RESEARCH ARTICLE

Journal of Sociotechnical Inquiry
Volume 1, Issue 1, 2026, pp. 3–12

Who Page Protection Locks Out: Access Control, Attention and Editing on English Wikipedia, 2023 to 2025

Altynai Asanova[1,*] and Emilbek Toktogonov[2]

[1]Department of Journalism and Communication, Osh State University, Kyrgyzstan
[2]Osh State University, Kyrgyzstan

ABSTRACT

English Wikipedia describes itself as the encyclopedia anyone can edit, yet its software lets administrators restrict editing of a page to accounts that pass age and edit-count thresholds. This study examines that access control, called page protection, as a working system. It asks how protection is applied by level, duration and stated reason; how it relates to bursts of reader attention; and how editing by unregistered and new editors changes once a page is locked. The data are the complete public protection log for 2023 to 2025 (93,642 events, including 50,024 new edit protections of articles) and a stratified sample of 450 article protections with daily pageviews and revision histories. Semi-protection made up 82.9% of new article protections and was mostly temporary, while 84.5% of extended confirmed protections were indefinite and 85.0% cited contentious-topic or arbitration rules. About one in five protections (19.1%, pool-weighted) followed a pageview spike, and spike-linked protections were shorter, not more often indefinite. In the seven days after protection, article edits by unregistered editors fell from 1,688 to 3 and edits by accounts under four days old fell from 298 to 3. Measured against vandalism, protection works almost completely. Measured by who is turned away, the picture differs: 42.9% of unregistered edits made in the week before protection were never reverted, and 55.9% of the month's page views arrived while the page was locked. The study argues that protection should be evaluated on both measures.

Keywords: Wikipedia, page protection, access control, peer production, breaking news, anonymous editing, content moderation

1. Introduction

English Wikipedia invites every reader to edit, and most pages carry an Edit tab that works without an account. The same software also lets administrators switch that invitation off for a single page. This feature is called page protection: a setting that limits who may edit a page to user groups that meet thresholds of account age and edit count (Wikipedia contributors, 2026). Semi-protection blocks unregistered editors and accounts that are not yet autoconfirmed, which on English Wikipedia means at least 4 days since the first edit and at least 10 edits. Extended confirmed protection raises the bar to 30 days and 500 edits. Full protection leaves editing to administrators alone.

Protection is therefore a piece of software that turns a social judgement about trust into a rule enforced at the moment someone presses Save. It matters most when many people arrive at once. Breaking news drives a large share of both editing and reading on Wikipedia (Keegan et al., 2013), and such surges bring vandalism, edit wars and well-meant corrections at the same time. A lock applied during a news peak stops the first kind of edit. It also stops the others, and it does so for the readers who are most likely to notice something new or wrong.

Research on Wikipedia has studied newcomer decline, quality control and governance in depth (Halfaker et al., 2013; Forte et al., 2009), but protection itself has received little direct study. Hill and Shaw (2015) called it a missing dimension of Wikipedia research and warned that analyses of editing which ignore protection may mistake a software restriction for a lack of interest. Work on account requirements shows that blocking unregistered contributors deters much more good work than bad (Hill & Shaw, 2021). Wikipedia's own record of protection consists of the log, the request noticeboard and the revision history, and each of these records the disruption that prompted a lock more fully than the editing that the lock prevented. Work on Tor users shows that contributions from blocked groups can be similar in quality to those of permitted newcomers (Tran et al., 2020). These findings raise a direct question about page protection: when the software closes a page, what does it close out?

This article answers three research questions using public logs and traffic data for 2023 to 2025. RQ1 asks how English Wikipedia applies protection to articles, by level, duration and stated reason. RQ2 asks how protection relates to attention: how often it follows a spike in page views, and whether spike-linked protections are more often left in place indefinitely. RQ3 asks how editing by unregistered and new editors changes after protection, and how much of the editing that protection stops had been kept by the community before the lock.

The analysis treats protection as the object of study, not as background. It reads the protection log as a record of what the access control did, links each sampled protection to the page views before and after it, and classifies every edit around it by the editor's account status at the time. The main finding is a gap between two ways of judging the system. Judged by the vandalism it stops, protection is close to perfect: unregistered and very new editors almost vanish from protected articles. Judged by who it turns away, it removes a stream of edits that the community had largely kept, and it does so while most of the page's readers are present. The journal's interest in systems that are evaluated on one set of measures and experienced on

another fits this case closely.

2. Related work and framing

2.1 Openness, newcomers and the costs of control

Wikipedia's early growth rested on low barriers to entry. Anthony et al. (2009) found that anonymous one-time contributors, whom they called Good Samaritans, made some of the most reliable contributions in their sample. Growth slowed after 2007 (Suh et al., 2009). Halfaker et al. (2013) traced the decline of active editors to the community's own response to popularity: tools and rules built to fight vandalism also rejected good-faith newcomers, who then left. Reverts of newcomers' work reduce the chance that they keep editing (Halfaker et al., 2011). TeBlunthuis et al. (2018) found a similar rise-and-decline pattern across hundreds of wikis, which suggests that the pattern follows from how peer production projects manage growth.

Access controls are one form of this management. Hill and Shaw (2021) studied 136 wikis that suddenly required accounts to edit. Account requirements cut low-quality contributions, but most of the contributions they deterred were of higher quality. Tran et al. (2020) studied edits that reached Wikipedia from the Tor anonymity network despite a block and found them broadly similar to edits by unregistered and new editors. Both studies point to a trade-off that page protection shares: a rule that filters by account status filters out good work too. Ford and Wajcman (2017) make the wider point that Wikipedia's infrastructure, including its rules and tools, shapes who can take part, so that the claim that anyone can edit describes a possibility more than a practice.

