

Accept, decline, postpone: How newcomer productivity is reduced in English Wikipedia by pre-publication review

Jodi Schneider*
INRIA Sophia Antipolis France
jodi.schneider@inria.fr

Bluma S. Gelley
NYU Polytechnic School of
Engineering
Brooklyn, NY
bgelley@nyu.edu

Aaron Halfaker
Research & Data
Wikimedia Foundation
aaron.halfaker@gmail.com

ABSTRACT

Wikipedia needs to attract and retain newcomers while also increasing the quality of its content. Yet new Wikipedia users are disproportionately affected by the quality assurance mechanisms designed to thwart spammers and promoters. English Wikipedia's Articles for Creation provides a protected space for drafting new articles, which are reviewed against minimum quality guidelines before they are published. In this study we explore how this drafting process has affected the productivity of newcomers in Wikipedia. Using a mixed qualitative and quantitative approach, we show how the process's pre-publication review, which is intended to improve the success of newcomers, in fact decreases newcomer productivity in English Wikipedia and offer recommendations for system designers.

Categories and Subject Descriptors

H.5.3 [Information Systems]: Group and Organization Interfaces—*computer-supported collaborative work*

Keywords

Wikipedia; newcomers; Articles for Creation; quality control; pre-publication review

1. INTRODUCTION

Starting in 2009, headlines decried the decline of Wikipedia. Volunteers were leaving the project in alarming numbers [24]. With articles published on the most common topics, and an increasing need to combat spam, the “easy work” of creating as many articles as possible gave way to an increasing focus on quality control. Wikipedia seemed to be becoming a more hostile environment, especially for newcomers.

*This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 246016.

OpenSym '14, August 27 - 29 2014, Berlin, Germany
Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3016-9/14/08...\$15.00
<http://dx.doi.org/10.1145/2641580.2641614>

From the point of view of a reader, Wikipedia has continued to improve: it has grown larger, more complete, and seemingly of better quality. Popularity has, however, challenged the work of Wikipedia's maintainers, as spammers, promoters, and a seemingly endless stream of “clueless” new users seek to add content that does not seem encyclopedic. Creating better encyclopedia articles seems to conflict with the goal of creating a larger encyclopedia. Protecting against an oversupply of unwanted content, however, threatens to constrict the supply of another, potentially valuable resource: the new volunteers.

New Wikipedia users are disproportionately affected by the quality assurance mechanisms designed to thwart spammers and promoters [9, 13, 15]. This is a potentially hazardous situation for sustaining an all-volunteer collective: the future capacity for production depends on integrating new volunteers. Retaining new volunteers may depend on where they make their first edit, and on whether that edit is accepted.¹

Creating new articles is particularly challenging for new editors, and this is a sizable problem since one editor in three contributes a new article immediately after registering an account. On new articles, quality assurance is particularly speedy and harsh. Within minutes, new articles in Wikipedia's Main namespace may be ‘patrolled’ by an experienced user. Offending articles are tagged for speedy deletion and typically deleted within 30 minutes even when they are actively under development [6]. As of 2011, 22% of all deletions (and 37% of all speedy deletions) were due to A7: No indication of importance [8]. New article creation is a battlezone where socializing newcomers seems to take a back seat to ensuring quality control. Until now, however, a vital component in newcomer socialization, the Articles for Creation (AfC, see Section 4.1 below) process, has been overlooked. In this paper, we explore the impact of the AfC process on newcomer retention.

2. MOTIVATION

2.1 Newcomers are needed

Wikipedia's ideal is that “anyone can edit”, and ongoing editing is needed for a well-maintained, growing encyclopedia. Because volunteers can leave the project at will, as in other open collaboration systems [7], the sustainability of the project has thus far relied on replacing exiting volunteers. Mathematical models of WikiProjects suggest that

¹http://en.wikipedia.org/Wikipedia:Wikipedia_Signpost/2011-04-04/Editor_retention

new and infrequent editors are particularly important for building critical mass and ensuring the continued existence of a content-creation community [22]. Further, more diverse content is desirable, and new editors might also help bridge the gender gap [4] and address cultural bias [3]. Yet in recent years, retaining newcomers has become more difficult. A large body of work (e.g. [9, 13, 14, 15, 17]) has implicated negative feedback and deleted contributions in the decline in newcomer retention.

2.2 Sudden decline in newcomer productivity

While there has been no substantial change in the proportion of new editors who create *some page*, we observed a substantial dip in the rate of *successful* article creation by newcomers in early 2011. Figure 1 shows the proportion of articles of each new page creator that survive (i.e. remain published in the Main article space) for at least 30 days. There is a clear drop in the survival rate in February 2011.




Figure 1: Surviving articles per new page creator. The number of surviving articles per newcomer page creator is plotted with loess fits before and after AfC’s era began in February 2011. CC BY-SA 3.0 Aaron Halfaker




Figure 2: The rising use of AfC. The proportion of newcomer (≤ 1 month tenure) articles created as AfC Drafts is plotted over time with a vertical line denoting pre-AfC and post-AfC epochs. CC BY-SA 3.0 Aaron Halfaker

Newcomers increasingly use a process called Articles for Creation to create draft articles. Figure 2 shows the rate at which drafts were created in AfC by newcomers (≤ 1 month since registration). A sharp increase in the rate is visible beginning in mid-2011. This increase is so marked that we refer to the epochs before and after it as “pre-AfC” and “post-AfC”.

The rising use of AfC correlates with the decrease in newcomer article survival rate at around that time. Is this by chance, or is something important happening here?

