

# Volunteer contributions to Wikipedia increased during COVID-19 mobility restrictions

Thorsten Ruprechter<sup>1,\*</sup>, Manoel Horta Ribeiro<sup>2</sup>, Tiago Santos<sup>1</sup>, Florian Lemmerich<sup>3</sup>, Markus Strohmaier<sup>3,4</sup>, Robert West<sup>2</sup>, and Denis Helic<sup>1</sup>

<sup>1</sup>Graz University of Technology, 8010 Graz, Austria

<sup>2</sup>EPFL, 1015 Lausanne, Switzerland

<sup>3</sup>RWTH Aachen University, 52062 Aachen, Germany

<sup>4</sup>GESIS – Leibniz Institute for the Social Sciences, 50667 Cologne, Germany

\*Corresponding author (ruprechter@tugraz.at)

**Wikipedia, the largest encyclopedia ever created, is a global initiative driven by volunteer contributions. When the COVID-19 pandemic broke out and mobility restrictions ensued across the globe, it was unclear whether Wikipedia volunteers would become less active in the face of the pandemic, or whether they would rise to meet the increased demand for high-quality information despite the added stress inflicted by this crisis. Analyzing 223 million edits contributed from 2018 to 2020 across twelve Wikipedia language editions, we find that Wikipedia’s global volunteer community responded remarkably to the pandemic, substantially increasing both productivity and the number of newcomers who joined the community. For example, contributions to the English Wikipedia increased by over 20% compared to the expectation derived from pre-pandemic data. Our work sheds light on the response of a global volunteer population to the COVID-19 crisis, providing valuable insights into the behavior of critical online communities under stress.**

Wikipedia is the world’s largest encyclopedia, one of the most prominent volunteer-based information systems in existence [18, 29], and one of the most popular destinations on the Web [2]. On an average day in 2019, users from around the world visited Wikipedia about 530 million times and editors voluntarily contributed over 870 thousand edits to one of Wikipedia’s language editions (Supplementary Table 1).

Amidst the COVID-19 pandemic and the “infodemic” [15] that ensued, Wikipedia played and continues to play an important role in supplying information about the COVID-19 crisis [9, 19, 45, 46]. Notably, the increase in access related to all kinds of articles—not only those related to the pandemic—suggests that Wikipedia’s role in this time of crisis transcends mere COVID-19-related information seeking [24]. However, page views are but a single aspect of the pandemic’s impact on Wikipedia, an aspect that ignores the fundamental contribution of editors, who perform unpaid volunteer work to maintain and develop content on the website. If the pandemic negatively impacted the productivity and number of editors on Wikipedia, the world’s largest online encyclopedia could be at peril [21, 33].

We can devise two competing hypotheses on how the COVID-19 crisis may have impacted editors on Wikipedia. First, the editor community may have shrunk in response to COVID-19 and corresponding mobility restrictions. As almost everyone, Wikipedia volunteers may have been affected by the negative economic and social ramifications of the pandemic [4, 7, 39], especially after most governments enforced mobility restrictions [11, 13, 54]. The challenges associated with this new reality may have led editors to withdraw from volunteer work for Wikipedia while focusing their efforts on personal issues and on dealing with the crisis. Alternatively, editors may have increased their volunteer work. This could be due to a personal response to the increased demand for high-quality information, as previously observed during locally confined disease outbreaks [1] and extraordinary events [52], or simply due to mobility restrictions resulting in individuals spending more time at home in front of computer screens [40] or on the Internet [12]. Whether Wikipedia editors withdraw from volunteering or increase their activity during distress determines the overall quality of information that Wikipedia serves to a global audience of readers. Therefore, understanding how editors responded to the COVID-19 pandemic and the accompanying mobility restrictions is crucial to assess Wikipedia’s capacity to act as a global information medium during worldwide disasters.

