Who Decides What Stays and Why? Participation, Authority, and Rationales in Chinese Wikipedia's Deletion Discussions

Tianyu Gao
Graduate Center, City University of New York
New York, NY, USA
tgao@gc.cuny.edu

Ping Ji
Hunter College & GC, City University of New York
New York, NY, USA
ping.ji@hunter.cuny.edu

Abstract

Chinese Wikipedia (zh-wiki), the largest Chinese-language encyclopedia not under state control, decides whether to keep or delete articles through open deliberations: editors argue by citing policies, and a *closer*, an editor who reviews the full discussion and renders a binding decision on the article's fate. But who actually participates in these discussions, how concentrated is the decision-making authority, what topics are contested, and what rationales drive the outcomes? We analyze 47,030 Articles for Deletion (AfD) discussion cases from Chinese Wikipedia spanning 2020 to 2024. We find that while hundreds of editors participate each year, closing authority is concentrated in a small and shrinking pool of closers, with a single individual responsible for up to 89.9% of substantive decisions in one year. Using diagnostic Unicode characters to classify editors by *script* preference (whether they write in traditional or simplified Chinese characters, a signal of regional origin and, by extension, whether they operate under internet censorship), we find that traditional-script editors outnumber simplified-script editors by 1.2–1.8:1, even though simplified Chinese speakers outnumber traditional Chinese speakers by roughly 45:1 globally. Yet deletion outcomes show no consistent directional bias favoring either community. Biography, media, list, and geography articles dominate nominations. Notability is the dominant rationale across all outcome types, and over a quarter of deletions contain no explicit rationale. These findings characterize a governance system where broad community participation coexists with highly concentrated decision-making authority, raising questions about the relationship between deliberative process design and the distribution of actual outcomes.

Keywords

content moderation, platform governance, Chinese Wikipedia, Wikipedia Articles for Deletion, participation inequality

1 Introduction

Over a billion people read and write Chinese, yet the largest Chinese-language encyclopedia not under state control is governed by a few hundred volunteer editors. Chinese Wikipedia (zh-wiki), the Chinese-language edition of Wikipedia hosted by the Wikimedia Foundation, operates without the editorial oversight of platforms like Baidu Baike (百度百科, "Baidu Encyclopedia"), a widely used alternative subject to Chinese government content regulations [8]. Instead, zh-wiki's content decisions are made through open deliberative processes among its editors. Yet zh-wiki's editor community operates under an unusual constraint: access from mainland China has been intermittently restricted since 2004 and fully blocked since 2015 (Section 2.2), meaning that participation is shaped by differential access across regions. Who shows up to edit, who takes on governance roles, and how decisions are made are all empirical questions, and none have been studied at scale in the context of deletion governance.

Among zh-wiki's governance processes, Articles for Deletion discussion (AfD discussion; 頁面存廢討論) is where content inclusion decisions are publicly debated: editors debate whether a nominated article should survive, and a *closer* (an editor, typically an administrator, who reads the full discussion and renders a final judgment) decides the outcome. Although every Wikipedia language edition maintains some form of AfD, each community defines its own rules for discussion periods, who may close, and how consensus is assessed, making deletion governance a locally shaped process rather than a uniform global one. Unlike corporate content moderation, which relies on automated filters or paid moderators, AfD discussion is deliberative, public, and argument-based: the closer evaluates the quality of policy arguments rather than counting votes (Section 2.1). This makes it a rich site for studying governance in practice: every deliberation, participant, and decision is recorded and empirically traceable.

Prior studies of Wikipedia's AfD have focused almost exclusively on English Wikipedia (Section 3).

In this paper, we present a quantitative analysis of all 47,030 zh-wiki AfD discussion cases from 2020 to 2024, a complete census of every deletion discussion in that period. Written Chinese has two script systems (Section 2.2): *simplified* characters, mandated in education and official use in mainland China (~1.4 billion people), and *traditional* characters, standard in Taiwan, Hong Kong, and Macau (~30 million). Because each region's school system teaches only one script, an editor's character choices are a reliable indicator of where they were educated and, by extension, which side of China's internet censorship regime they live on. Mainland China—home to nearly all simplified-script users—has blocked access to zh-wiki since at least 2015, with the block expanding to all Wikipedia language editions in April 2019. The resulting demographic inversion is notable: despite a 45:1 population ratio favoring simplified-script speakers, traditional-script editors outnumber simplified on zh-wiki by 1.2–1.8:1, and Taiwan and Hong Kong together account for nearly half of zh-wiki's page views [14]. This raises the question of whether the imbalance in who participates affects which articles survive. We address three research questions:

Free and Open Communications on the Internet 2026(2), 52–61

- **RQ1: Who discusses, who decides, and who dictates?** How concentrated is discussion participation? How concentrated is closing authority? Does the closer follow or override community votes?
- **RQ2: What gets deleted?** What topics are nominated, how do deletion rates vary, and do editors from different script communities (traditional vs. simplified Chinese; Section 2.2) target different content?
- **RQ3: Why are articles deleted?** What rationales dominate AfD discussions across different outcome types?

We find that closing authority is held by a small and shrinking pool of closers, that notability concerns drive outcomes across all categories, and that despite demographic imbalances in participation, deletion outcomes do not systematically favor either script community: articles nominated by simplified-script editors are not disproportionately kept, nor are traditional-script editors' nominations disproportionately deleted.

2 Background

2.1 The AfD Discussion Process

Chinese Wikipedia provides several mechanisms for content removal. Speedy deletion (快速刪除) handles clear-cut policy violations, and a copyright review track (疑似侵權), housed within the AfD system but focused on verifying infringement rather than debating inclusion, handles suspected copyright infringement. Both are primarily administrative: an administrator reviews and decides without the extended community deliberation that characterizes standard AfD.

When an article's inclusion is genuinely disputed, the community resolves it through **AfD discussion** (頁面存廢討論, literally "page survival discussion"), [1] a deliberative process in which editors write arguments, respond to each other, cite policies, and present evidence over a period of days, after which an administrator reads the entire discussion and renders a judgment. This is the process we study. It proceeds in three stages:

Nomination. Any autoconfirmed user (editors who meet minimum account age and edit count thresholds) may nominate an article by placing a deletion template on the page and opening a discussion on the daily AfD log page. The nominator must provide a rationale, typically citing *notability* concerns (whether the article's subject has received significant coverage in independent, reliable sources, a core inclusion criterion across all Wikipedia language editions[2]), lack of reliable sources, or advertising, which frames the subsequent discussion.