3. HYPOTHESES

This leads us to consider the AfC process, and to hypothesize that the non-wiki, “review first” process is a problem.

One potential problem is the delay that a pre-publication review enforces on draft creators. A long wait period between when an editor submits a draft for review and the completion of the review process may be discouraging. Such a long wait period would suggest that many good draft articles languish in a backlog which would explain the decreased level of productivity.

H1: The review process is too slow.

Another potential problem is that, the lack of visibility of drafts may reduce the capacity for efficient, stigmergic collaboration to take place. The process of loose coordination around articles in Wikipedia bears a striking similarity to open source software projects, which seem to operate under Linus’ Law—“given enough eyeballs, all bugs are shallow” [19]—since attracting ‘eyeballs’ seems to correlate with success [21]. If this phenomenon were also present in Wikipedia, then hiding drafts from the view of potential collaborators would result in missed opportunities for collaboration – therefore reducing productivity.

H2: The review process hides drafts from potential collaborators.

4. BACKGROUND & METHODS

4.1 Articles for Creation (AfC)




Figure 3: The process of drafting, reviewing, and publishing articles under the Articles for Creation (AfC) process. CC BY-SA 3.0 Aaron Halfaker

Articles can enter the AfC process in several ways, as shown in Figure 3. Editors create a draft, often through

the Article Wizard,² which provides step-by-step guidance in drafting an article, and deposits the article in a designated namespace. Drafts created in the User namespace, where each logged-in editor has their own space to create personal pages and drafts, may also be reviewed in AfC.

Drafts are shielded from the harsh page patrol environment. Instead, volunteer reviewers (experienced Wikipedians who have more than 500 accepted edits and more than 90 days tenure) provide feedback on essential quality measures such as reliable sources, as well as norms such as formatting. The reviewing process³ has three main steps, checking first that there are no major policy violations, second that the topic is suitable for an article (e.g. notability and verifiability), and finally that the article as written is sufficiently neutral and encyclopedic. Drafts passing all of these checks can be moved into Wikipedia article space. The bulk of articles are declined, with a rationale (such as, reads like an advertisement, does not meet a given notability criterion; or appears to be a test edit); though in severe cases (e.g. violating copyright and other major policies), drafts may be blanked and swiftly deleted. Article creators are generally notified about the outcome of the review, and declined articles can be edited and resubmitted for review. Articles which are not improved for 6 months are flagged for deletion under the “G13” speedy deletion category. Editors may search G13-eligible drafts, postponing deletion of content that has potential to be improved into an encyclopedic article.⁴

4.2 General terminology

On Wikipedia, there is a great deal of complexity and nuance to the jargon for various kinds of wiki pages and page creation processes. The following definitions are terms we will use in our quantitative analyses:

Page any wiki page in all namespaces. If we are referring to pages within a particular namespace only, we will specify the namespace, or (for the Main namespace) use *article*.

Article a page that, at some point, is *published* in the encyclopedia. Articles in the encyclopedia appear in the Main namespace (Namespace 0).

Published we refer to a page as published when it becomes visible in the encyclopedia. This is operationalized as the time of creation for *Direct to Main* articles or the time at which drafts are moved to the Main namespace.

Removed we refer to an article as removed when it is removed from the encyclopedia (i.e. the Main namespace). This is operationalized as either a page deletion or a move to a non-Main namespace.

AfC Draft an article created as a draft within AfC’s process. This includes all pages created in the “Wikipedia_talk” namespace (5) whose title is prefixed by “Articles_for_Creation/”. Drafts are *published* when they are moved to the Main namespace

²http://en.wikipedia.org/wiki/Wikipedia:Article_wizard

³http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation/Reviewing_instructions#Reviewing_workflow

⁴http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation/G13_rescue

Direct to Main a page created in the Main namespace (0). As of the first revision, *Direct to Main* articles are *published* and visible within the encyclopedia.

4.3 Approximations and assumptions

4.3.1 Publication, removal, original namespace

Due to a bug in the way that page moves are recorded⁵, we were unable to consistently associate move log items with an actual page move, so we identified page moves by looking for structured comments left in the revision history by MediaWiki’s page move functionality using a regular expression:

```
. *moved .*\\[[\\[[\\[\\]]+\\]]\\] to \\[[\\[[\\[\\]]+\\]]\\].*.*
```

With this strategy, we are able to identify *published* and *removed* events as well as identify the namespace of origin for all pages.

4.3.2 Time of deletion

Due to a bug in MediaWiki⁶, it’s not possible to consistently extract the time of deletion for pages. Given that the process for deleting pages on Wikipedia involves the addition of structured templates noting the pending deletion⁷, which MediaWiki registers as an edit, we expect that the last revision to a page will correspond closely with its actual deletion time. So we approximate the *time of deletion* using the last edit made to a page before it is deleted.

4.3.3 Success as article survival




Figure 4: Article lifetime of removed articles. The density of time between first and last edits is plotted for removed articles created between 2008 and 2013 in the English Wikipedia. CC BY-SA 3.0 Aaron Halfaker

Since Wikipedians work to ensure that problematic articles do not remain in the encyclopedia [8, 20, 28], we hypothesize that articles that reach a minimum level of quality and

⁵https://bugzilla.wikimedia.org/show_bug.cgi?id=57084

⁶https://bugzilla.wikimedia.org/show_bug.cgi?id=26122

⁷http://meta.wikimedia.org/wiki/Research:The_Speed_of_Speedy_Deletions

relevance will remain in the encyclopedia once published. If an effective review process is in place, new articles below such a quality threshold are removed in a *reasonable amount of time* and articles above the threshold should remain. In such a system, any article that survives a certain amount of time can be assumed to be a *successful* article creation. In order to identify how long a *reasonable amount of time* might be, we performed an analysis of the time between creation and removal. Figure 4 shows a strong cluster between one minute and one hour. While some removal events take more than one year, 87.3% of removal events occurred within 30 days of creation. Based on these observations, we operationalize a reasonable amount of time as 30 days, which leads to the following definition:

Surviving article an article that, once *published*, remains in the encyclopedia (not *removed*) for at least 30 days.