After careful quantitative analyses of large-scale edit logs on Wikipedia, we present robust evidence that *volunteer contributions significantly increased during the COVID-19 crisis* across many language editions. During the pandemic, the Wikipedia editor community not only generated many more edits than what we would expect given




Figure 1: **Edit volume in the English Wikipedia increased during COVID-19 mobility restrictions.** We visualize the rolling 7-day average edit volume in the English Wikipedia from January to October 2020 alongside the daily mean of 2019 and 2018, only considering non-bot edits to Wikipedia articles. Vertical lines mark major developments<sup>1</sup> during the COVID-19 pandemic in 2020. After the first Western countries (e.g., Italy) enforced mobility restrictions in early March, edit volume stagnated briefly before rising sharply—a trend that prevailed until late May, where the maximum difference in rolling 7-day average edit volume reached 20970. Although this initial sharp increase in edits declined, a surplus persisted until late September. Until September 31<sup>st</sup>, editors produced 8.4% (7.3%) more edits in 2020 than in 2019 (2018), an increase of 2.2 million (2 million) edits (Supplementary Table 2). Much of this edit surplus appears to stem from periods of mobility restrictions in the spring of 2020.

<sup>1</sup>Extracted from <https://wikimediafoundation.org/covid19/data>

historical baselines, but also acquired many more newcomers than in recent history, demonstrating the remarkable resilience of this online community in the face of adverse conditions.

Figure 1 depicts the increase in volunteer edits in the English Wikipedia during the COVID-19 timeline in 2020 compared to previous years. Whereas no increase in edit volume was apparent in early 2020, the mobility restrictions in Western countries seemed to first slightly dampen edit activity, before triggering a strong upward trend towards the end of March. In the weeks thereafter, a considerable edit surplus developed in comparison to previous years, which lasted until its peak in late May. As the pandemic subsided over the summer, the growth in edit volume also continuously decreased until fall. By October, the relative increase in edit volume, and thus volunteer contribution, from 2019 to 2020 (about 7.9%, or 2.1 million edits) was about double that from 2015 to 2019 (about 4.2%, or 1.5 million edits; see Supplementary Table 2). In summary, visual representation of edit volume in the English Wikipedia suggests a considerable contribution surplus in 2020.

Beyond the mere descriptive analysis of a single Wikipedia language edition, we systematically analyzed a varied sample of 12 Wikipedia language editions (“Wikipedias”), consisting of four large, medium, and small language editions each (Methods), with over 223 million edits spread through 24.6 million articles. In accordance with the descriptive analysis shown in Figure 1, our quasi-experimental difference-in-differences analysis finds a significant increase in edit volume after COVID-19 mobility restrictions came into effect for many of the Wikipedia editions, and an influx of new editors that is particularly salient for larger Wikipedias. Our study sheds light on the impact of the COVID-19 mobility restrictions on Wikipedia volunteer contributions and provides a reusable framework to measure user activity under stress. More broadly, the evident increase in edit volume and newcomers across most




Figure 2: **Edit volume during COVID-19 mobility restrictions.** We show edit volume findings in large (top), medium (middle), and small (bottom) Wikipedias during COVID-19 mobility restrictions, which in the figure we delineate using mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We show rolling 7-day average edit volume generated by human editors for 2018, 2019, and 2020 until October. After a slight retraction of editing around mobility changepoints in most Wikipedias, the number of contributions recovers to previous levels within a few days. Editors contribute substantially more in all large and some medium Wikipedias in the weeks after the mobility restrictions in 2020, compared to historical baselines. **b**, We depict the relative change in edit volume ( $ev$ ) as retrieved from DiD via  $\delta_{ev}$  (95% confidence interval as two standard deviations) and plot  $\delta_{ev}$  for 120 left-aligned seven-day-windows (see Methods), with the x-axis describing days after the respective mobility changepoint. We observe that editing in large and medium Wikipedias significantly increases after their mobility changepoint, while most small Wikipedias show neither significant increase nor decrease.

observed Wikipedias is a finding of interest not only to Wikipedia itself but also to researchers and managers of other online collaboration systems, as it provides valuable insight into user behavior during a global crisis.