Discussion. Discussions run for a minimum of seven days. Editors argue for deletion (刪除), keep (保留), merging (合併), or redirection (重定向) using structured vote templates, but are required to provide substantive reasoning. The policy explicitly states that "voting cannot replace discussion" (投票不能代替討論),[3] meaning that a well-reasoned policy argument from one editor can outweigh multiple opposing votes that lack policy basis. In practice, discussions range from a single comment to dozens of editors debating source reliability, article scope, naming conventions, and policy interpretation.

Closing. After the discussion period, an administrator closes the case by evaluating the policy-based arguments presented, not by counting votes. This gives the closer substantial discretion; the same set of editor comments can lead to different outcomes depending on which administrator closes the case. Possible outcomes include:

- **Delete** (刪除): consensus that the article fails deletion criteria (usually notability).
- **Keep** (保留): consensus supports retention, or issues were addressed during discussion.
- **No consensus** (無共識): editors actively debated but the closer cannot determine a clear policy-based consensus. The article is retained.
- **Merge / Redirect / Move** (合併 / 重定向 / 移動): article content merged into another article, replaced with a redirect, or moved to a different title or namespace.
- **Request removed** (撤回): the nominator withdraws their nomination, typically because the article was improved during discussion or the nominator was persuaded by counterarguments.
- **Relist** (重新提交): too few editors participated to form any consensus; the case is reopened on a future day's log page for another round. The relist template has existed since 2018 but was rarely used until mid-2022, when Shizhao, the most active closer (Section 5.1.2), began applying it systematically. An article may be relisted up to three times per policy. In practice, 95.7% of relist entries are bare template stubs with no new discussion; they are administrative deferrals, not extensions of deliberation.

Not all closings require administrator privileges. Outcomes that involve page deletion require admin rights, but outcomes like keep, no consensus, and request removed can be executed by any experienced editor who did not participate in the nomination or discussion.[4] We use the term *closer* for any editor who closes a discussion (whether or not they hold admin rights) and *administrator* specifically for users who have been granted the technical ability to delete pages via a community-run Request for Adminship (RfA) process, which requires a supermajority (65–75%) of community support.[5] In practice, the vast majority of AfD closings are performed by administrators, since most cases end in deletion. Appendix A shows a complete real AfD discussion case with English translations.

2.2 Traditional vs. Simplified Chinese and the Access Divide

Written Chinese uses two script systems, the result of a mid-twentieth-century writing reform. **Traditional Chinese** (繁體中文) predates the reform and remains the standard in Taiwan, Hong Kong, and Macau (~30 million people). **Simplified Chinese** (简体中文) was introduced by the People's Republic of China in the 1950s to promote literacy; it is now the standard in mainland China, Singapore, and Malaysia (~1.4 billion people). The two scripts encode the same spoken language but use different written forms for thousands of words (e.g., 國 vs. 国 for "country"), analogous to British vs. American English spelling but with deeper visual differences and much stronger regional associations. An editor's script

[1] https://zh.wikipedia.org/wiki/Wikipedia:頁面存廢討論
[2] https://zh.wikipedia.org/wiki/Wikipedia:注度
[3] https://zh.wikipedia.org/wiki/Wikipedia:投票不能代替討論
[4] https://zh.wikipedia.org/wiki/Wikipedia:關閉存廢討論指引
[5] https://zh.wikipedia.org/wiki/Wikipedia:申成管理

choice reliably indicates whether they are from a traditional-script region or a simplified-script region, and by extension, whether they access the internet under censorship or without it.

Zh-wiki employs an automatic content conversion system that renders article text in the reader's preferred script. However, this conversion applies only to article content. Editor-written discussion text in AfD debates retains the original script chosen by its author, making it a visible marker of each participant's script community.

China first blocked Chinese Wikipedia in June 2004 [11], and access was intermittently restricted over the following decade [3]. In May 2015, Chinese Wikipedia was fully blocked, with the block expanding to all Wikipedia language editions in April 2019 [1, 10]. As a result, zh-wiki's participation skews heavily toward traditional-script regions: Taiwan accounts for 31.6% and Hong Kong for 15.3% of page views [14], and a 2010–2011 editor survey found these two regions contributing 64% of edits vs. mainland China's 17.7% [15]. Simplified-script editors who do participate likely rely on circumvention tools (e.g., VPNs) to bypass the block, adding a self-selection barrier that further shapes the community's composition. Because AfD discussion text retains the author's original script, we use it as a demographic signal to study participation and outcomes across these two communities (Section 4.3).

3 Related Work

Wikipedia governance and participation inequality. Wikipedia's governance has been extensively studied in the English-language edition. Geiger and Ford [2] documented power-law participation patterns in Wikipedia's deletion processes, showing that a small fraction of editors produce most deliberative content. Halfaker et al. [5] traced the long-term decline of Wikipedia's active editor base. TeBlunthuis et al. [13] studied how governance structures affect editor retention. These studies establish that Wikipedia governance concentrates over time, but they focus on English Wikipedia and do not examine the specific role of closing administrators.

Modeling and analyzing AfD discussions. Several studies have examined the deliberative dynamics of Wikipedia's AfD discussions. Mayfield and Black [9] released a structured corpus of over 400,000 English Wikipedia deletion debates with timestamped votes, comments, and outcomes, and demonstrated that policy citation patterns can predict deletion outcomes. Huq and Ciampaglia [6] characterized opinion dynamics and group decision-making in Wikipedia content discussions, studying how editor positions evolve during deliberation. Kaffee et al. [7] created a multilingual dataset of AfD editor stances and developed models to jointly predict stance and policy rationale, finding that the rate at which editors explicitly cite Wikipedia policies (e.g., notability guidelines, verifiability requirements) in their deletion arguments varies widely across language editions, from 20% of English comments to as low as 2% in German and Turkish. Yam [16] applied discourse analysis to AfD discussions, conceptualizing them as language games with specific power dynamics in negotiating knowledge boundaries. These studies treat AfD as a site of deliberation and argumentation, but none examine who *closes* the discussions, how concentrated that authority is, or how closer discretion interacts with community votes.

Chinese Wikipedia and cross-strait politics. Research on Chinese Wikipedia has primarily focused on content disputes related to the political tensions between mainland China and Taiwan (the "cross-strait" relationship). Liao and Petzold [8] used cartographic and network analysis to study linguistic diversity on zh-wiki, revealing regional patterns in editorial activity. Zhang et al. [17] studied the effect of censorship shocks on Chinese Wikipedia, showing that China's blocking events reduced contributions from mainland editors while increasing contributions from editors outside the censored region. However, no prior work has examined the governance structure of zh-wiki's deletion process: who closes discussions and how concentrated that authority is.

Content moderation labor and design. Roberts [12] documented how commercial content moderation labor concentrates in small, often unrepresentative groups. Grimmelmann [4] analyzed the "virtues of moderation" in online communities, arguing that moderation system design shapes outcomes independently of moderator intent. Our findings extend these arguments to volunteer-driven, argument-based moderation systems where a single moderator's judgment can shape the majority of outcomes.