4.4 Tracking the AfC process

In order to track AfC’s process (see Figure 3) we looked for digital traces [10] in the form of templates used by AfC reviewers to track the status of *AfC drafts* and published articles that originate in AfC. To identify instances of status templates, we scanned the revision content of all current AfC drafts (as of Nov 13th, 2013) and all articles (including deleted articles) whose talk pages⁸ appear in any of the relevant categories (“Draft_AfC_submissions”, “Pending_AfC_submissions”, “Accepted_AfC_submissions” & “Declined_AfC_submissions”) using a combination of database queries, publicly available dumps and Wikipedia’s API. This status change dataset allows us to track the timing of submission, review and publication of AfC drafts using artifacts of AfC’s process. For more details on our extraction strategy, see our open repository⁹. When measuring collaboration activities on drafts, we filter edits that change the status of the draft since those edits tend not to involve substantial contributions to the draft’s content.

4.5 Qualitative Methods

We supplemented our quantitative investigation with qualitative observation. Using netnography [16], we observed and participated in AfC. We observed discussion spaces for reviewers and newcomers, AfC reviews, and newly published Wikipedia articles emerging from the AfC process. We also conducted AfC reviews and participated in ongoing discussions. We report on qualitative observations in Section 6 after first discussing quantitative results.

5. RESULTS AND DISCUSSION

5.1 H1: The AfC review process is too slow.

To look for evidence of substantial delays in AfC’s review process, we observed the time differences between when drafts were first submitted for review and when the first review action took place. Figure 5 plots the density of time between submission and review. The distribution seems to be bimodal with modes 8 hours and 1 week. While a substantial proportion of reviews take longer than 1 week (11%), the majority of reviews will take place within less than 24 hours of submission (65.5%).

⁸So as not to clutter an article, templates of accepted AfC drafts appear on the article’s associated Talk page

⁹<https://github.com/halfak/Articles-for-Creation>

There’s another part of AfC’s process that could be responsible for the delay. It could be that newcomer draft creators aren’t submitting their drafts for review in a timely manner. Figure 6 plots the density of time between draft creation and the first submission for review. Here we see a much more complex distribution. Most AfC drafts are submitted for review within minutes of their creation, but again, a substantial portion of AfC drafts are submitted for review more than a week after creation (7%).

Together, these results suggest that AfC’s review process is generally healthy – that most drafts will be submitted for review quickly and that reviews will happen in a timely manner.

However, there’s an additional concern here. Do all drafts get submitted for review? When we looked at drafts created more than a year before the end of our time horizon (November 2013), 29.4% of AfC drafts created by newcomers were never submitted for review or reviewed. Given this result, it seems more likely that any breakdown in AfC’s process is due to newcomers not understanding the draft submission process rather than because the review process is not fast enough once initiated.




Figure 5: Time to review. The density of time between when AfC submissions are first marked as “pending review” and when a review is completed is plotted. CC BY-SA 3.0 Aaron Halfaker




Figure 6: Time to submission. The density of time between when AfC submissions are created and when they are first marked as “pending review” is plotted. CC BY-SA 3.0 Aaron Halfaker

5.2 H2: AfC hides drafts from potential collaborators.

5.2.1 Quality through collaboration

The quality model for Wikipedia is collaboration. Once an article is created, each successive edit modifies the article directly. Sequential handoffs (wiki-like work patterns) are highly efficient and effective for producing high quality artifacts [1]. They allow large numbers of editors to collaborate [25], and in the best cases, good articles emerge. Extra effort and technical expertise [23] yield well-referenced articles suitable for mainstream consumption. When Wikipedia was young and less popular, readers and editors were a substantially overlapping population. So there was no perceived need to hide procedures or ‘rough draft’ articles: everything was in process. The minimum content standards have also changed over time; for instance, currently, articles must have inline references in order to be published through the AfC process.

But distinguishing mainstream consumers from ‘backstage’ (cf. [12]) editors, and enforcing high quality standards on articles for readers has a downside, because quality filtering and collaborative improvement conflict. Hiding information does not allow it to be iteratively improved, and instead the same or similar content may be created multiple times (see Section 6.4).

The collaborative model of wiki editing is that we will all do it together, and there is no time limit.¹⁰ By contrast, the pattern of review at AfC means waiting for an authorized person to put their “stamp of approval” on the article. While anyone can edit the article during that time, in practice, as shown in Figure 7, few people do. Rather, in AfC, there are two clearly separated roles: the reviewer, who makes a binary accept/decline decision for each draft, and the author, who must wait for the reviewer’s decision, then revise if necessary to reach a binary quality threshold.¹¹

In effect, the reviewing system places most of the responsibility for improving the article on the creator. For inexperienced creators, this is problematic. They have different perspectives from established volunteers [2] and less knowledge of policy.