2.2 Quality control and governance as sociotechnical work

Vandalism has been a constant pressure on Wikipedia. Priedhorsky et al. (2007) measured how often readers saw damaged articles and showed that most damage was repaired quickly. Shachaf and Hara (2010) described trolling behaviour that goes beyond simple vandalism. Geiger and Ribes (2010) followed the distributed work of tools, bots and people that led to the banning of one vandal, and showed that order on Wikipedia is maintained by heterogeneous assemblages more than by single decisions. Geiger and Halfaker (2013) found that when a major anti-vandalism bot went offline, people still caught vandalism but more slowly. Geiger (2017) argued that such algorithmic systems carry organisational values as well as technical functions.

Rules sit beside these tools. Butler et al. (2008) showed that Wikipedia's policies grew into a large and complex body that serves many roles. Forte and Bruckman (2008) and Forte et al. (2009) described governance that became more decentralised as the project grew, with local groups and administrators interpreting shared rules. Shaw and Hill (2014) found that peer production projects tend toward oligarchy as they age, with a small group holding formal power. Steinsson (2024) showed how disputes over ambiguous rules on English Wikipedia demobilised some groups of editors and changed the institution over time. Page protection connects these strands: it is a rule that only administrators can apply and that the software enforces on everyone else.

The study also draws on writing about content moderation on commercial platforms. Gillespie (2019) argued that moderation is a central part of what platforms do, not an occasional add-on. Gorwa et al. (2020) described how automated moderation shifts decisions into technical systems that are hard to see or contest. Protection is a simpler mechanism than the classifiers those authors study, but it shares their logic of acting on classes of users at scale. Star and Ruhleder (1996) noted that infrastructure becomes visible mainly when it fails or blocks someone. A locked Edit tab is a clear case: it is invisible to established editors and very visible to a first-time reader who finds an error.

2.3 Breaking news and attention

Breaking news is where these tensions peak. Keegan et al. (2013) showed that articles about breaking news events draw more editors and faster collaboration than other articles. Avieson (2019) described editors competing and coordinating to cover news as it happens. Twyman et al. (2017) traced how articles about Black Lives Matter events grew in bursts tied to offline events. Ciampaglia et al. (2015) found that bursts of reader attention to Wikipedia topics are followed by bursts of article creation, which links attention to production. García-Gavilanes et al. (2017) showed that attention to past events spikes again when related events occur. Kiene et al. (2016) studied a different community that survived a sudden surge of newcomers by leaning on existing norms and moderators, which suggests that surges can be absorbed without closing the door.

Conflict research adds a second reason for protection during news. Kittur et al. (2007) measured the growing share of Wikipedia work spent on conflict and coordination. Yasseri et al. (2012) modelled edit wars and found that some controversial articles stay contested for long periods. Protection is a standard response to both vandalism surges and disputes. Yet the literature offers little evidence on how often protection follows attention spikes, how long it lasts and who it stops. This study supplies that evidence for 2023 to 2025.

2.4 Framing: two ways to evaluate an access control

The analysis uses a simple frame. An access control can be evaluated by what it prevents and by whom it excludes. The first measure counts harms avoided, such as vandal edits that no longer need reverting. The second counts contributions that no longer arrive and the people who meet a closed door. Wikipedia's own processes mostly record the first: requests for protection cite vandalism or disruption, and success means the disruption stops. The second is largely unrecorded, because an edit that was never made leaves no trace in the revision history. This study estimates the second measure indirectly, using the edits made just before protection and the readers who arrived while the page was locked.

3. Data and methods

3.1 The protection log

Every protection action on English Wikipedia is written to a public log. The study downloaded the full protection log for 1 January 2023 to 31 December 2025 through the MediaWiki Action API (list=logevents, letype=protect), paging with continuation tokens, pausing between calls and using the maxlag parameter. The download returned 93,642 log events:

72,399 new protections, 9,859 modifications of existing protections, 3,939 removals and 7,445 records of protection carried over when a page was moved. Each event records the page, time, protection settings for editing, moving and page creation, expiry time and a free-text reason written by the administrator.

The main descriptive unit is a new edit protection of an article, meaning an event with action protect, namespace 0 and an edit restriction at the semi, extended confirmed or full level. There were 50,024 such events on 36,549 distinct articles. The analysis excludes 2,081 article protections that only blocked page creation, which is a separate use of the same tool, and five template-editor protections of articles. Protection of other namespaces is reported briefly for context.

3.2 Measures from the log

Level is the user group needed to edit: autoconfirmed for semi-protection, extended confirmed, or administrator for full protection. Set duration is the time between the protection and its stated expiry; protections without expiry are coded as indefinite. The effective end of a sampled protection is the earlier of its expiry and the next protection-log event for the same page, which captures early removal or a change of level.

Reasons were coded with documented keyword rules applied to the lower-cased reason text. Each rule looks for policy shortcuts and plain words. For example, vandal marks vandalism, and BLP or living person marks violations of the policy on biographies of living people. Sock, LTA, block evasion or ban evasion mark sockpuppetry and long-term abuse, and edit war or content dispute marks edit warring. Contentious-topic enforcement matches terms such as arbitration, ARBECR, 30/500, CTOP, PIA and Arab-Israeli. When a reason matched several rules, a fixed order assigned one primary category, with contentious topics first, then sockpuppetry, living persons, edit warring, vandalism, disruptive editing, unsourced content, spam and copyright. Reading random samples of uncoded reasons led to two added rules (spam and copyright) before the final run. After coding, 1.7% of new article protections fell into the residual category. The rules are in the analysis code and can be audited.

3.3 The event-study sample

A stratified random sample of 450 new article protections was drawn from the 50,024 events: 200 semi-protections, 150 extended confirmed protections and 100 full protections. Stratification gives enough cases at the rarer levels; results are therefore reported by level, and pooled shares are weighted by the size of each level in the log where stated. Fifteen sampled articles had since been deleted, which left 435 events (198 semi, 147 extended confirmed and 90 full).