## Results

### Edit volume during COVID-19 mobility restrictions

We observe an increase in edit volume (the number of edits made by non-bot users) on Wikipedia during the period of COVID-19 mobility restrictions in the spring of 2020, which is particularly evident in large and medium Wikipedias. Figure 2a depicts the rolling 7-day average edit volume for large (top), medium (middle), and small (bottom) Wikipedias in the context of COVID-19 mobility restrictions, which we delineate via automatically detected mobility (i.e., restrictions take effect) and normality (i.e., restrictions are lifted) changepoints (see Methods). We also report edit volumes for 2018 and 2019 as a reference for 2020. We observe substantial drops in edit volume around the mobility changepoint for almost all Wikipedias, indicating a shock to the Wikipedia ecosystem. In particular, larger Wikipedias experience a considerable short-lived decrease in edit volume but are able to recover quickly. English, Italian, German, French, Korean, and Japanese even clearly surpass their pre-shock volume levels, leading to an overall edit surplus. On the contrary, some smaller Wikipedias (e.g., Finnish) exhibit a steady decline in edit volume after the mobility changepoint. To better relate edit volume during the COVID-19 pandemic to reference values

from previous years and pre-pandemic periods, we employ a difference-in-differences regression (DiD) that controls for the year, period, and language, as well as their interactions. For all Wikipedias, we compute the effective change in edit volume ( $ev$ ) after the mobility changepoint from the three-way interaction of year, period, and language, and denote this effective change as  $\delta_{ev}$ . We apply the DiD analysis to a sequence of seven-day-windows post-changepoint, always retaining the 30-day pre-changepoint period, and plot the time series of logarithmic effects for edit volume according to  $\delta_{ev}$  in Figure 2b. We describe this DiD setup in more detail in Methods. The DiD analysis validates that all large and most medium Wikipedias significantly increase their edits following the mobility restrictions according to  $\delta_{ev}$  (95% confidence interval), while no general statement can be made for small Wikipedias.

For the rolling 7-day average edit volume in large Wikipedias, we identify an upward trend in 2020 immediately after mobility restrictions took place (Figure 2a, top). In the English, French, and Italian language editions, edit volume steadily increases for nearly two months after a dip around the date of the mobility changepoint, before slowly reverting to prior levels. The steady initial increase in edit volume leads to outstanding peaks—approximately 120 000 edits for English, 28 000 for French, 15 000 for Italian, and 24 000 for German, which exhibits a decline back to pre-crisis levels earlier than other large Wikipedias. DiD results confirm the edit volume surplus visible in the time series for large Wikipedias in 2020 (Figure 2b, top).  $\delta_{ev}$  for French, Italian, and English depicts an immediate relative increase in edits after the mobility restrictions take place, leading to over 100 days of significant increases for all three of these Wikipedias, whereas German declines earlier. Approximately 35 days after the mobility restrictions take effect, French ( $e^{0.337} = 144\%$ , a surplus of 44%), Italian (+42%), and German (+25%) reach their highest significant relative increase for edit volume. The higher short-term increases in French, German, and Italian may be related to more detailed reporting of local issues in these language editions. On the contrary, English shows a longer, sustained upward trend for  $\delta_{ev}$ , with a maximum significant increase of 23% after 69 days. In conclusion, edit volume significantly increases in large Wikipedias after mobility restrictions come into effect.