While stubs that are created directly as articles in the Main namespace are easily findable via Google search and links from other articles in the encyclopedia, AfC drafts are relatively hidden from view. The “Wikipedia_talk” namespace that is used for AfC drafts is not indexed by Google¹² and is not linked to from encyclopedia articles as a matter of policy.¹³

If AfC drafts aren’t visible to most potential contributors, they won’t benefit from the many-eyes principle and have their quality improved by many collaborators. Hypothesis 2 examines whether AfC’s lack of visibility to potential collaborators is one of the reasons newcomers struggle to successfully create articles via AfC.

¹⁰<https://en.wikipedia.org/wiki/Wikipedia:DEADLINE>

¹¹Inexperienced authors who attempt to circumvent the review process by moving AfC drafts to published Main space are found out and may be chastised.

¹²However, copies of AfC drafts (from mirrors outside wikipedia.org) can still appear in Google search results.

¹³Within Wikipedia, AfC pages can be found by Wikipedia’s “Everything” search .

5.2.2 Measuring collaboration

We use contributions by non-creating editors as a measure of collaboration. To look for evidence of decreased collaboration on AfC drafts, we measured contributions by non-creating editors for the first 4 weeks of Direct to Main articles and AfC drafts. Figures 7 to 10 plot the geometric mean¹⁴ revisions and bytes changed by non-creating users as well the number of unique non-creating users (registered and anonymous) involved in editing the article.

Overall, Direct to Main articles see more collaborators (both registered and anonymous) and more revisions than AfC drafts for all four weeks post-article creation. This suggests that AfC drafts see a decreased level of collaboration than Direct to Main articles.

The relationship between how an article is created and the amount of actual content changed by other editors is less clear. Figure 7 shows that, regardless of AfC epoch, articles that were started in AfC saw more bytes changed by non-creating users in their first week. However, from weeks 2 through 4, this relationship reverses: Direct to Main articles get more attention and the difference becomes more substantial. From these results, it seems clear that Direct to Main articles see more attention from other editors.

We began with two hypotheses about the Articles for Creation process, and found:

H1: *The AfC review process is too slow.* **FALSE**

H2: *AfC hides drafts from potential collaborators.* **TRUE**

Next we describe our qualitative analysis.

6. QUALITATIVE OBSERVATIONS

Our qualitative analysis uncovered four additional problems. From a reviewer’s perspective, rejecting articles is safest; and subject-specific expertise is needed. Further, the reviewing process confuses newcomers; drafts may be long-lasting and duplicated, reducing productivity.

6.1 Rejecting articles is safest

Reviewers are under pressure to decline articles, which sends them back to the original author with indications of the current flaws and suggestions for improvement. Evidence of such pressure comes from Wikipedians outside the reviewing population. Mistakes can be highly visible, and result in flak for the editors who approved an article. *The problem is, we have some very real issues with people accepting articles at AfC that do not meet notability standards ... we have a big, big problem with people accepting articles that have no place being accepted in the first place.*¹⁵

Direct communication to editors about articles they accepted may be harsh: *I’m also surprised that you rated this B-class and let it get past AFC with copyvio photos in it... (not to mention promotional language like “on (sic) of the three top journals”, for a journal that was established last year, really?) I realize that people at AfC are overburdened*

¹⁴We use the geometric mean due to control for the long-tailed distribution of revisions saved and unique editors per article. Compared to an arithmetic mean, geometric means are less susceptible to extremely large values.

¹⁵http://en.wikipedia.org/w/index.php?title=User_talk%3AJasonwilczak&diff=605102917&oldid=605092777




Figure 7: Revisions by others. The geometric mean revisions per week (excluding creating user and AfC status change edits) is plotted for AfC drafts and Direct to Main articles. CC BY-SA 3.0 Aaron Halfaker

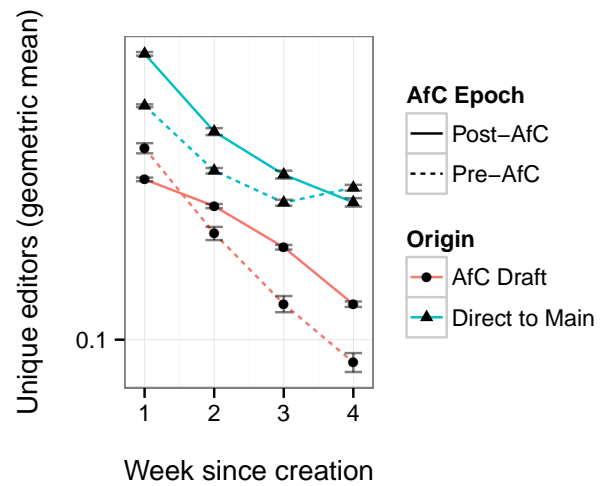


Figure 9: Unique non-creating users. The geometric mean number of registered editors per week (excluding creating user and AfC status change edits) is plotted for AfC drafts and Direct to Main articles. CC BY-SA 3.0 Aaron Halfaker




Figure 8: Bytes changed by others. The geometric mean bytes changed per week (excluding creating user and AfC status change edits) is plotted for AfC drafts and Direct to Main articles. CC BY-SA 3.0 Aaron Halfaker

