601.

Miquel-Ribé, M. and Laniado, D. (2020). The wikipedia diversity observatory: A project to identify and bridge content gaps in wikipedia. In *Proceedings of the 16th International Symposium on Open Collaboration*, pages 1–4.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.

Parker, G. G., Van Alstyne, M. W., and Choudary, S. P. (2016). *Platform revolution: How networked markets are transforming the economy and how to make them work for you*. WW Norton & Company.

Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029.

Scheidel, W. (2024). Beyond technology and wages: power and the history of inequality. *Oxford Open Economics*, 3(Supplement_1):i212–i216.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

Teh, T.-H. (2022). Platform governance. *American Economic Journal: Microeconomics*, 14(3):213–254.

Thompson, N. and Hanley, D. (2018). Science is shaped by wikipedia: evidence from a randomized control trial.

Thompson, N., Luo, X., McKenzie, B., Richardson, E., and Flanagan, B. (2024). User-generated content shapes judicial reasoning: Evidence from a randomized control trial on wikipedia. *Information Systems Research*.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

Von Hippel, E. (2006). *Democratizing innovation*. the MIT Press.

Wagner, C., Graells-Garrido, E., Garcia, D., and Menczer, F. (2016). Women through the glass ceiling: gender asymmetries in wikipedia. *EPJ Data Science*, 5:1–24.

Wessels, B. (2013). The reproduction and reconfiguration of inequality: Differentiation and class, status and power in the dynamics of digital divides. In *The digital divide*, pages 17–28. Routledge.

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

Zhang, X. and Zhu, F. (2011). Group size and incentives to contribute: A natural experiment at chinese wikipedia. *American Economic Review*, 101(4):1601–1615.

Zhu, K., Walker, D., and Muchnik, L. (2020). Content growth and attention contagion in information networks: Addressing information poverty on wikipedia. *Information Systems Research*, 31(2):491–509.

Web Appendix

A Data Source

A.1 Published Translations on Wikipedia

The primary data for our analysis comprises published translations from Wikipedia's Content Translation system. Wikimedia Foundation provides data access through Application Programming Interface (API) that lists published translations, offering comprehensive details essential for our study:

- **API Access**: The API provides lists of all published translations across languages, including details like translation IDs, translation date, source and target titles, source and target languages, URLs of the articles, and translation statistics indicate the extent of human versus machine translation.
- **Parallel Text Dump**: The tool also allows access to parallel text data, encompassing the source content and the final user-translated content.

This data is provided under an open license and represents an ever-improving resource for cross-linguistic research and machine translation development. For more information, refer to the MediaWiki page: Content Translation - Published Translations

A.2 Wikipedia Knowledge Gap Index

The Knowledge Gap Index developed by the Wikimedia Foundation is instrumental in quantifying knowledge disparities across various dimensions on Wikipedia. The dataset is structured to provide a clear mapping of content gaps across languages

- **Content Mapping:** The Knowledge Gap Index employs a content mapping system, categorizing Wikipedia articles into specific domains and themes. This system allows us to assess disparities in content availability and representation across diverse topics and regions.
- **Gender Gap Measurement:** the Knowledge Gap Index addresses the Gender Gap. This measure assesses the balance in representation of different genders across Wikipedia articles. It aids in quantifying and understanding the skewness in gender portrayal,
- **Geographic Gap Measurement:** Another useful aspect of the Knowledge Gap Index is its focus on the Geographic Gap. This measurement evaluates the extent of content coverage pertaining to different geographic regions.
- **Open Access and Continuous Updates:** The Knowledge Gap Index is publicly accessible and is regularly updated, reflecting the dynamic nature of content creation on Wikipedia.

The Knowledge Gap Index offers a detailed perspective on the existing content distribution and the potential areas where machine translation can have a significant impact. For further details about the dataset and its applications, please refer to the Wikimedia Meta-Wiki page on the Knowledge Gap Index and its Geographic Gap Measurement.

B Methodology

B.1 Classifying Topics in Wikipedia Articles

Our approach to classifying article topics across different languages in Wikipedia builds on the language-agnostic methodology developed by Johnson et al. (2021). This method involves several steps:

1. **Topic Taxonomy**: Utilizing a taxonomy of 64 topics derived from WikiProject tags, encompassing a broad range of subjects from science and technology to arts and culture.

2. **Link-Based Representation**: Articles are represented by Wikidata items of contained links, facilitating a language-independent content representation.

3. **Mapping Articles to Topics**: Machine learning algorithms map articles to the predefined topics based on link similarities.

4. **Ensuring Cross-Linguistic Consistency**: The focus on link-based features rather than textual content assures consistent classification across languages.

5. **Validation**: Classification accuracy is validated and refined through comparison with a manually classified article sample.

This methodology enables a comprehensive and consistent analysis of article topics across multiple language editions of Wikipedia, overcoming the limitations of language-specific models and ensuring uniform topic classification.