Gap. Prior work has studied AfD as a deliberative process, modeled its outcomes, and analyzed its argumentation patterns, but almost exclusively in English Wikipedia. No study has quantified who closes AfD discussions on Chinese Wikipedia, how concentrated that authority is, whether closers follow or override community votes, or whether outcomes differ by the script demographics of participants. This paper addresses all four questions.

4 Dataset and Methodology

4.1 Data Collection and Pipeline Overview

We collected the complete AfD discussion archives of Chinese Wikipedia from January 2020 through December 2024. We chose 2020 as the start date because it is the first full calendar year after the April 2019 expansion of China's block to all Wikipedia language editions [1], which further altered zh-wiki's editor demographics (Section 2.2). Studying the post-block period ensures a consistent access regime throughout the dataset.

Zh-wiki organizes AfD discussions as daily log pages, one per calendar date (e.g., the page for January 1, 2024 is titled 頁面存廢討論/記錄/2024/01/01), each containing all cases nominated on that date as structured wikitext (Wikipedia's markup language). We used the MediaWiki API's `action=query` endpoint with the `revisions` property to retrieve the full wikitext of each daily page, iterating over every date in the five-year range. Individual AfD cases appear as level-2 or level-3 section headings within each daily page; we split the raw wikitext at these headings to extract one file per case, preserving the complete discussion thread including the nomination rationale, all editor comments, and the closing statement. In total we retrieved 1,806 daily pages covering every day in the five-year period with no missing dates, yielding 47,030 individual case files (Table 4). From each case file we extract structured fields (Section 4.2), classify each editor's script preference (Section 4.3), and assign a topic category to each nominated article (Section 4.4).

4.2 Case Parsing

Each AfD discussion case is stored as a raw wikitext file containing structured markup from which we extract several fields (see Appendix A for a complete annotated example with English translations):

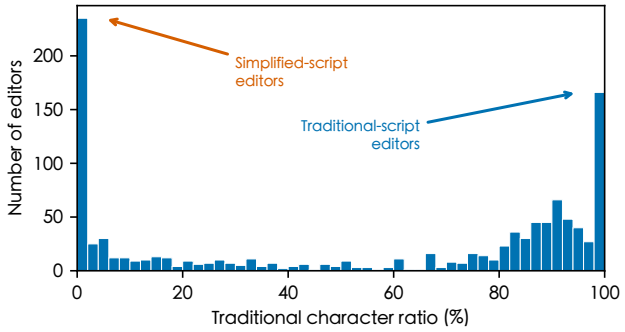

Figure 1: Distribution of traditional character ratios among 3,477 classifiable editors (2020–2024 combined). The strong bimodality validates the binary classification.

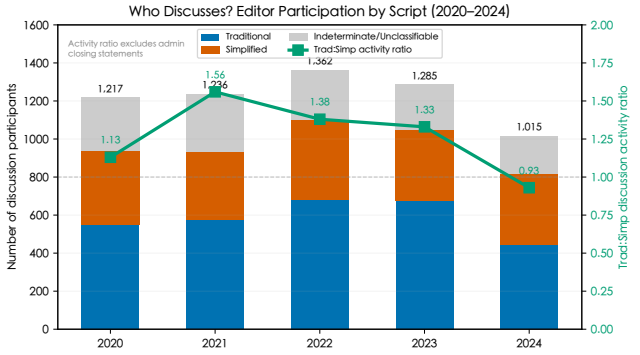

Figure 2: AfD discussion community demographics (2020–2024).

case title, nominator, nomination rationale, individual editor votes and reasoning, timestamps, closer identity, closing decision, and rationale. Each vote is attributed to a user via their wiki signature, a user link (`[[User:Name|...]]`) paired with a UTC timestamp. We aggregate fine-grained outcome codes into seven categories: delete, keep, no consensus, merge, redirect, move, and relist.

4.3 Editor Script Classification

We classified each editor's script preference based on *diagnostic characters*: Unicode codepoints that exist only in one script (Section 2.2). We compiled 648 traditional-only and 645 simplified-only diagnostic codepoints from standard conversion tables. For each editor, we attributed discussion lines via signatures, stripped boilerplate content that does not reflect the editor's own script choice (Appendix B), and counted diagnostic characters. Administrator closing statements are stripped so that activity metrics reflect discussion participation rather than administrative actions.

Editors with ≥3 diagnostic characters were classified as *traditional* (≥50% trad) or *simplified* (<50%); those with fewer were labeled *indeterminate* or *unclassifiable*. The classification is strongly bimodal (Figure 1): of 3,477 classifiable editors, 1,866 have a traditional ratio above 90% and 1,172 below 10%, with only 22 in the 45–55% range.

4.4 Topic Classification

We classified each case into one of 12 topic categories (biography, sports, media, geography, organization, event, science, politics, culture, list, internet, other) using a two-pass LLM approach (Claude 3.5 Haiku via Google Cloud Vertex AI). Pass 1 classified namespace-stripped titles in batches of 40; Pass 2 re-classified remaining "other" cases using up to 400 characters of discussion context alongside the title, in batches of 15 (full prompts in Appendix C). The two-pass approach reduced the "other" rate from 25% to 9.2%, rescuing ~7,400 cases into substantive categories. Manual validation by the first author on a stratified random sample of 60 cases (5 per category) yielded 95% agreement with the LLM labels.

5 Findings

5.1 RQ1: Who Discusses, Who Decides, and Who Dictates

5.1.1 Who discusses? Figure 2 shows the AfD discussion community over five years. Traditional-script editors consistently outnumber simplified-script editors by headcount (1.2–1.8:1), an inversion of the 45:1 population ratio that favors simplified-script speakers globally, driven by the Great Firewall (China's internet censorship infrastructure) blocking mainland Chinese users from accessing zh-wiki (Section 2.2). By discussion activity (excluding admin closing statements), the traditional advantage is moderate and stable (1.1–1.6:1 in 2020–2023), dropping to near parity (0.93:1) in 2024. The community is shrinking: total participants declined 25% from 1,362 (2022) to 1,015 (2024). Excluding bare relist stubs, the typical AfD discussion is sparsely attended: the mean and median number of participants per case are both 2.0 (IQR: 1–3), and 31% of cases have only a single participant (typically just the nominator).

Participation follows a steep power law (Figure 3): over 2,400 editors participated in only one AfD discussion across the five-year period, while only 27 participated in more than 500. The top 1% account for 31–40% of all case participations. Traditional-script editors account for a larger share at every participation level, but the gap narrows among the most active editors, where simplified-script power users are proportionally better represented.