For each event the study downloaded three further records. First, daily page views by human users (agent=user, all access methods) from the Wikimedia Pageviews REST API for 30 days before to 30 days after the protection date. Second, every revision of the article and of its talk page from 14 days before to 14 days after the protection time, with user, anonymity flags and change tags. Third, the registration date, user groups and current edit count of each registered editor in those histories, from the users list of the Action API.

3.4 Attention measures

The baseline attention of a page is the median of its daily views from day −30 to day −8, where day 0 is the UTC date of protection. The pre-protection peak is the highest daily count from day −7 to day 0. A pageview spike is defined as a peak at least five times the baseline and at least 1,000 views. The analysis repeats the main shares with three-fold and ten-fold thresholds. Exposure is the share of views from day 0 to day 30 that fell on days when the protection was still in effect.

3.5 Editor classes and edit measures

Each revision was assigned to an editor class at the time of the edit. Unregistered edits are those flagged as anonymous (IP address) or, from 4 November 2025, made by a temporary account; the sample contains 121 temporary-account revisions, the first on 4 November 2025. Registered editors were classed by account age at the time of the edit: new accounts (under 4 days), young accounts (4 to 29 days) and established accounts (30 days or more). Accounts in the bot group or with names ending in bot were set apart. Because the software counts autoconfirmed age from the first edit, which cannot come before registration, every new account was certainly unable to edit a semi-protected page. Edit counts at the time of an edit are not available, so the study adds one conservative check for extended confirmed protection: an established account whose current edit count is below 500 cannot have held the extended confirmed right when it edited.

Revert status uses the mw-reverted change tag, which MediaWiki adds when a later edit restores an earlier version or undoes the edit. An edit without the tag was not reverted in that way, although it may have been changed by later ordinary edits. The share of unreverted edits is therefore an upper bound on edits that the community kept intact and a reasonable measure of edits it did not treat as damage.

3.6 Analysis

RQ1 uses counts, shares and duration quantiles from the full log. RQ2 uses spike shares by level and a logistic regression of indefinite protection on spike status, with level and the five most common reason categories as controls and heteroskedasticity-consistent (HC1) standard errors. RQ3 compares edits in the seven days before and after protection for the 309 sampled events that stayed in force for at least seven days. For each editor class, a Poisson model of event-period edit counts with event fixed effects estimates the ratio of post to pre edits (an incidence rate ratio, IRR), with standard errors clustered by event. Daily means with 95% confidence intervals show the timing for the 252 events in force for at least 14 days. The design is descriptive. Protection is applied because editing activity has changed, so the before-and-after differences cannot be read as the causal effect of protection alone. All data handling and statistics were done in R.

4. Findings

4.1 How protection is applied (RQ1)

Protection on English Wikipedia is mainly a tool for articles. Of protections and modifications that restricted editing, 72.8% concerned articles, 11.4% project pages, 5.9% templates and 4.8% user talk pages. Among new article protections, semi-protection dominated with 41,474 events (82.9%), followed by extended

confirmed protection with 7,661 (15.3%) and full protection with 889 (1.8%). Protection also repeats: 22.6% of the 36,549 protected articles were protected two or more times in the three years. Table 1 summarises levels, durations and reasons.

Table 1. New edit protections of articles on English Wikipedia by level, 2023 to 2025: volume, set duration and coded reason.

Measure	Semi	Extended confirmed	Full	All
New protections (n)	41,474	7,661	889	50,024
Share of all new article protections	82.9%	15.3%	1.8%	100%
Indefinite	6.9%	84.5%	19.1%	19.0%
Median set duration if temporary, days (IQR)	14 (7 to 91)	91 (28 to 365)	4 (2 to 7)	14 (7 to 92)
Reason: disruptive editing	29.2%	3.0%	2.6%	24.7%
Reason: vandalism	28.0%	1.2%	1.6%	23.5%
Reason: contentious topic or arbitration	4.6%	85.0%	1.2%	16.8%
Reason: sockpuppetry or long-term abuse	17.4%	6.4%	3.8%	15.5%
Reason: unsourced content	9.2%	0.4%	0.6%	7.7%
Reason: living persons (BLP)	8.4%	1.2%	1.5%	7.2%
Reason: edit warring or content dispute	1.0%	1.0%	70.8%	2.3%
Reason: other or unstated	2.2%	1.8%	18.0%	2.4%
Reason text cites a request at RfPP	34.6%	39.7%	17.5%	35.1%

The three levels are used in different ways. Semi-protection is the everyday tool for vandalism and disruption, which together account for 57.2% of semi-protections, and it is usually short. Its median set duration was 14 days, and only 6.9% of semi-protections were indefinite. Full protection is a short pause in disputes: 70.8% cited edit warring or a content dispute, and the median temporary duration was 4 days. Extended confirmed protection is different in kind. Of these protections, 85.0% cited contentious-topic or arbitration rules and 84.5% had no expiry.

Figure 1 shows the monthly volume. New article protections ranged from 1,128 to 2,722 a month, with a median of 1,306. Semi-protections held steady at 13,477 to 14,296 a year. Extended confirmed protections rose from 1,181 in 2023 to 1,502 in 2024 and 4,978 in 2025, with 1,420 in February 2025 alone. Much of that rise was bulk enforcement of rules for the Arab-Israeli conflict topic area. One identical reason text, which records enforcement of that topic rule at a single noticeboard request, appears on 888 extended confirmed protections. Two other standard arbitration-enforcement texts for the same area appear 663 and 446 times. Together these three texts cover 26.1% of all extended confirmed protections in the window.