B.2 Extracting Attributes from Wikidata

We utilized Wikidata to extract demographic and geographical data for our study. We used the Wikidata Query Service, which allows for sophisticated and precise queries, to retrieve the required information efficiently and accurately

- **Biographical Articles**: Information such as gender, sexual orientation, and ethnicity of individuals was systematically extracted from their respective Wikidata entries. This process involved querying Wikidata's structured data to obtain these specific attributes for each biographical article.
- **Geographical Articles**: For articles pertaining to geographical locations, we extracted location data, including coordinates and administrative hierarchies. This was achieved by querying the corresponding Wikidata items linked to each geographical article.

This methodology ensured a standardized and accurate extraction of demographic and geographical information, useful for our analysis.

C Robustness for Aggregate Impact of Google Translate Deployment

C.1 Primary Design: Alternative DID Estimator

C.1.1 More Inforamtion about Matrix Completion Method Estimation

The matrix completion method provides additional validation of our main findings on Google Translate's impact. As shown in Table A.2, The matrix completion method approach estimates an Average Treatment Effect on the Treated (ATT) of 71.2 additional monthly translations ($p<0.01$), with a 95% confidence interval of [38.4, 104.1]. This result aligns with our primary difference-in-differences estimates and the two-way fixed effects model, providing robust evidence that Google Translate integration substantially increased translation volume across Wikipedia languages.

Table A.2: Matrix Completion Method: Estimated Impact of Google Translate Integration on Monthly Translation Volume

ATT	S.E.	CI.lower	CI.upper	P.value
71.2	16.8	38.4	104.1	0.0

Note: This table presents the Average Treatment Effect on the Treated (ATT) estimated using the matrix completion method. The ATT of 71.2 indicates the average increase in monthly translations attributable to Google Translate integration. S.E. represents standard error, CI shows 95% confidence interval bounds, and P.value indicates statistical significance.

C.1.2 Two-way Fixed Effect

This model estimates the Average Treatment Effect (ATE) of Google Translate's integration on the volume of translated articles. By incorporating both language and time fixed effects, this approach controls for time-invariant language characteristics and common temporal shocks across languages. The statistically significant treated coefficient of 61.4 ($p<0.01$)

in Table A.3 reinforces the substantial impact of Google Translate on content translation activities in Wikipedia, while accounting for potential confounding factors through the two-way fixed effect structure.

Table A.3: Two-way Fixed Effects Analysis of Google Translate Integration Impact

Dependent Variable:	Number of Translated Articles
Model:	(1)
Google Translate Integration	61.4***
	(14.9)
Fixed-effects	
Language	Yes
Year-Month	Yes
Fit statistics	
Observations	6,336
R^2	0.77

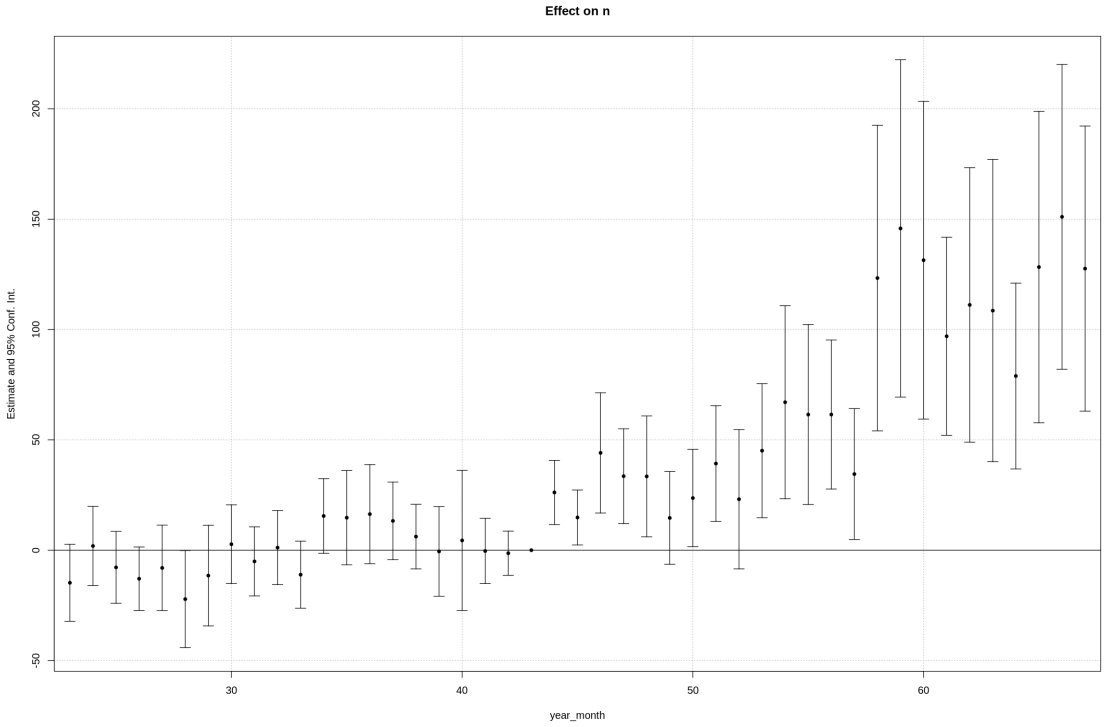
Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents results from a two-way fixed effects model examining the impact of Google Translate integration on monthly translation volume. The coefficient of 61.4 represents the average increase in monthly translations after integration, controlling for both language-specific and temporal effects. Standard errors are clustered at the language level.