5.1.2 Who decides? Table 1 reveals a higher degree of concentration at the decision-making level. Six administrators account for 85–100% of all substantive AfD decisions (excluding relists) across the study period. In 2023, a single administrator, Shizhao (a simplified-script user), closed 6,312 cases (17.3 per day), representing 89.9% of all substantive decisions. By 2024, Shizhao and Manchiu (a traditional-script user) together account for 86.5%.

The closer pool is not stable: of these six administrators, three (蟲蟲飛, Lanwi1, AT) left the role between 2021 and 2023, while Manchiu was inactive for two years (2022–2023) before returning in 2024. The two dominant closers use opposite scripts—Shizhao writes in simplified Chinese (associated with mainland China) and Manchiu in traditional Chinese (associated with Taiwan/Hong Kong)—meaning the script composition of closing authority oscillates based on who is more active: traditional-dominant when Manchiu led (2020–2021), simplified-dominant when Shizhao led

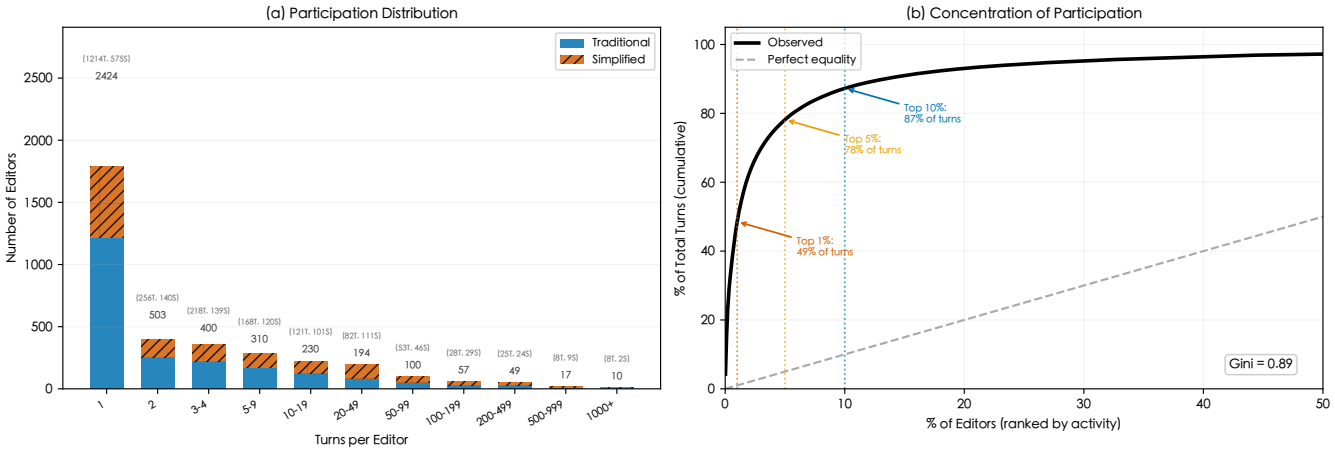

Figure 3: (a) Editor counts by participation level, colored by script preference, with per-bucket traditional/simplified breakdowns. The steep drop from one-time participants (2,424) to power editors illustrates the power-law distribution. (b) Lorenz curve showing concentration of participation: the top 1% of editors account for 49% of all discussion turns (Gini = 0.89).

Table 1: Substantive AfD closings by administrator (excluding relists). — = inactive that year.

Closer	Script	2020	2021	2022	2023	2024
Shizhao	simp	—	908	2,830	6,312	2,521
Manchiu	trad	1,996	1,647	—	—	1,239
AT	trad	256	933	1,197	—	—
蟲蟲飛	trad	1,268	879	—	—	—
Lanwi1	trad	1,035	1,084	—	—	—
淺藍雪	simp	—	—	937	312	—
All others		249	900	598	98	583
Total subst.		4,804	5,451	5,562	7,022	4,343
Top-2 share		67.9%	46.8%	72.4%	89.9%	86.5%

(2022–2023), near parity in 2024. Governance demographics are a function of individual activity levels, not a stable community property.

5.1.3 Who dictates? Zh-wiki policy states that closers evaluate argument quality, not vote counts. Figure 4 measures this discretion by comparing editor vote majorities against actual closing outcomes in 5,920 cases with ≥2 identifiable votes.

Closers align with the vote majority 57–67% of the time. Explicit overrides, where the closer chose a different outcome than the vote majority, range from 4.2% (2024) to 13.3% (2021). Overrides are directionally asymmetric: the dominant pattern is **keep→delete** (editors vote to keep, but the closer deletes: 196 cases across all years, vs. 104 in the reverse direction). The 2021 spike (102 keep→delete overrides) is largely attributable to Lanwi1, a closer with an 89.6% deletion rate. By 2023–2024, overrides are rare and roughly balanced.

The two current dominant closers override at similar rates (~7%) but in **opposite directions**: Shizhao leans toward deleting when editors vote keep (2:1 keep→delete), while Manchiu leans toward keeping when editors vote delete (4:1 delete→keep). Each closer's individual judgment thus introduces a directional tendency into the system. Since these two individuals process 86.5% of all decisions, their personal tendencies shape the aggregate outcome distribution.

5.1.4 Do deletion outcomes differ by script community? Despite the concentration documented above, we find no consistent pattern of one script community being disadvantaged in deletion outcomes.

At the aggregate level, simplified-script nominators' targets are deleted at equal or higher rates than traditional-script nominators' targets in every year (+0.2 to +9.1 percentage points; Table 5), suggesting simplified nominators are at least as effective at securing deletions. At the topic level, gaps exist in both directions (Section 5.2). We also classified each article's title by applying the same diagnostic-character method used for editors (Section 4.3): titles containing traditional-only characters are labeled "traditional-title," those with simplified-only characters "simplified-title," and those with no diagnostic characters (e.g., titles in English or using characters shared by both scripts) "neutral." Traditional-title articles are deleted at lower rates than simplified-title articles in 2020–2022 (Table 8). Whether this reflects differences in source availability, topic selection, or other factors cannot be determined from our data alone.

This pattern holds regardless of which closer dominated in a given year. We also tested whether closers disproportionately override one script community using two-sided Fisher's exact tests on 2×2 contingency tables (overridden vs. not × traditional-majority vs. simplified-majority). At first glance, overridden vote majorities are slightly more traditional-leaning than non-overridden ones (p=0.03). However, this reflects a compositional confound: keep-voters are more traditional-leaning (62.8% trad) than delete-voters (49.3%), and the most common override direction is keep→delete. When we compare within each direction separately, the script composition of overridden and non-overridden groups does not differ (p=0.45 for keep→delete, p=0.20 for delete→keep). That is, closers do not appear to weigh arguments differently based on the speaker's script: within each override direction, the probability of being overridden is independent of the vote majority's script composition.