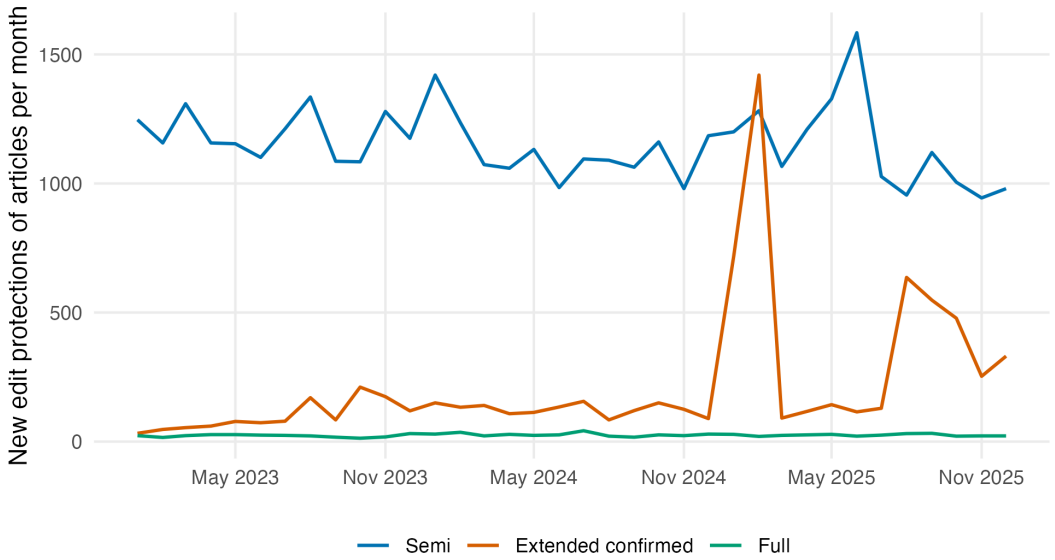

Figure 1. New edit protections of English Wikipedia articles per month by protection level, January 2023 to December 2025. Source: protection log, MediaWiki Action API.

Figure 2 shows set durations. Over a third of semi-protections (34.5%) were set for one to seven days and another 27.2% for eight days to a month. Full protections were even shorter, with 58.7% set for one to seven days. Extended confirmed protections stand apart, with almost every one either indefinite or set for months. Across all levels, indefinite protection tracks the reason more than the level: 83.8% of protections citing contentious topics were indefinite, against 4.4% of those citing vandalism and

5.2% of those citing living-person violations.

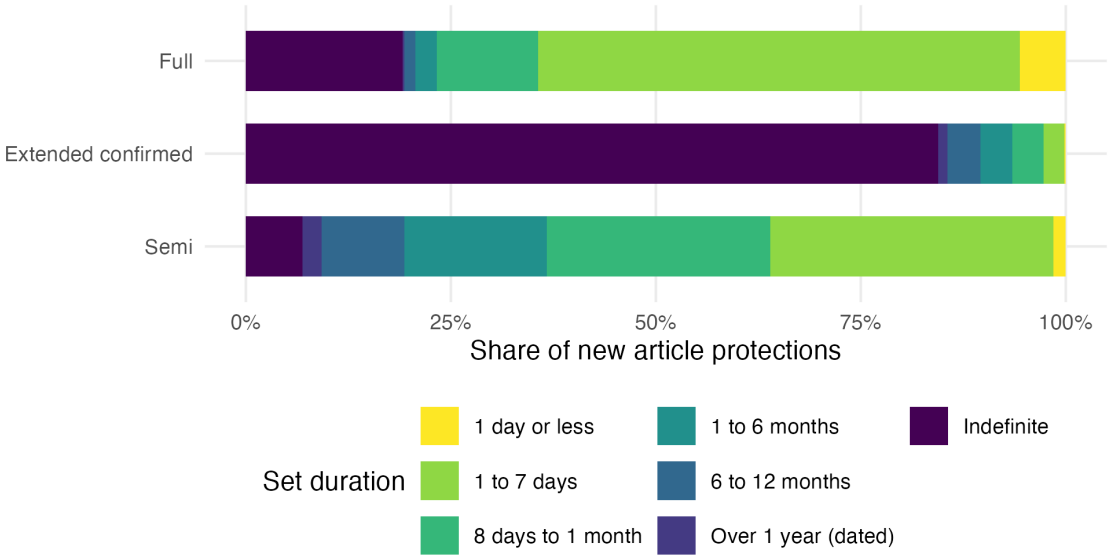

Figure 2. Set duration of new edit protections of English Wikipedia articles by protection level, 2023 to 2025 (n = 50,024). Durations are the time from protection to stated expiry. Source: protection log, MediaWiki Action API.

The answer to RQ1 is that English Wikipedia runs two protection regimes through one tool. The first is short semi-protection, usually requested on a noticeboard and justified by vandalism or disruption, and it lapses after days or weeks. The second is open-ended extended confirmed protection, justified by topic rules made by the Arbitration Committee, and it is applied to whole topic areas, often in batches. The second regime is the one that grew in 2025.

The log also shows how often protections are revised. Over the three years administrators recorded 9,859 modifications of existing protections and 3,939 removals, against 72,399 new protections in all namespaces. In the event sample, the next protection-log entry for the page came before the stated expiry for 4.0% of semi-protections, 14.3% of extended confirmed protections and 18.9% of full protections. Early change is thus most common for full protection, where a dispute may be settled on the talk page and the lock lifted, and for extended confirmed protection, where a page may be moved between levels. Most semi-protections simply run to their expiry. Because full protections are short, only 34.4% of sampled full protections stayed in force for a week, compared with 71.2% of semi-protections and 93.2% of extended confirmed protections. This matters for the analysis of editing in Section 4.3, which requires a full week of protection.

4.2 Protection and attention (RQ2)

Most protected articles were not experiencing a news-scale surge when they were locked. Using the main definition, 20.7% of sampled semi-protections, 11.6% of extended confirmed protections and 11.1% of full protections followed a pageview spike. Weighted by the size of each level in the log, about one in five new article protections (19.1%) followed a spike. The shares change little with the threshold: with a three-fold threshold they are 24.2%, 12.9% and 13.3%, and with a ten-fold threshold 17.2%, 10.2% and 10.0%. Pages also differed in baseline attention. The median semi-protected article had 245 daily views in its baseline weeks, compared with 134 for fully protected articles and 9 for extended confirmed articles, which reflects the bulk topic protections of little-read pages.