C.1.3 Sun and Abraham Weighted Interactions

The Sun and Abraham Weighted Interactions method addresses potential biases in staggered difference-in-differences designs by properly accounting for treatment effect heterogeneity. This approach computes cohort-specific treatment effects and combines them using weights that reflect the relative size of each cohort. The method highlights the temporal dynamics of the treatment effect, indicating a sustained positive impact of Google Translate over time on the translation volume. By explicitly modeling treatment effect heterogeneity across adoption cohorts, this approach provides more reliable estimates of the dynamic treatment effects.

Figure A.17: Dynamic Treatment Effects of Google Translate Integration Using Sun and Abraham Method



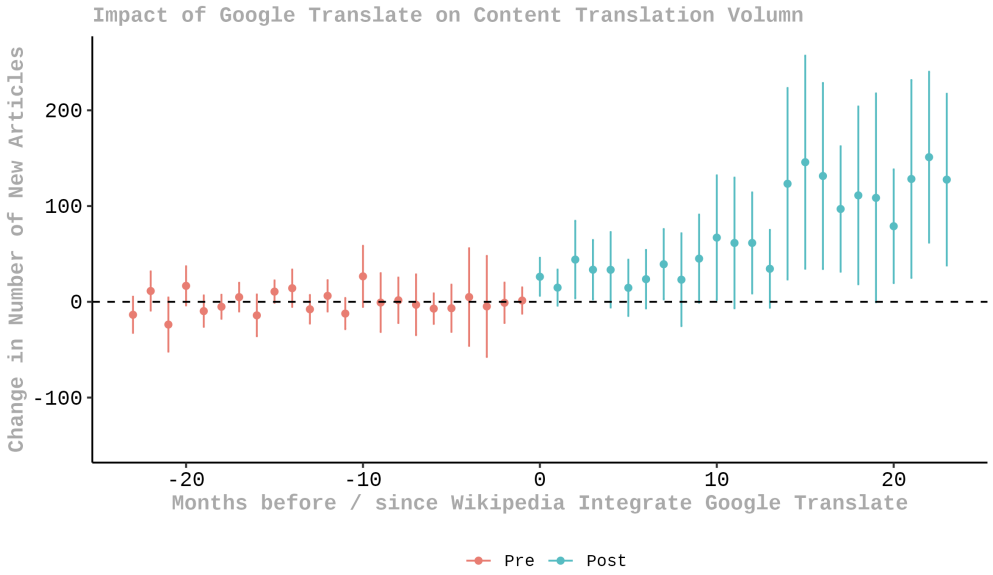
Note: This figure displays the dynamic treatment effects estimated using the Sun and Abraham weighted interactions method. The vertical axis shows the estimated effect on monthly translation volume, while the horizontal axis represents time relative to Google Translate integration. The solid line indicates point estimates, and the shaded area represents 95% confidence intervals. The results demonstrate persistent positive effects of integration across different time periods, accounting for heterogeneous treatment effects across adoption cohorts.

C.1.4 Callaway and Sant'Anna Estimators

We further assess the robustness of our findings using the Callaway and Sant'Anna Estimators, which are specifically designed for settings with staggered treatment adoption and heterogeneous treatment effects. This methodology offers several advantages: it allows for treatment effect heterogeneity both across groups and over time, remains valid under parallel trends assumptions, and is robust to negative weights that can arise in two-way fixed effects models. The estimators compute group-time average treatment effects and aggregate them into overall effects using carefully chosen weights. The Callaway and Sant'Anna Estimators

confirm the robustness of our main findings, showing consistent positive effects of Google Translate's integration on translation volume across various Wikipedia language editions.