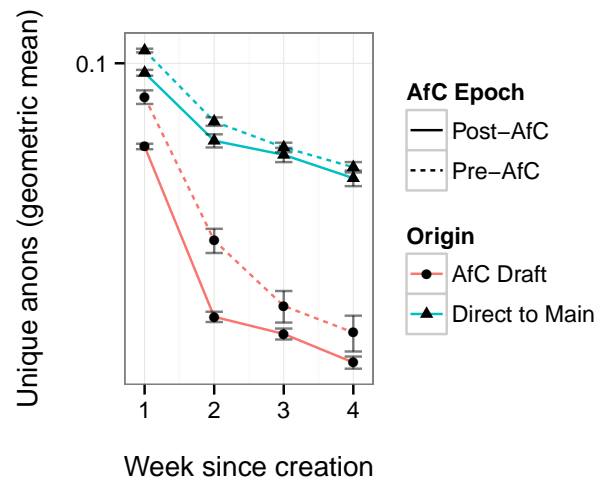


Figure 10: Unique non-creating anons. The geometric mean number of anonymous editors per week (excluding creating user and AfC status change edits) is plotted for AfC drafts and Direct to Main articles. CC BY-SA 3.0 Aaron Halfaker

and appreciate their efforts, but, really. . .”¹⁶

Reviewing is a hard job, made harder by the critical eye of other editors. No wonder attracting and retaining reviewers is difficult!¹⁷ Reviewing is high stakes compared to the iterative, wiki edit process in which each edit builds upon the next.

6.2 Subject-specific expertise is needed

While reviewers get flak from other editors for being too permissive in the drafts they publish, they also get complaints because they fail to publish other, encyclopedic articles. These false negatives may be due to misinterpreting the notability guidelines, which are subject specific and require significant judgement:

*My article was declined because of inadequate references, while all references were (correctly) cited medical journals. I had to go to the Wiki Medicine talk page to request a new reviewer (who accepted my article within an hour, without any changes to the references).*¹⁸

AfC reviewers are not experts on the subject matter of the drafts they review; rather, they are experts on Wikipedia policy. On the same day, a reviewer might consider articles on a World War I soldier, a children’s TV special from the 1980’s, a South African band, and an Indian village. Referencing and notability guidelines are specialized and appropriate sourcing can depend on the domain, so this can cause challenges.

Special rules in particular domains make things even more complicated. For instance, two albums by a mainstream recording company ensure that an artist is worth writing about, while having an impact factor assigned by the Institute for Scientific Information’s Journal Citation Reports emphatically qualifies an academic journal for a Wikipedia article. Reviewers unfamiliar with the subject, however, may be unaware of these rules and see only that the article fails Wikipedia’s General Notability Guideline¹⁹.

In addition, AfC submissions are often evaluated solely on what the draft contains, rather than on the notability of the draft’s topic itself. Reviewers’ work is to sort articles, not to write them. For instance, an article about an Emmy-award winning TV show from the 1980’s was twice declined at AfC, before finally being published 15 months after the draft was started. But it can be time-consuming to source this kind of information, especially when notability is not asserted, and for every notable article with notability not sufficiently asserted, there are several more that are hopelessly non-notable. The problem increases for material that would demand offline or pay-per-view sources, such as ‘historic’, pre-1990 topics, and for material with mostly non-English-language sources.

6.3 Reviewing procedures confuse newcomers

Reviewing opens up new things for newcomers to be confused about. The bulk of queries incoming to the AfC Help desk²⁰ revolve around a few issues: *is my article really submitted for review?* (which may be either a technical problem with the submission templates or a desire to understand how the queuing system works); *why did you reject my article?*; *how can I improve my article?*; *when can I expect a review?*; *where did my draft go?* This shows newcomers’ confusion with the process.

These indicate some further problems newcomers have in interacting with the review process:

- Newcomers do not understand the review process, including how to submit articles for review and the expected timeframe for reviews.
- Newcomers cannot always find the articles they created. They may recreate drafts, so that the same content is created and reviewed multiple times. This is worsened by having multiple article creation spaces (Main, userspace, Wikipedia.talk, and the recently-created Draft namespace²¹).

6.4 Drafts can be long-lasting and duplicated

It is both a success and a failure of the AfC process that drafts can be long-lasting. One poignant example is a draft that entered AfC in February 2011²². One failure is clear: Over three years later (as of May 2014), it remains unpublished. Review is not the problem: the draft has received 2 reviews, once 11 hours after it entered AfC. The success is that the draft might have been deleted under other procedures: In fact, the first AfC reviewer nominated it for speedy deletion due to no assertion of importance (A7). Within minutes, another reviewer removed the speedy deletion tag (implying that the topic is worth writing an improved article about).

What is going wrong? Collaboration on the draft has been limited: the creator has edited the page approximately one day a year, for instance adding a date of death when that became necessary. Three reviewers have looked at the draft, noting its flaws, but making no content changes. The only other edit was made by an IP editor (it is likely that the IP editor is the same creator, who logged in and edited the draft one minute later). The draft is forgotten, but not gone.

In fact, this the second of three attempts the creator has made to write about the same content. It is identical to a previous version (Draft1) started in the less-visible user namespace by the same editor 6 days after registering, in October 2010. Draft1 suffers from several problems and might have been deleted for copyright issues if it had more prominence, since it consists of uncited quotes from several sources. Draft1 was edited three times, all by the creator, with no other edits and no reviews.

A third attempt²³, Draft3, composed in the User/sandbox²⁴

¹⁶http://en.wikipedia.org/w/index.php?title=Talk\%3ANational_Security_Law_Journal&diff=605013877&oldid=604644067

¹⁷For example, see reviewers’ discussions of the need to recruit editors at https://en.wikipedia.org/wiki/Wikipedia_talk:WikiProject_Articles_for_creation/2014_2.