Figure 3 shows the typical path of attention. On the day of protection the median article had 1.53 times its baseline views at the semi level, 1.89 times at the extended confirmed level and 1.97 times at the full level. The upper quartile was far higher, at 6.4 to 8.6 times the baseline, and attention returned to the baseline within about a month. Protection thus tends to arrive at the top of an attention peak, not before it.

The study expected that news-driven protections might be left in place indefinitely once the surge passed. The data show the opposite. Among sampled protections that followed a spike, 22.1% were indefinite, compared with 36.0% of the rest. Within levels, 2.4% of spike-linked semi-protections were indefinite against 4.5% of other semi-protections, and the median effective length of spike-linked semi-protections was 7 days against 21 days. None of the 10 spike-linked full protections was indefinite. In the logistic regression with level and reason as controls, a spike was associated with lower odds of indefinite protection (odds ratio 0.26, 95% CI 0.07 to 0.90), an average difference of 6.7 percentage points. The interval is wide because only 68 sampled protections followed a spike.

Short protections do not mean short exposure. Across the sample, 55.9% of all views in the 31 days from protection onward fell on days when the protection was in effect. For the 68 spike-linked protections, 9.19 million of 15.99 million views in that period (57.5%) came while the page was locked. The answer to RQ2 is that about a fifth of protections respond to attention spikes, that these protections are kept short, and that they still cover most of the readers the spike brings.

4.3 Who stops editing (RQ3)

Table 2 compares editing in the seven days before and after protection for the 309 sampled protections that stayed in force for at least a week. The software does what it is designed to do. Unregistered editors made 1,688 article edits in the week before protection and 3 in the week after (IRR 0.002). New accounts made 298 edits before and 3 after (IRR 0.010). Young accounts, which may or may not have passed the semi-protection threshold, fell from 70 edits to 14 (IRR 0.20). Established accounts continued, with 2,462 edits before and 1,900 after (IRR 0.77, 95% CI 0.51 to 1.18).

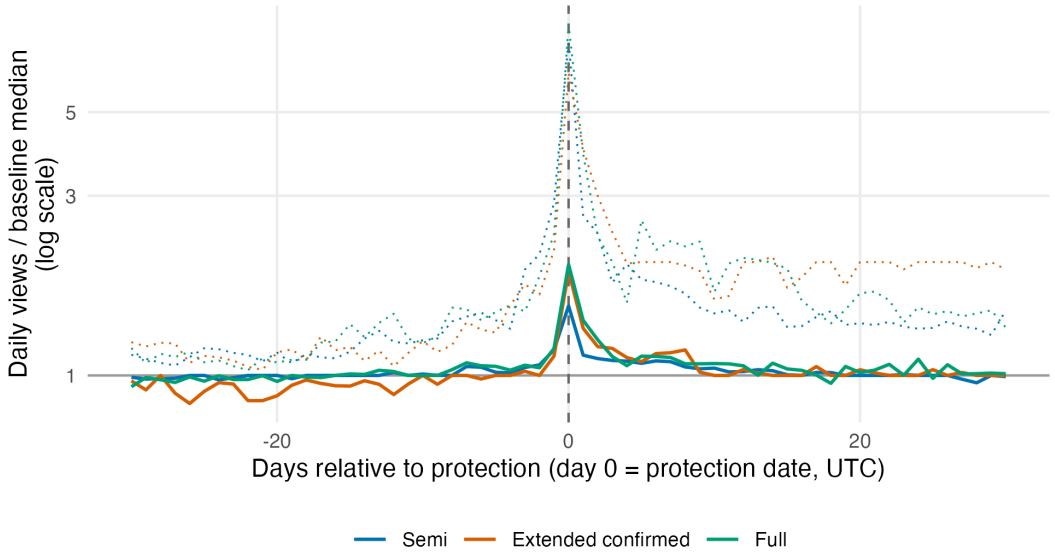

Figure 3. Daily user page views relative to each article's baseline (median of days −30 to −8), 30 days before to 30 days after protection, by level (n = 435 sampled protections). Solid lines show the median ratio and dotted lines the 75th percentile; the vertical line marks the protection date. Source: Wikimedia Pageviews REST API.

Table 2. Article and talk-page edits in the seven days before and after protection by editor class, sampled English Wikipedia protections in force for at least seven days (n = 309).

Page and editor class	Edits before	Edits after	IRR (95% CI)	Share of pre edits reverted
Article: unregistered	1,688	3	0.002 (0.000 to 0.013)	57.1%
Article: new account (under 4 days)	298	3	0.010 (0.002 to 0.045)	68.2%
Article: young account (4 to 29 days)	70	14	0.20 (0.07 to 0.58)	39.9%
Article: established account (30 days or more)	2,462	1,900	0.77 (0.51 to 1.18)	20.0%
Talk page: unregistered	52	85	1.63 (0.98 to 2.73)	
Talk page: new account	19	30	1.58 (0.40 to 6.29)	
Talk page: established account	395	766	1.94 (1.00 to 3.76)	

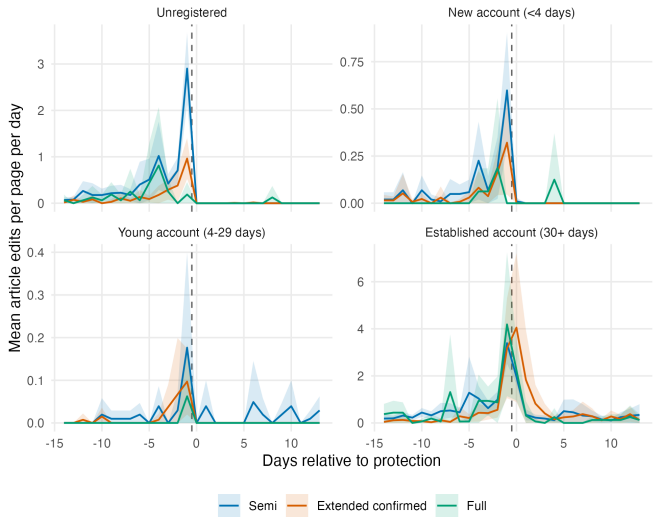

Figure 4 shows the timing. Edits by unregistered and new editors rise sharply in the days before protection, peak the day before and drop to zero at day 0. Established accounts also edit most around day 0, partly because their edits include the reverts and clean-up that come with the protection request.