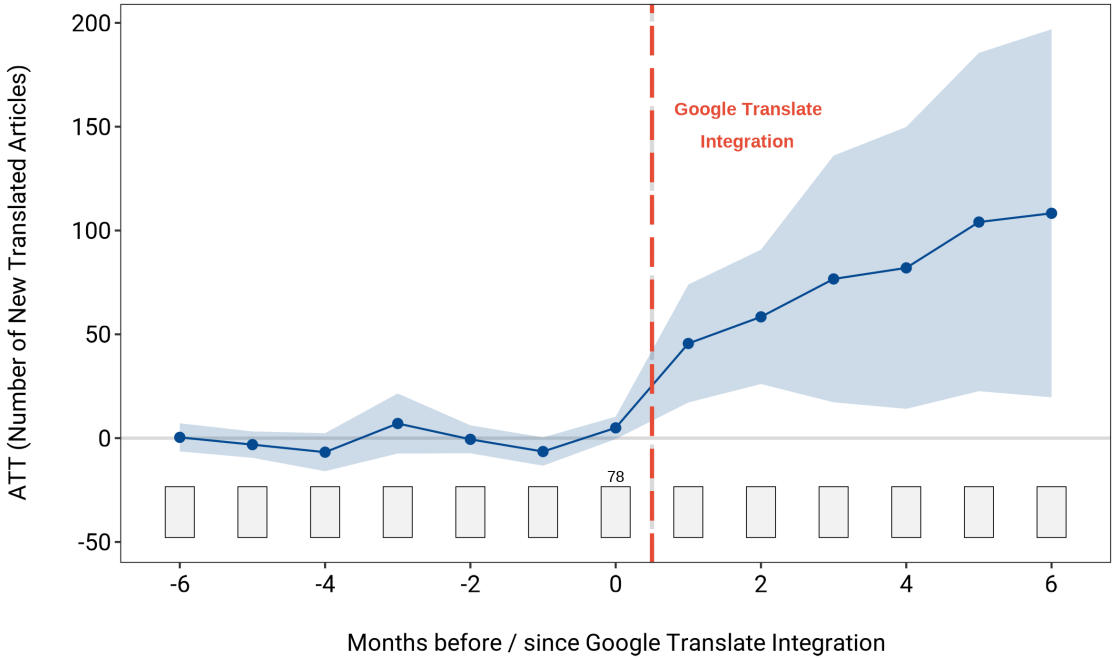
Figure A.18: Group-Time Average Treatment Effects Using Callaway and Sant'Anna Estimators



Note: This figure presents the estimated group-time average treatment effects using the Callaway and Sant'Anna method. The vertical axis shows the estimated effect on monthly translation volume, while the horizontal axis represents time relative to Google Translate integration. The solid line shows point estimates, and the shaded area represents 95% confidence intervals. The results demonstrate robust positive effects across different adoption groups and time periods, avoiding potential biases from negative weighting in traditional two-way fixed effects models.

C.2 Alternative Identification: Time-Shift DID Design

Figure A.19: Effect of Google Translate Integration on Monthly Translation Volume



Note: Event study plot showing changes in translation volume relative to Google Translate integration (t=0). The black line represents point estimates with 95% confidence intervals (shaded area).

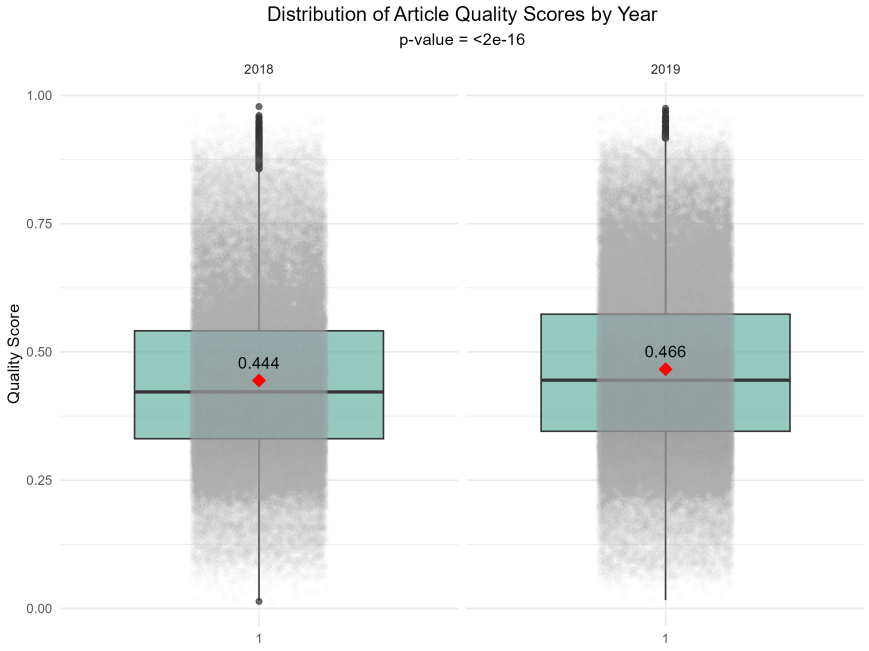
D Additional Supporting Results

The results presented in this section provide detailed model estimates and statistical evidence supporting the key findings discussed in the main paper. While the main text focuses on the high-level implications and interpretation of our analyses, here we present tables showing the full model specifications, robustness checks, and additional statistical tests.

D.1 Article Quality Analysis

We analyze the evolution of article quality between 2018 and 2019. The results show a modest but statistically significant improvement in article quality scores over this period, with the mean quality score increasing from 0.444 in 2018 to 0.466 in 2019. The distribution of scores, as visualized through both box plots and density plots, reveals that while there is considerable variation in article quality within each year, there is a consistent upward shift in the overall quality distribution from 2018 to 2019.

Figure A.20: Distribution of Article Quality Scores by Year (2018-2019)



Note: This figure shows the distribution of article quality scores comparing 2018 and 2019. The box plots display the median, quartiles, and range of quality scores, with the underlying data points shown as a density plot in gray. The red diamond indicates the mean quality score for each year (0.444 in 2018 and 0.466 in 2019). The extremely low p-value (<2e-16) suggests a statistically significant difference between the two years' distributions. The plot demonstrates a slight improvement in average article quality from 2018 to 2019.