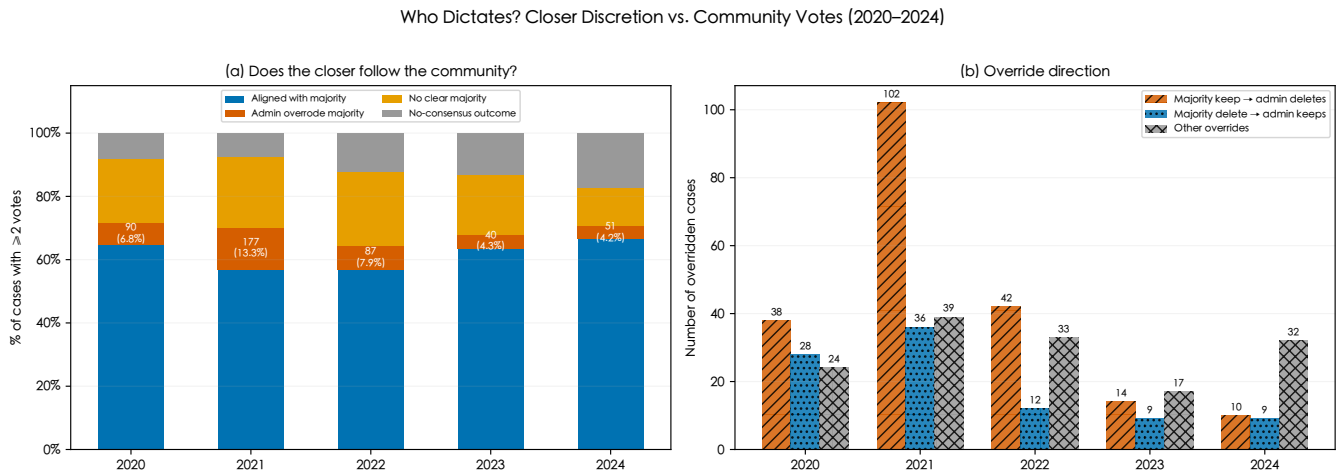

Figure 4: Closer override analysis. (a) Alignment between vote majority and closing outcome; legend ordered bottom-to-top to match the stacked bars. (b) Direction of explicit overrides; bars ordered left-to-right to match the legend. Both panels use hatching patterns for accessibility.

5.2 RQ2: What Gets Deleted

5.2.1 Topic distribution and deletion rates. Figure 5(a) shows the distribution of AfD nominations across 12 LLM-classified topic categories. Biography dominates in volume (11,405 cases, 24.3%) but has a mid-range deletion rate (78.7%). List articles and sports have the highest deletion rates (84.8% and 84.2%), reflecting clear-cut notability failures. At the other end, science (73.7%) and event (73.2%) are hardest to delete (Table 6). The "other" category (Wikipedia-internal pages) has the highest deletion rate of all (93.5%).

5.2.2 Nomination patterns by script community. Figure 5(b) compares deletion rates by the nominator's script preference across topics. Traditional and simplified nominators nominate nearly the same topic mix: biography accounts for ~30% of nominations from both groups, and no topic differs by more than 2 percentage points.

However, deletion rates differ by nominator script within individual topics. Simplified nominators' targets are deleted more often in biography (+8.7 pp), internet (+10.5 pp), and organization (+5.8 pp), while traditional nominators' targets are deleted more in culture (+11.9 pp) and geography (+4.1 pp). These gaps likely reflect differences in the specific articles nominated rather than closer prejudice: the direction of the gap reverses across topics (simplified nominators' targets deleted more often in biography, internet, and organization; traditional nominators' targets deleted more in culture and geography), which is inconsistent with a systematic preference by closers for either script community.

5.3 RQ3: Why Articles Are Deleted

5.3.1 Notability as the dominant rationale. We extracted rationales from AfD discussion text using keyword matching across 17 policy-related patterns. Table 2 shows the rationale distribution for each outcome type. A single case may match multiple rationales; percentages can sum over 100%.

Notability (關注度/关注度) is the dominant rationale across every outcome category, but its role shifts depending on the result:

- **In deletions** (n=23,272): notability is cited in 23.2% of cases and no reliable sources in 12.5%. Advertising (8.0%) and outdated/deprecated content (7.1%) form a second tier. Notably, 25.7% of deleted cases contain **no explicit rationale**: the closer wrote a bare "删除。" ("Delete.") with no explanation.
- **In keeps** (n=7,247): notability was *raised but successfully rebutted* in 35.3% of cases. Winning keep arguments include sources found during discussion (7.9%), article can be improved (9.4%), and other-language Wikipedia has the article (6.6%). AfD thus functions partly as a **source-discovery mechanism**.
- **In no-consensus** (n=1,922): notability is *disputed* in 39.8%, the highest rate of any outcome, confirming these are genuinely ambiguous cases.
- **In merges/redirects/moves** (n=2,049): merge target exists (50.0%) is the dominant driver; these are structural decisions about article organization rather than content disputes.

The pattern reveals that zh-wiki AfD discussion operates primarily as a **notability adjudication system**: the same question ("does this article have enough independent, reliable source coverage?") drives deletions, motivates keeps, produces deadlocks, and triggers merges.

6 Discussion

6.1 Many Voices, Few Judges

Zh-wiki's AfD system draws from a broad editor base, with traditional-script editors outnumbering simplified-script editors despite the much larger simplified-script population globally (Section 2.2). Yet while many editors participate in deliberation, the authority to close discussions and render final judgments rests with a small and shrinking pool of administrators. Of six administrators who each processed over 900 substantive closings during the study period, three departed between 2021 and 2023. The gap between who argues and who decides is the central structural feature of the system.