¹⁸http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation/feedback

¹⁹<http://en.wikipedia.org/wiki/Wikipedia:Notability>

²⁰https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation/Help_desk

²¹<http://blog.wikimedia.org/2013/12/20/new-draft-feature/>

²²http://en.wikipedia.org/wiki/Wikipedia_talk:Articles_for_creation/Dwight_K._Shellman,_Jr

²³http://en.wikipedia.org/wiki/Wikipedia_talk:Articles_for_creation/Dwight_Shellman

²⁴A userspace area intended for test edits and drafting.

in April 2012, is substantially rewritten and perhaps more promising. But in October 2013, the creator requested reviews for Draft2 and Draft3 on the same day. Draft3 was moved immediately into the namespace AfC uses, where reviews generally take place, and at that time, the moving editor made the only content edit any of these drafts have seen: the Wikidata person infobox was edited.

Since then the fate of both Draft2 and Draft3 have been similar. Both received a decline review, and both were flagged as postponed from G13 (see Section 4.1), indicating that there might be savable content. Despite the attention of 7 editors, only one content edit was made. The most clear feedback was given when declining Draft3, “because it is largely written about something else other than <the article title>”. So far no one has come along to make enough of the desired edits to make either of the AfC drafts a publishable article.

When we view AfC as a filtering process, it has some success. Content that needs significant work is protected from the encyclopedia’s harsh quality controls. AfC also protects the encyclopedia from problematic content, which is processed more effectively than un-patrolled userspace drafts.

Yet for helping and training newcomers AfC seems inadequate: the review paradigm of AfC puts the onus on the creator to update and fix content. Multiple types of expertise are needed [26], and perhaps drafts created by ‘substantive experts’ who do not have editing expertise could be improved by ‘technical editors’. In contrast to A7 deletion, which might remove an article within an hour, the AfC draft still has potential to be edited into a functional article, however this extra time is of limited use unless there is a mechanism for editors to find interesting drafts.

7. DESIGN IMPLICATIONS AND FUTURE WORK

Now we discuss design implications and future work. These mainly aim at improving collaboration, since AfC hides drafts from potential collaborators, and at creating more consistency in draft creation and quality control procedures, since current AfC procedures are confusing to newcomers and off-putting to experienced Wikipedians not yet familiar with them.

7.1 Determine which drafts need collaboration

As discussed above (Section 5.2.1), AfC drafts attract far fewer contributors than articles created in the Main (article) namespace. Since collaborators are the lifeblood of any open collaboration system (as discussed in Section 2), this situation needs to be corrected if we want to overcome the defects of the review-first model exemplified by AfC. This suggests **Design implication 1: Determine which drafts need collaboration.**

One possible solution is to make AfC drafts more visible to potential editors. Currently, the Wikipedia:talk namespace, where AfC drafts are located, is generally invisible to search engines (see footnote 13). Making draft articles available in searches would attract more pageviews, and, hopefully, more edits. A major concern about this approach, however, is that a large percentage of AfC drafts are promotional, or contain spam links. Making them more visible would degrade the quality, and potentially the reputation, of Wikipedia. A solution, therefore, would have to route only

potential collaborators, not all users, to AfC drafts that need improvement.

Intelligent task routing (ITR) could help [5]. ITR aims to match available workers to tasks that need attention in a way that maximizes output and worker satisfaction and minimizes wasted effort. Wikipedia, with its legion of available volunteers and multitude of tasks that need to be done, is a prime candidate for intelligent task routing. Its current ITR systems include SuggestBot²⁵, which uses ITR to recommend articles to be worked on to editors who opt-in, and Huggle²⁶ and STiki²⁷ [27], which help editors find possible vandalism to revert.

To use ITR, what’s missing is a model of drafts that are promising and could benefit from improvement by other Wikipedians. To determine whether a draft has potential, we propose a variation on Gelley and Suel’s [11] model for predicting whether an article will be deleted. Our model would use training data from the AfC reviews to identify drafts on intrinsically notable topics. Each review specifies the reason for decline (e.g. notability), so we know not only whether or not a draft was accepted, but also why it was declined.²⁸

With such a model, high-potential drafts would be passed to an intelligent task router that would identify editors who would be best qualified to improve them. We have begun preliminary discussions with the maintainers of SuggestBot about integrating AfC drafts into the recommendations that they make to users interested in working on articles in their areas of expertise. Integrating AfC drafts into SuggestBot’s recommendations would help bring promising AfC drafts to the attention of contributors with the right expertise to improve them. We note this as

Design implication 2: Recommend drafts to potential collaborators, based on both user interests and the likelihood of creating an encyclopedic article from a given draft.

7.2 Re-envision draft review workflows

Maintaining two separate article creation processes leads to several problems, as discussed in Section 6. We suggest unifying the default creation process for articles, and making Draft the default space for article creation for all users.

The community has a high tolerance for deleting material quickly,²⁹ and one frequent rationale is that article spaces are heavily linked to from outside the encyclopedia. Deleting new pages [8] has so far been a key technique for ensuring quality for readers. This has probably contributed to the belief that “the more curmudgeonly old-timers should be kept away from the newcomers until they have gained some experience in the system” [15], hence AfC. Yet newcomers seem to respond well to personalized constructive criticism from the New Page Patrol, [29] suggesting that perhaps a separate draft space for newcomers is not needed, after all. One area for future work is to understand why newcomers do not respond to article reviews, especially for declined articles.