D.2 Language Community Characteristics

Using the difference-in-differences estimator described in the Data and Methods section, we examine how the impact of Google Translate varies across different dimensions of language community characteristics. The results in Table A.4 show heterogeneous effects across language tiers based on article count, editor count, view count, and Google Translate quality. When languages are categorized by article count, medium-tier languages saw an increase of 78.3 articles per month and high-tier languages experienced an increase of 181.3 articles compared to the baseline, both statistically significant. Similar patterns emerge when languages are categorized by editor and view counts, with high-tier languages showing the largest gains (256.6 articles for both metrics). Interestingly, when considering Google Translate quality, only medium-tier languages show a significant positive effect of 49.4 articles, while high-tier languages show no significant increase. These findings suggest that community size and engagement are crucial factors in leveraging machine translation technology, potentially more so than the underlying quality of the translation system.

Table A.4: Heterogeneous Effects of Google Translate Integration Across Language Community Characteristics

Dependent Variable:	Number of New Translations			
Language Tier Based On:	Article Count	Editor Count	View Count	GT Quality
Model:	(1)	(2)	(3)	(4)
Variables				
GT Integration × Low-tier	13.3	0.6	4.3	-10.2
	(9.5)	(5.8)	(6.9)	(20.5)
GT Integration × Medium-tier	78.3**	56.6***	52.9***	49.4**
	(30.2)	(19.9)	(19.8)	(23.8)
GT Integration × High-tier	181.3***	256.6***	256.5***	-19.9
	(59.7)	(73.0)	(73.0)	(29.0)
Fixed-effects				
Language	Yes	Yes	Yes	Yes
Year-Month	Yes	Yes	Yes	Yes
Fit statistics				
Observations	6,048	6,048	6,048	3,744
R^2	0.777	0.790	0.789	0.722
Within R^2	0.059	0.112	0.110	0.015

Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents difference-in-differences estimates examining how Google Translate's impact varies across language community characteristics. Each column represents a different categorization of languages: by article count, editor count, view count, and translation quality. The coefficients show the differential impact across low, medium, and high-tier languages, with standard errors clustered at the language level. The results demonstrate that community size and engagement significantly influence the effectiveness of machine translation integration.

D.3 Content Domain Analysis

Our difference-in-differences analysis reveals varying impacts across content domains, as shown in Table A.5. Culture and Geography articles saw the largest increases (25.2 and 21.4 articles per month respectively), while STEM and History/Society topics experienced more modest growth (7.9 and 5.6 articles). This pattern suggests that Google Translate may be more effective at handling cultural and geographical content compared to technical or

Table A.5: Differential Impact of Google Translate Integration Across Content Domains

Dependent Variables:	Culture	Geography	STEM	History and Society
Model:	(1)	(2)	(3)	(4)
GT Integration	25.2***	21.4***	7.9***	5.6**
	(5.9)	(5.4)	(2.1)	(2.3)
Fixed-effects				
Language	Yes	Yes	Yes	Yes
Year-Month	Yes	Yes	Yes	Yes
Fit statistics				
Observations	6,336	6,336	6,336	6,336
R^2	0.76	0.64	0.62	0.58
Pre-integration Mean	20.8	12.5	7.2	4.3

Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents difference-in-differences estimates showing how Google Translate's impact varies across different content domains. The coefficients represent the average monthly increase in translations for each content type after integration. Standard errors are clustered at the language level, and pre-integration means are provided for context. The results reveal stronger effects for cultural and geographical content compared to technical and historical topics.

historical material.

D.4 Gender Analysis

The difference-in-differences estimates in Table A.6 reveal persistent gender disparities in biographical translations after Google Translate integration. Male biographies increased by 7.2 articles monthly compared to 3.8 for female biographies, while non-binary biographies saw minimal growth (0.06). This suggests that machine translation may reinforce existing gender biases in content creation.

Table A.6: Differential Impact of Google Translate Integration on Biographical Content by Gender

Dependent Variables:	Male Biography	Female Biography	Non-binary Biography
Model:	(1)	(2)	(3)
GT Integration	7.2***	3.8***	0.06***
	(2.5)	(1.3)	(0.02)
Fixed-effects			
Language	Yes	Yes	Yes
Year-Month	Yes	Yes	Yes
Fit statistics			
Observations	6,288	6,288	6,288
R^2	0.72317	0.54092	0.26982
Pre-integration Mean	12.5	7.8	0.06