The levels differ in who remains. On semi-protected articles, unregistered editors made 43.5% of edits in the week before protection and established accounts 46.5%; in the week after, established accounts made 89.2% and bots most of the rest. Edits by established accounts on these articles also fell (IRR 0.53, 95% CI 0.36 to 0.78), which suggests that part of their earlier activity was the work of reverting others. On fully protected articles, edits by established accounts fell to a fifth of the earlier level (IRR 0.20, 95% CI 0.12 to 0.34), which is consistent with a pause in an edit war. On extended confirmed articles, established accounts edited somewhat more after protection (IRR 1.53, 95%

CI 0.94 to 2.50). Extended confirmed protection also excludes a group that the account-age classes hide: older accounts with fewer than 500 edits. Before protection, 83 of 721 edits by established accounts on these articles (11.5%) came from accounts whose edit count is below 500 even today. After protection, such accounts made 2 of 1,106 edits (0.2%).

Some activity moves to the talk page, where protected articles direct readers to request changes. Unregistered talk-page edits rose from 52 to 85 (IRR 1.63, 95% CI 0.98 to 2.73), and on extended confirmed articles from 6 to 29 (IRR 4.83, 95% CI 1.94 to 12.04). These numbers are small. The 85 talk-page edits by unregistered editors after protection are about one-twentieth of the 1,688 article edits they made in the week before. The talk page receives a trickle of what the article had received.

The last column of Table 2 addresses what was stopped. In the week before protection, 57.1% of article edits by unregistered editors carried the reverted tag (95% CI 49.3% to 64.8%, clustered by event), as did 68.2% of edits by new accounts. These rates are high and explain why administrators protect. They also mean that 1,284 of 2,990 unregistered edits (42.9%) and 150 of 472 new-account edits (31.8%) were not reverted. A third of analysed protections (33.8%) were applied to articles that had received at least one unreverted unregistered edit in the preceding week. The answer to RQ3 is that protection removes unregistered and new editors almost entirely and only partly redirects them to talk pages, and that the stream it removes was mixed: mostly reverted, but with a large minority of edits that stayed.

5. Discussion

5.1 Two measures of the same access control

The findings describe one system that looks different depending on what is counted. Measured against the problem stated in its log entries, page protection works almost perfectly. The 1,688 unregistered article edits in the week before protection, more than half of which were reverted, become 3 in the week after. Administrators who protect a page see the disruption stop, and nothing in the revision history records what replaced it. This is the measure on which the system is evaluated, in the sense that a protection request succeeds when the reverting stops.

Measured by who is turned away, the same action looks different. Before the lock, 42.9% of unregistered edits were not reverted, and on a third of protected articles at least one such edit had been made in the previous week. After the lock, these contributors cannot edit the article, and only a small fraction appear on the talk page. Meanwhile 55.9% of the following month's views reach the article while it is locked. A reader who arrives after a news event and spots an error meets a closed Edit tab and an instruction to make a request on the talk page. This is the measure on which the system is experienced by the people it excludes, and Wikipedia does not record it. The gap matches the trade-off Hill and Shaw (2021) found for account requirements, where most deterred contributions were of higher quality, and the earlier evidence that one-time anonymous contributors can be reliable (Anthony et al., 2009).

The data do not show that protection is wrong in a given case. A reverted share near 60% is a heavy load for volunteers, and Geiger and Halfaker (2013) showed what happens when vandal-fighting capacity drops. The point is narrower: the record that Wikipedia keeps captures the cost protection removes and not the cost it creates. A protection decision weighs both, but only one is visible in the log, the revision history and the noticeboard archive.

5.2 Two regimes in one tool

The second finding concerns the shift toward extended confirmed protection. Semi-protection fits the picture of a temporary response to a surge: most are set for days or weeks, few are indefinite, and spike-linked semi-protections are even shorter. Extended confirmed protection fits a different picture. It is applied to whole topic areas under rules set by the Arbitration Committee, often in batches recorded with identical reason texts, and it rarely expires. Its median baseline of 9 daily views shows that it often covers pages that few people read, which suggests a preventive logic based on topic, not on observed disruption at the page.

This regime moves the threshold for participation from four days and ten edits to thirty days and five hundred edits for a growing set of pages. In this sample, 11.5% of pre-protection edits by established accounts on these pages came from accounts that still have fewer than 500 edits. The barrier therefore excludes registered users with modest editing histories as well as anonymous users. Shaw and Hill (2014) argued that peer production projects concentrate power in a small group over time. Extended confirmed protection writes a version of that concentration into software for the topics where disputes are most intense. Steinsson (2024) argued that such rule enforcement can make Wikipedia more reliable by removing editors who push fringe views. The present data cannot judge reliability, but they show the scale at which the rule now operates: 7,661 new protections in three years, 4,978 of them in 2025.

5.3 Breaking news

The results qualify a common worry about news events. The study found no sign that attention spikes lead to permanent locks: spike-linked protections were less often indefinite, with an odds ratio of 0.26. Administrators appear to treat news surges as temporary and set short expiries. Yet a short lock still covers most of the surge, because attention peaks around day 0 and returns to normal within weeks. Keegan et al. (2013) showed that breaking news articles are written by large, fast-forming groups of editors. Semi-protection during a surge narrows that group to accounts that already exist and have some history. Kiene et al. (2016) found that a community could absorb a surge of newcomers through norms and active moderation; Wikipedia's protection tool offers a quicker, cruder alternative that admits only people who are already members.