Edit volume in most medium and small Wikipedias slightly drops around the respective mobility changepoints in 2020. However, virtually all Wikipedias quickly recover from the initial shock, with most maintaining a stable edit volume in the ensuing weeks and some even generating an edit surplus. While Figure 2a (middle) shows that medium Wikipedias do not homogeneously increase their edit volume, Korean and Japanese surpass their pre-mobility-restriction levels about a month post changepoint, peaking at about 5 400 and 14 500 edits, respectively. For small Wikipedias, edit volume only decreases slightly right after the mobility changepoint (Figure 2a, bottom). Afterward, edit volume recovers to previous baselines within thirty days, before following similar trends and levels as in previous years. DiD analysis and corresponding values for  $\delta_{ev}$  reveal that, in fact, medium Wikipedias experience varying periods of significant relative increases in edit volume (Figure 2b, middle). For example, when compared to pre-pandemic years around the same time period, the Korean and Dutch Wikipedias produce a consistent relative increase (peaking at +40%), whereas Swedish and Japanese exhibit shorter significant periods (+30% and +38% in maximum, resp.). Furthermore, the relative change for small Wikipedias (Figure 2b, bottom) signals brief periods of substantial relative increases for Danish and Norwegian (peaks of +69% and +43%, resp.). Most notably, Serbian exhibits a considerate increase during the first month after mobility restrictions take place, with volume nearly tripling (logarithmic effect of 1.03). Lastly, we note that out of our twelve investigated Wikipedias only Finnish shows a significant decrease in  $\delta_{ev}$  over longer stretches of the observed period. In any case, small and medium Wikipedias are mostly resilient to the initial shock to edit volume triggered by COVID-19, with some even surpassing their pre-pandemic baselines after a few weeks.

## Newcomers during COVID-19 mobility restrictions

We find that all large and medium Wikipedias acquire considerably more newcomers (the number of registered users who made their first edit) for most of the study period, while the remaining Wikipedias exhibited resilience and do not decrease their levels significantly. We visualize the 7-day rolling averages for newcomer counts during the COVID-19 pandemic for large (top), medium (middle), and small (bottom) Wikipedias in Figure 3a, while also showing values for previous years as well as mobility and normality changepoints. Newcomer counts plummet around the mobility changepoint, in particular for large Wikipedias, but this attenuation in newcomer recruitment only persists for a brief period. Shortly thereafter, newcomer counts increase considerably in all but a few medium and small Wikipedias (e.g., Swedish or Finnish). Again, we build a DiD model for newcomers ( $nc$ ) to quantify effective changes during the period of COVID-19 mobility restrictions in spring 2020, again controlling for year, period, and language. We again perform our DiD analysis for a sequence of seven-day windows after the mobility changepoint (see Methods) and show the logarithmic effects for newcomers ( $\delta_{nc}$ ) in Figure 3b. This newcomer DiD analysis confirms that while all large Wikipedias acquire significantly more new editors after mobility restrictions take effect, some medium and small Wikipedias seem to be resilient and exhibit no significant long-term changes (95% CI).

Large Wikipedias appear to recover rapidly from the initial negative effect of mobility restrictions in terms of




Figure 3: **Newcomers during COVID-19 mobility restrictions.** We visualize newcomer results in large (top), medium (middle), and small (bottom) Wikipedias during COVID-19 mobility restrictions, which we delineate via mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We depict rolling 7-day average newcomer counts until October of 2018, 2019, and 2020. For many Wikipedias, newcomer acquisition strongly declines right around their mobility changepoint, but then quickly rises to or even exceeds pre-pandemic baselines. **b**, We investigate the relative change in newcomers ( $nc$ ) via  $\delta_{nc}$  as computed from DiD analysis (95% confidence intervals as two standard deviations) and plot  $\delta_{nc}$  for 120 left-aligned seven-day windows (see Methods), starting with the respective mobility changepoint. In large Wikipedias, considerably more newcomers join in the weeks after mobility restrictions come into effect, relative to before the changepoint and previous years. Results for medium and small Wikipedias are non-conclusive, with some showing increases in the number of newly acquired editors and others not significantly changing their values.