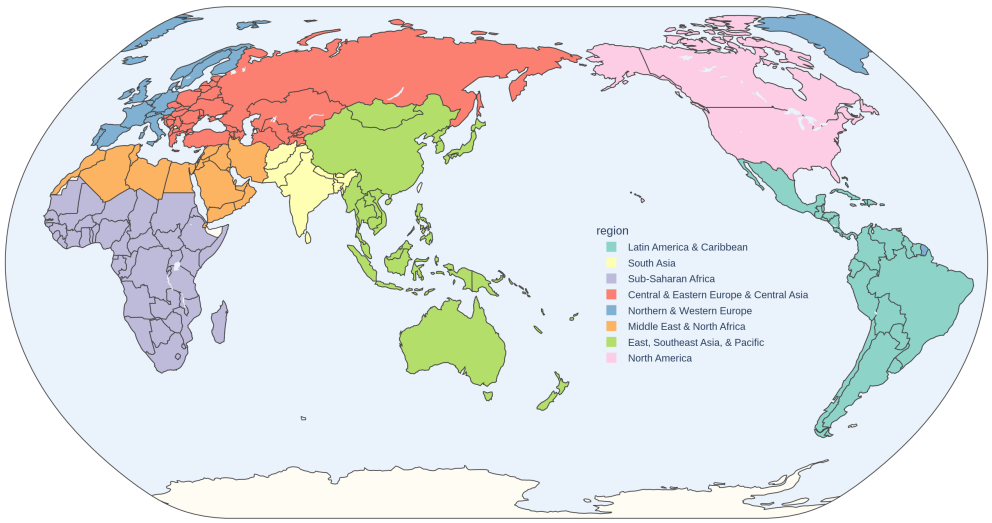
Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents difference-in-differences estimates showing how Google Translate integration affects the translation of biographical articles across gender categories. The coefficients represent the average monthly increase in translations for each gender category after integration. Standard errors are clustered at the language level, and pre-integration means are provided for context. The results reveal a persistent gender gap, with male biographies receiving substantially more translations than female or non-binary biographies.

D.5 Geographic Analysis

Figure A.21 illustrates the geographical framework we use to analyze regional patterns in Wikipedia translations. By adopting the Wikimedia Foundation's official regional classification system, which divides the world into eight distinct regions, we can systematically examine how Google Translate's impact varies across different geographical and cultural contexts. This standardized regional breakdown enables us to identify potential disparities in translation activity and technological adoption across diverse global regions.

Figure A.21: Wikimedia Foundation's Regional Classification System for Global Content Analysis



Note: This world map illustrates the eight regional classifications used by the Wikimedia Foundation for content analysis: Latin America & Caribbean (LAC), South Asia (SA), Sub-Saharan Africa (SSA), Central & Eastern Europe & Central Asia (CEECA), Northern & Western Europe (NWE), Middle East & North Africa (MENA), East/Southeast Asia & Pacific (ESEAP), and North America (NA). These regions form the basis for analyzing geographical patterns in Wikipedia translation activity and content distribution.

Using our difference-in-differences framework, we examine geographical coverage across eight major world regions defined by the Wikimedia Foundation (shown in Figure A.21). The results in the following table show that Central & Eastern Europe Central Asia (CEECA) and East/Southeast Asia & Pacific (ESEAP) regions saw the largest increases in geographical article translations (8.6 and 7.3 articles monthly), while regions like Middle East & North Africa (MENA) and Sub-Saharan Africa (SSA) experienced minimal growth (0.3 and 0.4 articles). This regional disparity suggests that machine translation may be reinforcing existing geographical biases in content coverage.

Table A.7: Effects of Google Translate Integration on Monthly Geography Translations by Regions

DVs:	ESEAP	SA	MENA	SSA	NA	SCA	CEECA	NWE
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
GT Integration	7.3***	5.3***	0.3**	0.4**	1.7***	1.3**	8.6***	5.4***
	(2.4)	(1.8)	(0.1)	(0.2)	(0.5)	(0.5)	(2.6)	(1.4)
Fixed-effects								
Language	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Observations	6,336	6,336	6,336	6,336	6,336	6,336	6,336	6,336
R^2	0.49	0.44	0.54	0.26	0.47	0.24	0.55	0.60
Pre-integration Mean	4.7	2.8	0.2	0.2	1.0	0.6	5.0	3.4

Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents difference-in-differences estimates showing the impact of Google Translate integration on monthly article translations across eight Wikimedia regions. ESEAP: East/Southeast Asia & Pacific; SA: South Asia; MENA: Middle East & North Africa; SSA: Sub-Saharan Africa; NA: North America; SCA: Latin America & Caribbean; CEECA: Central & Eastern Europe & Central Asia; NWE: Northern & Western Europe. The coefficients represent the average monthly increase in translations for each region after integration. Standard errors are clustered at the language level, and pre-integration means are provided for context.

D.6 Geographical Wikipedia Articles by Continents

The Table A.8 presents the difference-in-differences analysis results, capturing the impact of Google Translate's roll-out on the volume of articles across different continents. The results show that regions like Asia and Europe experienced a significant increase in article volume, with coefficients of 9.0 and 7.8, respectively, suggesting a more pronounced adoption of machine translation in these areas. America and Africa also saw increases, albeit at a lower magnitude, while Oceania demonstrated a modest but significant rise. These findings reflect the variable impact of machine translation across diverse geographical contexts.