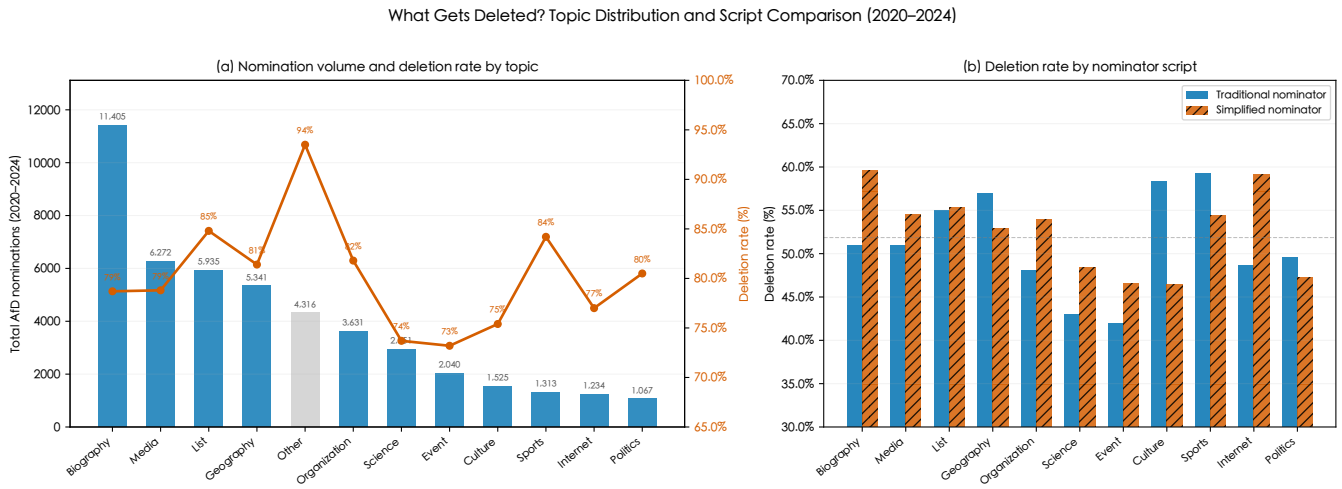

Figure 5: (a) Nomination volume and deletion rate by topic. (b) Deletion rate by nominator script preference across topics.

Table 2: Top rationales mentioned in AfD discussions, by outcome (% of cases). Each row represents a topic mentioned in discussion; in deletions, these are *grounds for removal*, while in keeps, the same concerns were *raised but successfully rebutted*. In no-consensus cases they remain *disputed*, and in merges they motivate *structural reorganization*. A case may match multiple rationales. "—" = not applicable.

Rationale	Del. (23,272)	Keep (7,247)	NC (1,922)	Merge (2,049)
No rationale given	25.7	—	—	—
Notability	23.2	35.3	39.8	27.7
Sources	12.5	17.1	23.6	11.4
Advertising	8.0	9.8	7.8	3.8
Outdated	7.1	—	8.4	3.4
Original research	4.0	6.2	5.1	—
Not encyclopedic	3.3	—	4.0	3.3
Duplicate	1.6	—	—	8.0
Stub/too short	1.6	—	—	—
Can improve	—	9.4	7.0	3.7
Merge target	—	4.2	6.6	50.0
Other wiki has art.	—	6.6	5.6	5.0

This concentration is deepening. The decline in unique closers (from 203 in 2022 to 127 in 2024) and in discussion participants (from 1,362 to 1,015) shows the community shrinking at both the deliberative and decision-making levels simultaneously. The relist mechanism adds to the workload: relisting grew from under 5% to 47.9% of all cases, increasing the volume of decisions without generating substantive new discussion. Because becoming an administrator requires passing a community-run Request for Adminship process (Section 2.1), and because the workload for active closers is substantial (the top closer averaged 17.3 closings per day in 2023), few editors take on the role.

Closer identity correlates with outcomes. In 2020, delete rates ranged from 53% to 90% depending on which administrator closed the case. The two most active closers override community vote majorities at similar rates (~7%) but in opposite directions: one leans toward deletion, the other toward retention. When 86–90% of decisions are concentrated in one or two individuals with measurably different deletion rates, aggregate outcome distributions become sensitive to which closers happen to be active in a given period.

Despite this concentration, neither our outcome-level nor our override-level analysis (Section 5.1.4) reveals consistent differential treatment by script community. However, the absence of detected bias does not constitute proven neutrality; the current balanced results depend on the specific individuals who are active and could shift if the closer pool changes.

6.2 Implications for Content Moderation Design

Our findings contribute to a growing literature on governance concentration in volunteer-driven content moderation systems [4, 12]. The gap between who *argues* and who *decides* is the key structural feature: participation inequality in discussions (Gini coefficient 0.83–0.86, where 1.0 represents perfect concentration) is unremarkable for online communities, but closer concentration (Table 1) is considerably more pronounced.

Notability as the organizing principle. Our rationale analysis (Section 5.3) reveals that a single policy question organizes the entire system: whether an article meets zh-wiki's sourcing threshold for inclusion. Notability is explicitly cited in 23.2% of deletions, 35.3% of keeps, and 39.8% of no-consensus outcomes; related rationales (lack of reliable sources, 12.5% of deletions; advertising, 8.0%; and substub status) also turn on the sourcing question. Even the 25.7% of deletions with no recorded rationale are consistent with this framing: closers who process dozens of cases per day may omit the reasoning, though we cannot determine the cause from our data. The same question drives all four outcome types in different ways: it motivates deletions, is rebutted in keeps (where editors find sources during discussion), produces deadlocks when sourcing is ambiguous, and triggers merges when content belongs elsewhere. A practical consequence is that AfD functions not only as a removal mechanism but also as an incentive for source discovery, improving the articles that survive.

Transparency gaps. Despite AfD's deliberative design, over a quarter of deleted cases contain no explicit rationale (Section 5.3). This practice, concentrated in 2020–2021, raises questions about accountability in an otherwise transparent system, regardless of whether outcomes are unbiased.

Burden of proof asymmetry. Keeping an article requires 1.2–2.5× the discussion effort of deleting one (Table 7), suggesting that AfD's design structurally favors deletion. Under-sourced topics, which may correlate with topics underrepresented in independent Chinese-language media (e.g., mainland Chinese topics that lack coverage outside state-controlled outlets), face a structural disadvantage independent of closer bias. The finding that keep-voters are disproportionately traditional-script users raises the question of whether this asymmetry reflects differential access to sources across regional media ecosystems.

6.3 Limitations

Script as demographic proxy. We use traditional Chinese as a proxy for Taiwan/Hong Kong origin and simplified Chinese as a proxy for mainland Chinese origin, but this mapping is imperfect. Some editors may use the non-dominant script of their region by personal preference, and editors in Singapore or Malaysia also use simplified Chinese. Our classification identifies script preference in AfD commentary, not geographic location or political affiliation.

Boilerplate contamination. Zh-wiki's AfD templates are written primarily in traditional Chinese, which could inflate an editor's traditional-character count if not removed. We stripped identified boilerplate (Appendix B) but residual contamination may remain. We validated against 20 editors whose script preference is publicly known from their user pages, confirming correct classification for all 20 after stripping (see Appendix B).

Deleted article content. When an article is deleted on Wikipedia, its full text is permanently removed from public view and accessible only to administrators. This means we cannot examine the actual content, quality, or sourcing of deleted articles, limiting our analysis to metadata (title, topic, discussion text) rather than the articles themselves.