²⁵<http://en.wikipedia.org/wiki/User:SuggestBot>

²⁶<https://en.wikipedia.org/wiki/Wikipedia:Huggle>

²⁷<https://en.wikipedia.org/wiki/Wikipedia:STiki>

²⁸http://en.wikipedia.org/wiki/Category:Declined_AfC_submissions

²⁹http://en.wikipedia.org/wiki/Wikipedia:Deletion_of_newly_created_pages

With a standard draft space for everyone, standardized, discrete checks would be feasible. Visible indications on drafts could show when and by whom they were last checked for various aspects such as copyright, notability, and point of view issues. Since the quality of articles and drafts changes over time, and because some checks are more specialized than others (as discussed above in Section 6.2) discretizing these checks would allow for more distributed work.

Creating all drafts articles in a single namespace would have the added benefit that an article would be created only once. Clear markings indicating that drafts are in a workspace with uneven quality could reduce concerns that reader-consumers might stumble across unfinished articles.

Experienced users could immediately move their drafts to the Main article space, and articles could be moved back to Draft as needed, with notification to the creator. This might reduce the application of certain speedy deletion rationales, such as A7, and ideally tighten their scope.

Design implication 3: Start all articles in the Draft namespace.

Design implication 4: Design a series of independent, discrete quality checks for drafts.

7.3 Support newcomers' first contributions

Problems around article creation for new and less experienced users are not likely to go away on their own; they need specific design attention. Some problems with newcomers writing articles are that:

- Newcomers often do not understand the standards of articles, especially notability, independence of sources, and neutral point of view.
- A significant number of new users 'drop content' and do not return. Even when an AfC review is completed within 4 minutes of its request, the original creator may not return to edit the article.
- Newcomers may have difficulty finding their drafts, and for instance may not be aware that they or anyone else can consult their contribution history [6].
- Newcomers have a variety of other technical problems, especially using inline references or reference templates, saving blank articles, etc.
- Many newcomers have additional problems, such as promotional usernames and conflicts of interest.

Writing a new article is one of the most visible possible actions on Wikipedia. In effect, Wikipedia's current interface pushes newcomers into registering accounts and writing articles, even when they have limited awareness of what makes an appropriate article. This does not seem beneficial. Merely having more accounts is not helpful for writing Wikipedia, especially when many accounts are quickly forgotten or abandoned: nudging people into registering on Wikipedia would not seem to have benefits. What is needed is more active editors, whether logged in or not.

Forefronting other contributions, besides creating articles, might help attract new editors to areas where they can truly collaborate. This suggests making contributions either much easier to make, or somewhat harder to make.

7.3.1 Make it easier to request articles.

Requesting articles is cumbersome. For instance, the Article Wizard, one of the most supportive tools for new editors drafting articles, has a prominent option, to "Request an article be written on a topic". Yet this takes the reader to a page with the standard information for Wikipedians: a screen full of instructions.³⁰ After following several steps to verify that the page doesn't already exist and is an encyclopedic topic, the requester would be asked to find the right category to which to add the topic they are requesting. This greatly deprives requesting articles over drafting them. A simpler request process could save time for reviewers as well as for creators who would prefer to deposit text rather than to iteratively write an article.

Design implication 5: Make it easier to request articles.

7.3.2 Collaborate with newcomers

Articles reach high quality faster when a cohesive, centralized network of collaborators are already working on them [18]. Social capital is important in order to find collaborators [18]: at least a few experienced users seem to be needed in order to write the best-quality articles (judging by German Wikipedia) [23]. To promote social capital, some systems such as Snuggle [14] and Teahouse [17] have designed for pro-social behavior. Observing the successes of these systems will help design more universal approaches to supporting and collaborating with newcomers.

Support for collaboration between newcomers and experienced Wikipedians might include:

- *Working together on process.* We recommend inviting all newcomers to Teahouse [17] when they first start drafting an article; Teahouse can provide newcomers with support negotiating complex processes.
- *Co-editing based on shared interest in a topic.* Newcomers could get personalized suggestions of people who have recently edited articles on similar topics. And experienced Wikipedians could get personalized recommendations, for instance by extending Suggest-Bot to include newcomers' drafts (#2 above).
- *Co-editing based on draft status.* Some editors have favorite tasks (copyediting, article rescue...). Making newcomers' first articles more prominent could encourage co-editing of newcomers' articles.

Design implication 6: Create mechanisms for newcomers to get help on drafts.

7.3.3 Improve the article creation advice

Currently, there are several ways to create a draft: The Article Wizard, a userspace draft, or starting with an empty textbox. In each case, some instructions are provided, ranging from five bullet points to a screen's length to multiple screens. This advice should be unified and consolidated. Messages to users have been found to be of great importance (e.g. [9]) and A/B testing of these widely used messages would have high impact.

More study is particularly needed of the Article Wizard, which provides scripted assistance in the steps to create an article. Can its successes and failures be identified and then

³⁰<https://en.wikipedia.org/wiki/Wikipedia:RA>

quantified? Are there obvious improvements that could be made in describing policies, or in providing intermediate steps for checking? For instance, would it be helpful to provide users with commentary about the overall quality of their sources, before they write an article? Many of AfC's articles come from the Article Wizard, and this is an obvious place to improve advice to newcomers.

Design implication 7: Improve the article creation advice, especially the Article Wizard.