Table A.8: Impact of Google Translate Roll-out on Article Volume by Geography Regions

Dependent Variables:	America	Europe	Asia	Oceania	Africa
Model:	(1)	(2)	(3)	(4)	(5)
Google Translate Integration	3.0***	7.8***	9.0***	0.5***	1.5**
	(1.0)	(2.0)	(3.3)	(0.2)	(0.6)
Fixed-effects					
Language	Yes	Yes	Yes	Yes	Yes
Year-Month	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	6,336	6,336	6,336	6,336	6,336
R^2	0.42505	0.58168	0.42917	0.29379	0.45671
Within R^2	0.00453	0.00887	0.00744	0.00226	0.00466

Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents difference-in-differences estimates showing the impact of Google Translate integration on monthly article translations across five continental regions. The coefficients represent the average monthly increase in translations for each continent after integration. Standard errors are clustered at the language level. The results reveal substantial variation in the effect across continents, with Asia and Europe showing the largest increases in translation volume.

E Analysis of Potential Displacement Effects

This section examines potential displacement effects of Google Translate integration on non-translated content creation and consumption. We analyze whether increased machine-translated content affects the production and readership of native content through resource reallocation or shifts in reader attention.

Data Collection We constructed a comprehensive panel dataset tracking monthly article creation across both translated and non-translated content in treated and control Wikipedia language editions from 2017 to 2020. This dataset captures content creation through all methods across more than 100 Wikipedia language communities, extending beyond articles created through the Content Translation system. Additionally, we collected monthly aggregated pageview statistics sourced from billions of records across Wikipedia language editions to analyze content consumption patterns.

Content Creation Analysis Using this rich dataset, we employed two model specifications examining:

1. The relationship between translated and non-translated content using fixed-effects regression controlling for language and time fixed effects

2. The impact of Google Translate integration on non-translated article creation using difference-in-differences, where the treatment effect is measured as the interaction between treatment group and post-integration period

For each specification, we analyzed both standardized article counts and log-transformed article counts to ensure robustness of our findings.

Table A.9: Analysis of Translated vs. Non-Translated Article Creation

	Dependent Variable: Non-Translated Article Creation			
	Fixed Effects Model		Difference-in-Differences	
	Standardized	Log	Standardized	Log
Variables	(1)	(2)	(3)	(4)
Number of Translated Articles	-0.05	0.04		
	(0.04)	(0.03)		
Google Translate Integration			-0.18	-0.30
			(0.18)	(0.18)
Language Fixed Effects	Yes	Yes	Yes	Yes
Year-Month Fixed Effects	Yes	Yes	Yes	Yes
Observations	5,856	5,856	5,856	5,856
R^2	0.20	0.84	0.20	0.85

Notes: Table reports regression results examining whether increased translated content displaces non-translated article creation. Standardized variables have mean zero and standard deviation one. Standard errors (in parentheses) are clustered at the language level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Content Consumption Analysis To examine potential displacement effects in article consumption, we analyzed the relationship between translated and non-translated article pageviews using similar methodological approaches:

1. A fixed-effects regression examining the direct relationship between translated and non-translated article pageviews

2. A difference-in-differences specification analyzing how Google Translate integration affects non-translated article pageviews

The analysis incorporated both standardized and log-transformed pageview counts to ensure robust findings across different specifications.

Table A.10: Analysis of Translated vs. Non-Translated Article Consumption

	Dependent Variable: Non-Translated Article Pageviews			
	Fixed Effects Model		Difference-in-Differences	
	Standardized	Log	Standardized	Log
Variables	(1)	(2)	(3)	(4)
Number of Translated Pageviews	0.12***	0.18***		
	(0.03)	(0.05)		
Google Translate Integration			0.02	0.03
			(0.012)	(0.04)
Language Fixed Effects	Yes	Yes	Yes	Yes
Year-Month Fixed Effects	Yes	Yes	Yes	Yes
Observations	5,856	5,856	5,856	5,856
R^2	0.99	0.99	0.99	0.99

Notes: Table reports regression results examining whether increased translated content affects pageviews of non-translated articles. Standardized variables have mean zero and standard deviation one. Standard errors (in parentheses) are clustered at the language level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Our analysis reveals several findings. For content creation, we find no significant negative effects across all specifications, indicating that increased translated content does not displace native content production. For content consumption, we observe a positive and significant relationship between translated and non-translated article pageviews in the fixed effects model (columns 1-2), suggesting complementarity rather than substitution. The difference-in-differences analysis (columns 3-4) shows no significant negative effect of Google Translate integration on non-translated article pageviews. When examining combined normalized pageviews of source and target articles, we observed an increase from a baseline of 1.0 to approximately 1.5, indicating that translations attract new readers while maintaining

the readership of original content.

These findings demonstrate that the integration of Google Translate expands Wikipedia's content ecosystem rather than creating displacement effects. Machine translation appears to complement rather than substitute for native content creation and consumption, contributing to the overall growth of Wikipedia's knowledge base across languages.

F Increased Participation vs. Editor productivity

We examine individual editor contributions following the integration of Google Translate and examine the dual forces of participation and productivity among Wikipedia editors. This analysis aims to dissect the increase in content production and discern whether it stems from a growth in the editor base or improved editor productivity.