The design of the tool shapes this outcome. Protection acts on classes of accounts, not on edits or people, so it cannot distinguish a reader who fixes a date from a vandal who blanks the page. The thresholds that define those classes were set for other purposes: autoconfirmation also unlocks page moves and

file uploads, and extended confirmation was created as a tool for enforcing topic rules. When an administrator chooses a level, the choice is among these existing groups, and the software gives no option in between. This is a clear case of a technical system that allocates participation by turning social trust into simple counters. The counters are easy to apply at scale, which is why 50,024 article protections could be issued in three years by a small group of administrators. They are also blunt, which is why the edits they stop are a mix of damage and contribution.

5.4 Implications

Three practical points follow. First, protection requests and reviews could record what is lost as well as what is stopped, for example the share of unreverted edits by unregistered and new editors in the week before the request, which can be computed from the revision tags used here. Second, Wikipedia already offers pending changes, a separate setting that lets unregistered edits wait for review instead of blocking them; this study did not analyse its log, but its trade-offs deserve the same measurement. Third, research that counts edits or editors on Wikipedia should treat protection status as part of the data, as Hill and Shaw (2015) urged, because a sharp drop in unregistered editing after protection is a software effect, not a change in public interest.

6. Limitations

The study has clear limits. The design is observational: protection is applied because editing has changed, and the before-and-after comparisons describe what happened around protection, not what would have happened without it. The event sample is a stratified random sample of 450 protections, so estimates for spike-linked protections and for full protection rest on small numbers, as the wide intervals show. Fifteen sampled articles had been deleted and could not be analysed.

Several measures are proxies. Reason categories come from keyword rules on short free-text reasons and assign one primary category, so they may miss or misclassify some cases. Account age is measured from registration, while the software counts from the first edit, and edit counts at the time of an edit are not available, so the classes are conservative bounds for semi and extended confirmed access. The mw-reverted tag misses damage removed by ordinary edits, so the unreverted share overstates the share of edits that the community kept. Page views count all human readers, including logged-in editors, and page titles changed by later moves may split some view counts.

Most importantly, the logs cannot show the people protection turns away. An edit that was not made leaves no record, and neither does a reader who clicked Edit, saw a lock and left. The study estimates this cost from edits made just before the lock and readers who arrived during it; it does not observe it directly. The analysis also covers only English Wikipedia, whose size, administrator numbers and arbitration rules differ from those of other language editions, so the levels and durations reported here may not hold elsewhere. Interviews with readers and newcomers, or server-side counts of attempts to edit protected pages, would be needed to measure it. The introduction of temporary accounts in November 2025 affects only the last weeks of the window and was handled by treating temporary accounts as unregistered.

7. Conclusion

Page protection turns a judgement about trust into a software rule based on account age and edit count. On English Wikipedia from 2023 to 2025 it was applied to 50,024 articles, mostly as short semi-protection against vandalism and disruption, and increasingly as open-ended extended confirmed protection for contentious topics. About one in five protections followed a spike in page views, and these were kept short. In the week after protection, unregistered and new editors almost disappeared from the protected articles, and only a small fraction of their activity reappeared on talk pages. By the measure Wikipedia records, the stopping of disruption, protection is close to fully effective. By the measure it does not record, it turns away a stream of contributions that was more than 40% unreverted, during the weeks when most readers arrive. Both measures belong in any account of how this access control governs the encyclopedia that anyone can edit.

Future work can extend the second measure. Counts of attempts to open the editor on protected pages, if Wikimedia were to release them in aggregate, would show how many readers meet the lock. A study of edit requests on talk pages could show how many of the excluded contributions are later made by others and how long that takes. Comparing English Wikipedia with language editions that protect pages less often would show whether the same surges can be handled with more open doors.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

This research received no external funding.

Ethics statement

This study is a secondary analysis of publicly available data from the English Wikipedia protection log, revision histories, user records and page-view statistics. No human participants were contacted. The analysis reports aggregate results only and does not name any editor or administrator. No ethics approval was required.

Data availability statement

All data are public. The protection log, revision histories and user records come from the MediaWiki Action API for English Wikipedia (https://en.wikipedia.org/w/api.php), accessed October 2026. Daily page views come from the Wikimedia Pageviews REST API (https://wikimedia.org/api/rest_v1/metrics/pageviews/per-article/), accessed October 2026. The protection thresholds are described in the English Wikipedia protection policy (https://en.wikipedia.org/wiki/Wikipedia:Protection_policy), accessed October 2026. The analysis code will be deposited in a public repository on acceptance; it is withheld during review only to keep review anonymous.

References

Anthony, D., Smith, S. W., & Williamson, T. (2009). Reputation and reliability in collective goods: The case of the online encyclopedia

Wikipedia. Rationality and Society, 21(3), 283–306. https://doi.org/10.1177/1043463109336804

Avieson, B. (2019). Breaking news on Wikipedia: Collaborating, collating and competing. First Monday, 24(5). https://doi.org/10.5210/fm.v24i5.9530

Butler, B., Joyce, E., & Pike, J. (2008). Don't look now, but we've created a bureaucracy: The nature and roles of policies and rules in Wikipedia. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. https://doi.org/10.1145/1357054.1357227

Ciampaglia, G. L., Flammini, A., & Menczer, F. (2015). The production of information in the attention economy. Scientific Reports, 5, Article 9452. https://doi.org/10.1038/srep09452

Ford, H., & Wajcman, J. (2017). 'Anyone can edit', not everyone does: Wikipedia's infrastructure and the gender gap. Social Studies of Science, 47(4), 511–527. https://doi.org/10.1177/0306312717692172