newcomer counts (Figure 3a, top). Most notably, Italian registers a downright newcomer surge until late April, recruiting over 150 newcomers on a rolling 7-day average. English and French show similar patterns of perpetual increases, reaching respective peaks of approximately 2100 and 330 new editors. Although German exhibits nearly 200 newcomers shortly after the mobility changepoint, the surplus in 2020 seems not as considerable as for other large Wikipedias. We further note that newcomer counts for large Wikipedias start to steadily decline in May. However, this seasonal trend also appears to be prevalent in previous years. Our DiD analysis, which captures the change in newcomers via  $\delta_{nc}$ , for the most part confirms these findings (Figure 3b, top). During the first two to three weeks past the mobility changepoint, large Wikipedias steadily recover from the COVID-19 shock without significant overall gains according to  $\delta_{nc}$ . However, right after this recovery phase significant peaks arise for English ( $e^{0.283} = 130\%$  of previous levels), French (138%), and German (139%). For the Italian Wikipedia, which belongs to a region with particularly strict mobility restrictions, we confirm an even stronger newcomer surge, leading to a 80% relative increase. Furthermore, English generates a notably stable, significant long-term growth in newcomers that is possibly owed to editors from all over the world joining this language edition during mobility restrictions in their regions, as the English Wikipedia serves as a global repository of knowledge. Ultimately, positive effects prevail for large Wikipedias and solidify a newcomer surplus after the mobility restrictions come into force.

Similar to large Wikipedias, most medium and small Wikipedias experience a decline in newcomers right around their mobility changepoints before then increasing their counts to previous baselines (Figure 3a, middle and bottom). Some of these Wikipedias (e.g., Norwegian, Finnish, Danish, Swedish) recover to previous levels within the first

month and exhibit no long-term effects afterwards. However, others recruit a surplus of newcomers during this crisis. Japanese, Dutch, Korean, and Serbian show short-term newcomer influxes about one to two months after the initial mobility restrictions take effect, with maximum respective values of approximately 170, 60, 50, and 30 daily newcomers. We also observe these effects in  $\delta_{nc}$  as captured by DiD (Figure 3b, middle and bottom), which confirms brief relative increases for Japanese (+30%), Dutch (+47%), Korean (+28%), and Serbian (+179%). Finally, the newcomer DiD analysis corroborates that some medium and most small Wikipedias do not significantly deviate from baselines prior to the mobility restrictions over much of the observed time span.

## Discussion

As the COVID-19 pandemic erupted on a global scale, it was unclear how this incisive event would affect Wikipedia’s volunteer community. Over the course of the last few years, both human editing [38] and newcomer recruitment [22] on Wikipedia have stagnated or even decreased (Supplementary Table 2). Accordingly, the pandemic could have accelerated the decline of the online encyclopedia as the hardships of this global crisis may even further decrease volunteer activity. However, our study, in which we analyze 223 million edits from 12 Wikipedia language editions, reveals that the COVID-19 pandemic and its accompanying mobility restrictions have substantially boosted volunteer activity on Wikipedia. By performing a difference-in-differences analysis, we show that edit volume as well as the influx of newcomers has generally *increased* after COVID-19 mobility restrictions went into effect. In what follows, we discuss the implications and limitations of this finding.

**Mechanisms behind contribution growth.** We observe significant increases in edit volume and newcomers during the COVID-19 pandemic across multiple Wikipedias, making it their most active period in at least the last three years. While our quantitative study sheds light on the extent of contribution growth, there are several possible mechanisms behind this effect, which may or may not impact the collaborative structure of editor communities.

Firstly, Wikipedia received significantly more page views during the COVID-19 crisis [24]. The increase in edits and newcomers may partially be due to the prior increase in Wikipedia readership, as a certain proportion of readers turns into contributors because of various motivational factors [37, 48]. In addition, we theorize that increased screen time and Internet exposure [12, 40] during the mobility restrictions lead to Wikipedia readers spending more time editing, possibly increasing the reader-to-editor turnover rate. Tracing the transformation of readers into editors during this pandemic in more detail is a promising avenue for future work.