8. CONCLUSION

Even small changes in the creation workflow are important, and the default choices make a huge difference. In this article we showed how newcomer productivity is reduced in English Wikipedia by pre-publication review. We showed that the review process is fast but has certain problems, most notably that AfC hides drafts from potential collaborators and that reviewing decisions can be hard to make correctly. We then suggested 7 design implications, motivated largely by our finding that *AfC hides drafts from potential collaborators*. Since the community has recently implemented a new Draft namespace, our design implications are particularly important for considering future processes for article creation and quality assurance. Wikipedians will have to determine priorities between wide coverage, high quality, and inclusiveness of future community members.

9. REFERENCES

- [1] P. André, R. E. Kraut, and A. Kittur. Effects of simultaneous and sequential work structures on distributed collaborative interdependent tasks. In *CHI '14*, pages 139–148.
- [2] S. L. Bryant, A. Forte, and A. Bruckman. Becoming Wikipedian. In *GROUP '05*, pages 1–10.
- [3] E. S. Callahan and S. C. Herring. Cultural bias in Wikipedia content on famous persons. *JASIST*, 62(10):1899–1915, 2011.
- [4] B. Collier and J. Bear. Conflict, criticism, or confidence: an empirical examination of the gender gap in Wikipedia contributions. In *CSCW '12*, pages 383–392.
- [5] D. Cosley, D. Frankowski, L. Terveen, and J. Riedl. Using intelligent task routing and contribution review to help communities build artifacts of lasting value. In *CHI '06*, pages 1037–1046.
- [6] H. Ford and R. S. Geiger. Writing up rather than writing down: Becoming Wikipedia literate. In *WikiSym '12*. Article No.16.
- [7] A. Forte and C. Lampe. Defining, understanding, and supporting open collaboration lessons from the literature. *American Behavioral Scientist*, 57(5):535–547, 2013.
- [8] R. S. Geiger and H. Ford. Participation in Wikipedia's article deletion processes. In *WikiSym '11*, pages 201–202.
- [9] R. S. Geiger, A. Halfaker, M. Pinchuk, and S. Walling. Defense mechanism or socialization tactic? Improving Wikipedia's notifications to rejected contributors. In *ICWSM '12*, pages 122–129.
- [10] R. S. Geiger and D. Ribes. Trace ethnography: Following coordination through documentary practices. In *HICSS '11*, pages 1–10.
- [11] B. S. Gelley and T. Suel. Automated decision support for human tasks in a collaborative system: the case of deletion in Wikipedia. In *WikiSym '13*. Article No. 28.
- [12] E. Goffman. *The presentation of self in everyday life*. Doubleday, 1959.
- [13] A. Halfaker, R. S. Geiger, J. Morgan, and J. Riedl. The rise and decline of an open collaboration system: How Wikipedia's reaction to sudden popularity is causing its decline. *American Behavioral Scientist*, pages 664–688, 2013.
- [14] A. Halfaker, R. S. Geiger, and L. G. Terveen. Snuggle: Designing for efficient socialization and ideological critique. In *CHI '14*, pages 311–320, 2014.
- [15] A. Halfaker, A. Kittur, and J. Riedl. Don't bite the newbies: How reverts affect the quantity and quality of Wikipedia work. In *WikiSym '11*, pages 163–172.
- [16] R. V. Kozinets. *Netnography: Doing ethnographic research online*. Sage Publications, 2010.
- [17] J. T. Morgan, S. Bouterse, H. Walls, and S. Stierch. Tea and sympathy: Crafting positive new user experiences on Wikipedia. In *CSCW '13*, pages 839–848.
- [18] K. Nemoto, P. Gloor, and R. Laubacher. Social capital increases efficiency of collaboration among Wikipedia editors. In *HT '11*, pages 231–240.
- [19] E. Raymond. The cathedral and the bazaar. *Knowledge, Technology & Policy*, 12(3):23–49, 1999.
- [20] J. Schneider, K. Samp, A. Passant, and S. Decker. Arguments about deletion: How experience improves the acceptability of arguments in ad-hoc online task groups. In *CSCW '13*, pages 1069–1080.
- [21] C. M. Schweik, R. C. English, M. Kitsing, and S. Haire. Brooks' versus Linus' Law: An empirical test of open source projects. In *dg.o '08*, pages 423–424. Extended version: NCDG Working Paper No. 07-009.
- [22] J. Solomon and R. Wash. Critical mass of what? Exploring community growth in WikiProjects. In *ICWSM '14*.
- [23] K. Stein and C. Hess. Does it matter who contributes: a study on featured articles in the German Wikipedia. In *HT '07*, pages 171–174.
- [24] B. Suh, G. Convertino, E. H. Chi, and P. Pirolli. The singularity is not near: slowing growth of Wikipedia. In *WikiSym '09*. Article No. 8.
- [25] F. B. Viégas, M. Wattenberg, and K. Dave. Studying cooperation and conflict between authors with history flow visualizations. In *CHI '04*, pages 575–582.
- [26] H. T. Welsler, D. Cosley, G. Kossinets, A. Lin, F. Dokshin, G. Gay, and M. Smith. Finding social roles in Wikipedia. In *iConf '11*, pages 122–129.
- [27] A. G. West, S. Kannan, and I. Lee. Stiki: an anti-vandalism tool for Wikipedia using spatio-temporal analysis of revision metadata. In *WikiSym '10*. Article No. 32.
- [28] A. G. West and I. Lee. What Wikipedia deletes: Characterizing dangerous collaborative content. In *WikiSym '11*, pages 25–28.
- [29] H. Zhu, A. Zhang, J. He, R. E. Kraut, and A. Kittur. Effects of peer feedback on contribution: a field experiment in Wikipedia. In *CHI '13*, pages 2253–2262.