Forte, A., & Bruckman, A. (2008). Scaling consensus: Increasing decentralization in Wikipedia governance. In Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008). IEEE. https://doi.org/10.1109/HICSS.2008.383

Forte, A., Larco, V., & Bruckman, A. (2009). Decentralization in Wikipedia governance. Journal of Management Information Systems, 26(1), 49–72. https://doi.org/10.2753/MIS0742-1222260103

García-Gavilanes, R., Mollgaard, A., Tsvetkova, M., & Yasseri, T. (2017). The memory remains: Understanding collective memory in the digital age. Science Advances, 3(4), Article e1602368. https://doi.org/10.1126/sciadv.1602368

Geiger, R. S. (2017). Beyond opening up the black box: Investigating the role of algorithmic systems in Wikipedian organizational culture. Big Data & Society, 4(2). https://doi.org/10.1177/2053951717730735

Geiger, R. S., & Halfaker, A. (2013). When the levee breaks: Without bots, what happens to Wikipedia's quality control processes? In Proceedings of the 9th International Symposium on Open Collaboration. ACM. https://doi.org/10.1145/2491055.2491061

Geiger, R. S., & Ribes, D. (2010). The work of sustaining order in Wikipedia: The banning of a vandal. In Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work. ACM. https://doi.org/10.1145/1718918.1718941

Gillespie, T. (2019). Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media. Yale University Press. https://doi.org/10.12987/9780300235029

Gorwa, R., Binns, R., & Katzenbach, C. (2020). Algorithmic content moderation: Technical and political challenges in the automation of platform governance. Big Data & Society, 7(1). https://doi.org/10.1177/2053951719897945

Halfaker, A., Geiger, R. S., Morgan, J. T., & Riedl, J. (2013). 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. https://doi.org/10.1177/0002764212469365

Halfaker, A., Kittur, A., & Riedl, J. (2011). Don't bite the newbies: How reverts affect the quantity and quality of Wikipedia work. In Proceedings of the 7th International Symposium on Wikis and Open Collaboration. ACM. https://doi.org/10.1145/2038558.2038585

Hill, B. M., & Shaw, A. (2015). Page protection: Another missing dimension of Wikipedia research. In Proceedings of the 11th International Symposium on Open Collaboration. ACM. https://doi.org/10.1145/2788993.2789846

Hill, B. M., & Shaw, A. (2021). The hidden costs of requiring accounts: Quasi-experimental evidence from peer production. Communication Research, 48(6), 771–795. https://doi.org/10.1177/0093650220910345

Keegan, B., Gergle, D., & Contractor, N. (2013). Hot off the wiki: Structures and dynamics of Wikipedia's coverage of breaking news events. American Behavioral Scientist, 57(5), 595–622. https://doi.org/10.1177/0002764212469367

Kiene, C., Monroy-Hernández, A., & Hill, B. M. (2016). Surviving an "Eternal September": How an online community managed a surge of newcomers. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. ACM. https://doi.org/10.1145/2858036.2858356

Kittur, A., Suh, B., Pendleton, B. A., & Chi, E. H. (2007). He says, she says: Conflict and coordination in Wikipedia. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. https://doi.org/10.1145/1240624.1240698

Priedhorsky, R., Chen, J., Lam, S. K., Panciera, K., Terveen, L., & Riedl, J. (2007). Creating, destroying, and restoring value in Wikipedia. In Proceedings of the 2007 International ACM Conference on Supporting Group Work. ACM. https://doi.org/10.1145/1316624.1316663

Shachaf, P., & Hara, N. (2010). Beyond vandalism: Wikipedia trolls. Journal of Information Science, 36(3), 357–370. https://doi.org/10.1177/0165551510365390

Shaw, A., & Hill, B. M. (2014). Laboratories of oligarchy? How the iron law extends to peer production. Journal of Communication, 64(2), 215–238. https://doi.org/10.1111/jcom.12082

Star, S. L., & Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. Information Systems Research, 7(1), 111–134. https://doi.org/10.1287/isre.7.1.111

Steinsson, S. (2024). Rule ambiguity, institutional clashes, and population loss: How Wikipedia became the last good place on the internet. American Political Science Review, 118(1), 235–251. https://doi.org/10.1017/S0003055423000138

Suh, B., Convertino, G., Chi, E. H., & Pirolli, P. (2009). The singularity is not near: Slowing growth of Wikipedia. In Proceedings of the 5th International Symposium on Wikis and Open Collaboration. ACM. https://doi.org/10.1145/1641309.1641322

TeBlunthuis, N., Shaw, A., & Hill, B. M. (2018). 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. ACM. https://doi.org/10.1145/3173574.3173929

Tran, C., Champion, K., Forte, A., Hill, B. M., & Greenstadt, R. (2020). Are anonymity-seekers just like everybody else? An analysis of contributions to Wikipedia from Tor. In 2020 IEEE Symposium on Security and Privacy (SP). IEEE. https://doi.org/10.1109/SP40000.2020.00053

Twyman, M., Keegan, B. C., & Shaw, A. (2017). Black Lives Matter in Wikipedia: Collective memory and collaboration around online social movements. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing. ACM. https://doi.org/10.1145/2998181.2998232

Wikipedia contributors. (2026). Wikipedia:Protection policy. Wikimedia Foundation. Retrieved October 2026, from https://en.wikipedia.org/wiki/Wikipedia:Protection_policy

Yasseri, T., Sumi, R., Rung, A., Kornai, A., & Kertész, J. (2012). Dynamics of conflicts in Wikipedia. PLoS ONE, 7(6), Article e38869. https://doi.org/10.1371/journal.pone.0038869

Corresponding author: Altynai Asanova, Osh State University. Email: altynai.asanova2026@gmail.com

Received: 12 June 2026. Accepted: 15 July 2026. Published: 22 July 2026.

© 2026 The Author(s). This is an open access article distributed under the terms of the CC BY 4.0 licence (https://creativecommons.org/licenses/by/4.0/), which

permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.