For this investigation, we collect editor IDs for each article created using the Content Translation system in our sample, enabling us to pinpoint the respective editors behind translated articles[8]. We used this information to create a language-month panel and calculate the number of new editors who used translation to create an article for the first time, the number of unique active editors contributing at least one article in a specific language and month, and the productivity of each editor, measured as the total articles created in the given month. For simplicity and robustness, we present results from analysis focused on language editions that integrated Google Translate in January 2019 and utilize two-way fixed effect models to analyze the editor panels in Table A.11.

As detailed in Table A.11, our analysis reveal a 49% growth in monthly new editors (an increase of 2.1 new editors beyond the pre-integration baseline of 4.3), a 66% increase in monthly active editors (an increase of 6.3 from the pre-integration baseline of 9.5), and a 69% increase in number of new translated articles per editor[9] (an increase of 4.5 from the pre-integration baseline of 6.5). Therefore, the magnitude of these effects indicate that the increase in translation volume stems from both existing editors becoming more active and productive as well as an expansion of the editor base. This shift towards greater editor participation and productivity underscores Google Translate's role in empowering Wikipedia editors.

Moreover, Figure A.22 presents the temporal trend of editor productivity, providing a model-free and intuitive visual substantiation to our econometric estimates. The figure

[8]For privacy consideration, only editor ID information is accessible on the website.

[9]Model (3) is analyzed at the editor level, as opposed to target language, and hence has a different number of observations

Table A.11: Editor Engagement and Productivity Analysis Post-Google Translate Integration

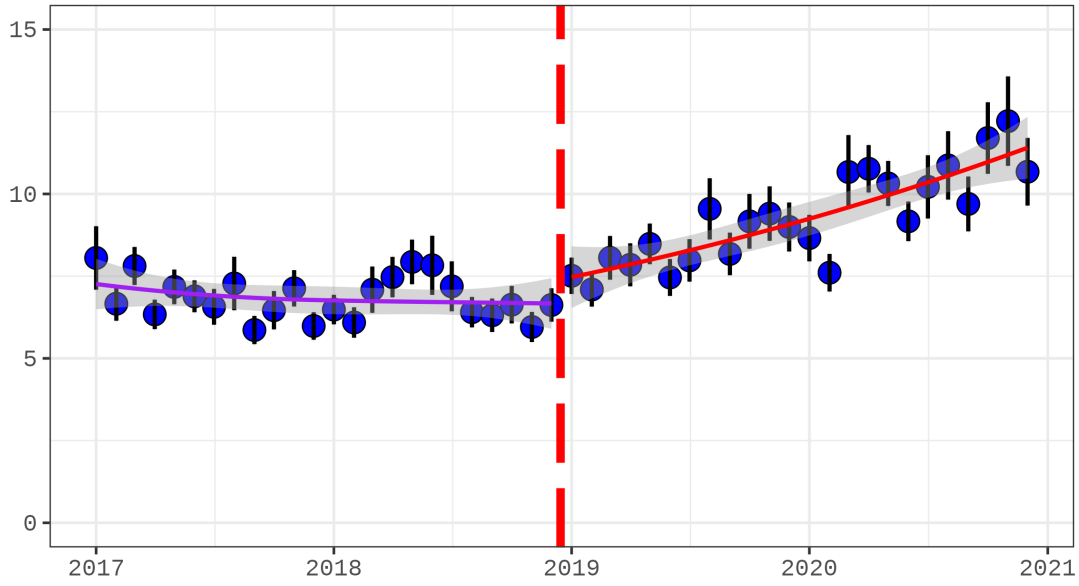
Dependent Variables:	New Editor	Active Editor	Article/Editor
Model:	(1)	(2)	(3)
GT Integration	2.1***	6.3***	4.5***
	(0.6)	(1.5)	(1.3)
Fixed-effects			
Language	Yes	Yes	Yes
Year-Month	Yes	Yes	Yes
Fit statistics			
Observations	6,336	6,336	3,979
R^2	0.83264	0.91247	0.13120
Pre-integration Mean	4.3	9.5	6.5

Clustered (Language) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

illustrates a steady temporal pattern in the monthly number of new translated articles per editor leading up to January 2019, after which there is a discernible upward trend. Individual editor productivity increases from an average of 6.8 article per editor/month in the pre-integration period to around 12 articles toward the end of our sample period. The error bars suggests variability in the data, yet the fitted LOESS curve shows that, on average, editors have been producing more articles each month since the integration of Google Translate.

Interestingly, the Interrupted Time Series analysis reveals a modest 10% increase in the level at the time of interruption for editor productivity ($\beta_2 = 0.7$ more articles relative to a 6.8 article pre-interruption baseline). However, there is a substantive positive shift between the post-interruption slope and pre-interruption slope ($\beta_3 = 0.19$ articles per editor/month). These findings imply that the full extent of the enhancement unfolds gradually as users become more adept with using the technology. Hence, while the tool lowers the skill barrier for editors to contribute, it should be complemented with strategic efforts to improve editor skill and experience with the new technology.

Figure A.22: Evolution of Editor Productivity on Wikipedia Over Time



Notes: This graph presents the trends in monthly number of new translated articles per editor on Wikipedia, illustrating a steady rise in productivity following the rollout of Google Translate in 2019. The fitted lines and shaded areas result from a piecewise LOESS regression.