Other limitations. Outcome detection relies on regex parsing of wiki markup. Our line-attribution heuristic fails to identify any participant in 19% of substantive cases (typically template-only nominations with no signed discussion), which are excluded from per-editor analyses. Closers may self-select into different case types. Rationale extraction is keyword-based and misses cases where no explicit rationale was given. This study is observational and cannot establish causal claims. The absence of consistent directional bias today does not guarantee its absence in the future.

7 Future Work

Cross-language comparison. Our methodology (parsing structured AfD markup, classifying editor demographics, and measuring closer concentration) is directly applicable to other Wikipedia language editions. A natural extension is to test whether the concentration patterns we observe are unique to zh-wiki or a general feature of Wikipedia's deletion governance. Prior work on other language editions has studied AfD as a deliberative process (modeling argumentation, stance prediction, and opinion dynamics; Section 3) but none have quantified closer concentration or override behavior, leaving the structural governance question open. Each Wikipedia language edition maintains its own deletion policies with subtle but consequential differences (varying thresholds for notability, different discussion periods, and distinct norms around closer discretion), making cross-language comparison not just a test of generalizability but an opportunity to characterize how policy design shapes governance outcomes in different language communities.

Other deliberative spaces on Wikipedia. AfD discussion is only one of several community processes on zh-wiki where editors deliberate and administrators decide. For example, requests for adminship (RfA), featured article reviews, move discussions, and arbitration cases all follow similar patterns: open community input followed by a closing judgment. RfA is particularly interesting because it determines *who becomes an administrator*, the very population whose concentration we document in AfD. Studying whether the same concentration patterns appear in the process that selects decision-makers would close a feedback loop in our analysis.

Pre- and post-censorship temporal analysis. Our dataset covers 2020–2024, entirely within the post-block period. Extending the analysis backward to include AfD discussions from before the 2015 full block and the 2019 all-language block would enable a comparison of governance patterns under different censorship regimes. If closer concentration, script demographics, or deletion rates shifted measurably around these blocking events, it would strengthen the link between access restrictions and governance outcomes.

Causal analysis. Our study documents longitudinal trends but cannot establish causation. Future work could exploit natural experiments, such as closer returns after periods of inactivity, to test whether outcome distributions shift measurably when the dominant closer changes.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback, which substantially improved the clarity and rigor of this paper. AI tools (Claude) were used to assist with generating figures and revising grammar; all writing, analysis, and intellectual contributions are solely the work of the authors.

References

[1] BBC News. Wikipedia blocked in China in all languages. https://www.bbc.com/news/technology-48269608, May 14, 2019.

[2] R. Stuart Geiger and Heather Ford. Participation in Wikipedia's article deletion processes. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11)*, pages 201–202. ACM, 2011.

[3] GreatFire.org. Wikipedia blocked in China. https://en.greatfire.org/wikipedia.org, 2019. Accessed: 2025-01-15.

[4] James Grimmelmann. The virtues of moderation. *Yale Journal of Law & Technology*, 17(1):42–109, 2015.

[5] Aaron Halfaker, R. Stuart Geiger, Jonathan T. Morgan, and John Riedl. The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *American Behavioral Scientist*, 57(5):664–688, 2013.

[6] Khandaker Tasnim Huq and Giovanni Luca Ciampaglia. Characterizing opinion dynamics and group decision making in Wikipedia content discussions. In *Companion Proceedings of the Web Conference 2021 (WWW '21 Companion)*, pages 597–601. ACM, 2021.

[7] Lucie-Aimée Kaffee, Arnav Arora, and Isabelle Augenstein. Why should this article be deleted? Transparent stance detection in multilingual Wikipedia editor

discussions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
[8] Han-Teng Liao and Thomas Petzold. Analysing geo-linguistic dynamics of the World Wide Web: The use of cartograms and network analysis to understand linguistic development in Wikipedia. *Cultural Science*, 3(2), 2010.
[9] Elijah Mayfield and Alan W. Black. Analyzing Wikipedia deletion debates with a group decision-making forecast model. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):Article 206, 2019.
[10] Open Observatory of Network Interference (OONI). The blocking of Wikipedia in China. https://ooni.org/post/2019-china-wikipedia-blocking/, May 4, 2019.
[11] Philip P. Pan. Reference tool on Web finds fans, censors. *The Washington Post*, February 20, 2006.
[12] Sarah T. Roberts. *Behind the Screen: Content Moderation in the Shadows of Social Media.* Yale University Press, 2019.
[13] Nathan TeBlunthuis, Aaron Shaw, and Benjamin Mako Hill. Revisiting "The Rise and Decline" in a population of peer production projects. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, Paper 355. ACM, 2018.
[14] Wikimedia Foundation. Wikimedia traffic analysis report — page views per country. https://stats.wikimedia.org, 2022. Accessed: 2025-01-15.
[15] Wikimedia Foundation. Wikipedia editor survey 2011. https://meta.wikimedia.org/wiki/Editor_Survey_2011, 2011.
[16] Shing-Chung Jonathan Yam. Negotiating boundaries of knowledge: Discourse analysis of Wikipedia's Articles for Deletion (AfD) discussion. *Communication and Critical/Cultural Studies*, 13(1):34–51, 2016.
[17] Ark Fangzhou Zhang, Danielle Livneh, Ceren Budak, Lionel P. Robert Jr., and Daniel M. Romero. Shocking the crowd: The effect of censorship shocks on Chinese Wikipedia. In *Proceedings of the 11th International AAAI Conference on Web and Social Media (ICWSM)*, 2017.

A Anatomy of an AfD Discussion

Figure 6 shows a complete real AfD discussion case with the original wikitext markup and English translations of each field.

B Editor Script Classification Details

Boilerplate stripping. The following template-generated content is removed before diagnostic character counting, as it does not reflect the editor's own script choice:

Table 3: Stripped boilerplate patterns.

Pattern	Source
請將新留言置於本提示下方	Relist notice (trad chars)
为了能更深入的讨论并获得明确的共识	Relist header (simp)
当处理完相关页面后	Admin consensus notice (simp)
若您对当前已形成的共识有不同意见	Follow-up notice (simp)
機器人發現被提刪頁面	Bot messages (trad)
理據：	Nomination template (trad)
提交的維基人及時間	Nomination header (trad)

We validated by manually checking 20 editors whose regional origin or script preference is publicly known from their user pages, confirming that all 20 were classified correctly after stripping.

Sensitivity analysis: The trad:simp ratio remains stable (±5%) across thresholds of ≥1, ≥2, and ≥3 diagnostic characters.

C Topic Classification Prompts

We classified each AfD case into one of 12 topic categories using Claude 3.5 Haiku via Google Cloud Vertex AI in two passes (Section 4.4). The prompts are reproduced below.