Secondly, the increase in contributions may be due to the rapidly changing information and new knowledge that the COVID-19 pandemic generates about the world. Past literature has suggested that Wikipedia growth is constrained by the amount of knowledge available, as editors have already contributed most of the easily obtainable and verifiable information [38]. The fact that volunteers have been “running out of easy topics” to contribute to has made it difficult for non-specialists to provide new content with little effort [21]. As the COVID-19 pandemic dramatically changes the status quo of our world today, it is generating new knowledge about many fields and thus may provide fresh opportunities for both novel and veteran editors to contribute to Wikipedia.

Moreover, the observed edit surplus may have been caused by the high-intensity activity of a core group of editors rather than the broader editor population. We therefore investigate the number of editors active on any given day according to their activity level: 1 to 4, 5 to 24, 25 to 99, or more than 99 daily edits (Methods). The DiD analysis for editor counts depicts increases across all activity levels after mobility changepoints for all large and most medium Wikipedias, while small Wikipedias show non-conclusive effects (Supplementary Figs. 1, 2, 3, and 4). These findings indicate that the editor population as a whole intensified their contribution during the COVID-19 pandemic, causing the overall increase in volunteer activity.

Finally, we detected a contribution disparity with respect to Wikipedia size, meaning that the smaller Wikipedias we studied did not benefit to the same degree as larger or medium Wikipedias. The observed discrepancy in edit and newcomer increases for large, medium, and small Wikipedias may stem from a difference in community size and structure, or these Wikipedias’ specific rules [5, 21, 26, 34, 49]. Moreover, the amount of content for certain topical categories diverges due to cultural contextualization in different language editions [32]. Specifically, a strong (hypothetical) affinity for topics not directly related to the pandemic (e.g., *Sports*) in medium or smaller Wikipedias might change the effect of this crisis on their edit volume, in comparison to larger Wikipedias. As an example, in case such Wikipedia language editions focused more on updating sports articles, edit volume would decrease more during the pandemic. The magnitude of such an effect may further depend on a region’s more (e.g., Italy) or less strict (e.g., Sweden) mobility restrictions. Future research may explore language-specific collaboration mechanisms in more detail, for example by attempting to topically analyze Wikipedia contributions during the pandemic.

**Resilience of Wikipedia communities.** Although we did not find the same surplus in contributions across large, medium, and small Wikipedia language editions, volunteer communities in all studied Wikipedias demonstrated

resilience by quickly recovering from the initial negative impact of the pandemic on their contributions. While slow response to negative events or other shocks causes severe problems in social-ecological systems [35, 36], resilient systems are adaptable and manage to withstand such shocks, even bearing the capacity to cross previous performance thresholds [14]—a behavior observed in this study. The strongest resilience and subsequent crossing of earlier thresholds in large Wikipedias during the pandemic may be partially explained by the difference in community size [51]. For example, in larger communities it may not be as problematic that leaders are limited due to the pandemic, as a greater number of other veteran members can take over their work. This conjecture borrows from critical mass theory, in the sense that a critical mass of core members is the fundamental source of content [48]. Future research might investigate the aspect of Wikipedia resilience during the pandemic in more detail, for example by considering threat rigidity [51] or building a model [47] that considers COVID-19 as an attack on the community structure.