Pass 1 system prompt (title-only classification in batches of 40):

```
You are a topic classifier for Wikipedia article titles.
Given a numbered list of article titles (in any
language), classify each into exactly one category.

Categories: biography, sports, media, geography,
organization, event, science, politics, culture, list,
internet, other

Rules:
- Respond ONLY with a JSON array of strings, one
category per title, in the same order.
- Use "biography" for any article about a specific
person.
- Use "list" for articles that are clearly lists,
indexes, or timelines.
- Use "other" only if no category fits.
- Do not include any explanation, just the JSON array.
```

Pass 2 system prompt (re-classification of "other" cases in batches of 15, with expanded rules):

```
You are a topic classifier for Wikipedia article titles.
Given a numbered list of article titles (in any
language: English, Chinese, Russian, German), classify
each into exactly one category.

Categories: biography, sports, media, geography,
organization, event, science, politics, culture, list,
internet, other

Rules:
- Respond ONLY with a JSON array of strings, one
category per title, in the same order. The array length
MUST equal the number of titles.
- Titles about a specific person (any name in any
language) → "biography"
- Titles about sports teams, athletes, matches,
tournaments, leagues → "sports"
- Titles about films, TV shows, music, albums, songs,
books, anime, manga, video games → "media"
- Titles about places, cities, villages, buildings,
roads, stations, bridges → "geography"
- Titles about companies, schools, universities,
churches, political parties → "organization"
- Titles about battles, wars, elections, disasters,
ceremonies, protests → "event"
- Titles about scientific concepts, diseases,
medicines, technology, math → "science"
- Titles about political topics, laws, government,
military units → "politics"
- Titles about art, religion, language, traditions,
food, philosophy → "culture"
- Titles that are lists, timelines, indexes, categories
→ "list"
- Titles about websites, memes, social media, YouTubers,
online communities → "internet"
- Use "other" ONLY if it truly doesn't fit ANY category
above.
- IMPORTANT: Return EXACTLY as many labels as there are
titles.
```

In Pass 2, the user message included up to 400 characters of discussion context (nomination rationale and early comments) alongside each title.

D Supplementary Tables

Table 4: zh-wiki AfD dataset summary (2020–2024).

Year	Cases	Articles	Relist %	Cases/Art.
2020	7,051	6,684	5.3%	1.05
2021	9,685	9,258	4.3%	1.05
2022	8,595	8,377	2.4%	1.03
2023	12,157	8,091	35.0%	1.50
2024	9,542	4,961	47.9%	1.92
Total	**47,030**	—	—	—

Field	Raw Wikitext (Chinese)	English Translation
Outcome	{{delh|rr}}	rr = request removed (nomination withdrawn)
Nomination	{{删除}}理據：近乎substub，幾乎無實際內容	"Nearly a substub, almost no real content"
Nominator	AINH 2022年3月4日	AINH — Mar 4, 2022
Vote: keep	{{保留}}：有潛力成為完整條目 — 中文維基百科20021024	Zhongwen20021024: "Has potential to grow"
Reply	→ AINH：請「成為完整條目」後才來投票	→ AINH: "Come back after it becomes a real article"
Discussion	Easterlies: 标题有些狭隘，可考慮「乌克兰加入欧盟」	Easterlies: "Title too narrow, try Ukraine Joining the EU"
Discussion	→ Txkk:「加入欧盟进程」？	→ Txkk: "Should the title be EU Accession Process instead?"
Discussion	→ Ericliu: 其他條目怎麼命名？	→ Ericliu: "How are similar articles titled?"
Discussion	Txkk: 乌克兰语维基存废走完了，结果是保留	Txkk: "Ukrainian Wikipedia AfD finished — result: keep"
Vote: keep	{{保留}}：現已有兩百多字，提刪理由消失 — HTinC23	HTinC23: "Now 200+ chars, nom. reason gone"
Decision	保留：請求理由消失	Keep: nomination reason no longer applies
Rationale	頁面已非substub	Page is no longer a substub
Closer	LuciferianThomas 2022年3月27日	LuciferianThomas — Mar 27, 2022

Parsed: Title: 乌克兰申请加入欧洲联盟 (Ukraine EU Application) | Outcome: keep (request removed) | Closer: LuciferianThomas
Participants: 8 | Keep votes: 2 | Delete votes: 0 | Discussion-only comments: 6 | Duration: 23 days

Figure 6: Anatomy of an AfD discussion: 乌克兰申请加入欧洲联盟 (Ukraine EU Application), March 2022. The article was nominated as a substub (an article too short to provide useful content), improved during the 23-day discussion, and ultimately kept when the closer judged the nomination reason no longer applied. 8 participants, 2 keep votes, 0 delete votes.

Table 5: Deletion rate by nominator script (non-relist cases).

Year	Trad nom. del% (n)	Simp nom. del% (n)	Δ
2020	54.0 (3,588)	54.2 (3,041)	+0.2
2021	58.7 (5,615)	67.8 (3,569)	+9.1
2022	53.1 (4,432)	58.7 (3,619)	+5.6
2023	69.1 (3,547)	75.9 (4,542)	+6.8
2024	52.8 (2,519)	57.6 (2,471)	+4.8

Table 6: Outcomes by topic (2020–2024, substantive closings). Sorted by deletion rate.

Topic	n	Del %	Keep %	NC %
Internet	956	90.0	6.8	1.0
List	4,722	85.8	8.8	1.9
Sports	1,107	85.5	7.5	3.3
Geography	4,295	84.3	7.9	1.7
Organization	2,831	82.8	10.8	2.4
Politics	847	81.7	9.3	2.8
Biography	8,743	80.4	11.7	3.2
Media	4,802	79.8	11.9	2.4
Science	2,369	78.3	10.7	1.2
Culture	1,137	75.5	12.0	2.4
Event	1,647	73.5	10.7	7.6

Table 7: Discussion effort for deletions vs. keeps by topic.

Topic	Del sigs	Keep sigs	Sig ratio	Char ratio
Internet	7.1	12.5	1.8×	2.2×
Sports	5.2	8.9	1.7×	2.5×
Geography	5.2	7.8	1.5×	2.2×
Culture	7.5	10.9	1.5×	1.6×
Politics	10.0	14.3	1.4×	1.6×
Event	9.3	12.9	1.4×	1.5×
Biography	6.6	7.2	1.1×	1.2×

Table 8: Deletion rate by article title script (non-relist cases).

Year	Trad title del% (n)	Simp title del% (n)	Neutral del% (n)
2020	45.2 (994)	51.0 (1,014)	55.6 (4,941)
2021	41.3 (1,664)	52.8 (1,061)	67.7 (6,770)
2022	33.1 (1,853)	50.2 (868)	63.5 (5,512)
2023	57.7 (1,528)	54.8 (701)	78.4 (5,948)
2024	50.0 (1,311)	47.3 (1,084)	60.5 (2,731)