**Revert rate during COVID-19 mobility restrictions.** The observed simultaneous increase in newcomers and edits may suggest that the edit surplus was partially caused by low-quality edits by first-time editors. Frequently, veteran editors or bots would then completely undo (i.e., identity revert) these newcomer revisions, which represents a common behavioral pattern on Wikipedia [21, 22, 53], in turn generating further revisions. To investigate whether an increase in such reverts occurred, we performed a cursory analysis of the revert rate, which is defined as the ratio of reverted edits to edit volume (see Methods and Supplementary Information). Supplementary Figure 5a visualizes the rolling 7-day average revert rate ( $rr$ ), while Supplementary Figure 5b plots the relative change in revert rate ( $\delta_{rr}$ ) as captured by a DiD analysis (Methods). Interestingly, we detect a significant increase of the revert rate in only one language (Korean). By contrast, several Wikipedias exhibit significantly decreased revert rates shortly after the mobility restrictions come into force. For example, the large Italian, French, and German Wikipedias all show reduced revert rates by about one quarter. This suggests that less valuable revisions, possibly made by newcomers, and their immediate reversal do not cause the reported increase in edit volume. Furthermore, potential misbehavior or conflict on Wikipedia, such as vandalism or edit wars, is prominently characterized by large numbers of identity reverts, as they undo these unwanted contributions [28, 41, 50]. Therefore, reduced revert rates may indicate that editors refrain more from confrontational behavior and thus demonstrate higher levels of solidarity during the pandemic, which is a common phenomena within collectives during crises [16]. However, a decline in revert rate could also imply that bots and administrators may be unable to keep up with the influx of edits, leaving low quality or malicious edits undetected and thus diminishing quality in the long term. We see the detection and analysis of behavioral patterns and collaborative structure of online communities as a promising path for future research. In addition, it may be valuable to further study the treatment and retention of newcomers [8, 22, 34] during and after the pandemic once more longitudinal data is available.

**Contribution to COVID-19 articles.** One might speculate that the increase in edit volume is mostly due to edits in articles that are strongly related to COVID-19. However, many of those articles were protected from public editing early in the pandemic to prevent spread of misinformation [27], and we find that only a negligibly small fraction of edits (at most 1% for most Wikipedias) goes towards articles with a primary focus on COVID-19 (see Methods) between January 1<sup>st</sup> and September 31<sup>st</sup> 2020 (Supplementary Table 4 and Supplementary Fig. 6). A clear outlier in that regard is German, where 2.5% of edits performed in 2020 by the end of September concern themselves with such articles. This may indicate higher coverage of local COVID-19 outbreaks in German than in other languages. We consequently repeat our DiD analysis for edit volume, this time excluding edits to articles strongly related to COVID-19 (Supplementary Fig. 7). The results support the previous findings and confirm that the reported edit volume increase is not due to COVID-19 articles. In this way, our work extends previous studies, which focused on a smaller subset of pandemic-related articles [19, 27].

**Other limitations.** Even though our work covers a large portion of Wikipedia’s content and editor population, it comes with several limitations. First, we do not consider a variety of different Wikipedias associated with languages widely spoken in the global south, including Spanish, Portuguese, Arabic, Hindi, or any African Wikipedias (see Methods for how we chose language editions). Future work analyzing these Wikipedias could improve our understanding of the impact of the pandemic on volunteer contribution in other parts of the world. Second, content on Wikipedia is predominantly edited by white males between the ages of 17 and 40 [10, 23]. It may be that the COVID-19 crisis has disparately impacted contributors of less represented demographics, as certain racial or socioeconomic groups are particularly disadvantaged by the pandemic [3, 6, 25]. In addition, bots have an important role in the creation and management of Wikipedia content [43, 53]. We excluded bots from our analysis as we specifically focused on edits performed by human volunteers. Nevertheless, other studies may choose to consider bot activities as valid contributions to Wikipedia.

In conclusion, our study provides evidence for a substantial surplus of volunteer contributions to multiple Wikipedia language editions during COVID-19 mobility restrictions, which shines light on the resilience of the Wikipedia community under times of stress. The methodological framework used in this work can easily be adapted



for similar domains. We believe that our work provides valuable insights into contributor behavior on online platforms during the COVID-19 pandemic and illustrates a plethora of possibilities for future work.

## Methods

**Data procurement and preprocessing.** We utilize the openly available MediaWiki history dataset dumps to analyze a varied sample of 12 Wikipedia language editions (“Wikipedias”).

*Wikipedia language editions.* We investigate 12 Wikipedias (Supplementary Table 3), consisting of languages primarily spoken in European countries that were exposed to the outbreak of COVID-19 in the spring of 2020, as well as two Asian Wikipedias. Our choice of language editions takes into consideration: (i) the size of the Wikipedia edition, (ii) whether the language is spoken in relatively few countries, and (iii) the mobility restrictions imposed in these countries—three criteria that are often very difficult to simultaneously satisfy. Overall, we aim to capture relevant Wikipedias that represent different attitudes towards the crisis, preferably from languages easily attributable to a single country or region. Accordingly, our sample contains regions with strict (e.g., Italian, Serbian, or French) and less stringent mobility restrictions (e.g., Japanese, Korean, or Swedish). Although it can not be attributed to a single country, we include English as it is the largest language edition. We employ the number of edits in 2019 as a metric to categorize the 12 Wikipedias we studied as either *large* (English, French, German, Italian, with more than 5 million edits), *medium* (Swedish, Korean, Japanese, Dutch, with 1.5 million to 5 million edits), or *small* (Serbian, Norwegian, Danish, Finnish, with less than 1.5 million edits).

*MediaWiki history dataset dumps.* We retrieve the monthly updated MediaWiki history dataset dumps<sup>1</sup> provided by the Wikimedia Foundation (WMF) and perform additional preprocessing before computing as well as plotting our results. The denormalized MediaWiki history dumps are generated from the full history logs stored in the WMF’s MediaWiki databases. During their generation, WMF’s automatic scripts reconstruct and enrich user and page history with additional data, and also automatically validate the dumps to prevent errors. After WMF’s preprocessing, the dataset contains fields with precomputed standard metrics, such as revert information, bot users, number of user contributions, or time since a user’s last revision. The technical documentation on Wikitech<sup>2</sup> closer describes the dataset dumps’ schema and contained fields. Overall, each entry in the dump consists of 70 fields with event information. Fields are grouped into entities, which bear information about either **revision**, **page**, or **user**.

*Preprocessing.* In the MediaWiki history dataset, we only consider edits to articles by excluding all pages not in the Wikipedia article namespace (“ns0”), thus removing revisions to talk pages or other content. Furthermore, we utilize corresponding dataset fields to distinguish human editors (anonymous or registered) from bots and mark certain revisions as reverts. Moreover, we convert MediaWiki history timestamps from Coordinated Universal Time (UTC) to the timezone of the local Wikipedia language edition. For Wikipedias in which languages can not be attributed to a single timezone (e.g., French), we choose the timezone with the highest volunteer population for the given Wikipedia. We do not apply timestamp conversion for the English Wikipedia. Lastly, we detect articles which are strongly related to COVID-19 via an algorithm by Diego Sáez-Trumper<sup>3</sup>, which recognizes COVID-19 articles based on their Wikidata [44] links to the main COVID-19 pages.

**Metrics.** To make sense of which exact data fields in the MediaWiki history dumps we utilize to compute our metrics, please refer to the code repository (see Code availability).

*Edit Volume.* We define edit volume as the number of daily revisions to pages in the article namespace (“ns0”) by non-bot users (anonymous or registered).

*Newcomers.* For each Wikipedia language edition and day, we specify the amount of newcomers as the number of registered editors which perform their first article edit in that Wikipedia language on the given day. Through recognizing new editors by their first edit, we measure the exact day they become a contributor in a language edition. Note that the number of daily registered users is generally much higher than the number of newcomers as computed in this work. However, as our study aims to quantify volunteer contribution, we choose to identify newcomers by their first actual contribution in a given Wikipedia.

*Revert rate.* Editors and bots revert article revisions to undo changes which they deem unwarranted. Frequently, these reverts correct revisions which arise from conflicts, edit wars, or vandalism [50]. Additionally, literature shows that revisions by newcomers are more likely to be reverted than those of veteran editors [22]. For this research, we only consider reverts to articles that undo all changes and subsequently create a new revision which exactly matches a previous article version (i.e., identity reverts). We calculate the daily revert rate by dividing the number

---

<sup>1</sup>[https://dumps.wikimedia.org/other/mediawiki\\_history/readme.html](https://dumps.wikimedia.org/other/mediawiki_history/readme.html)

<sup>2</sup><https://w.wiki/uzW>

<sup>3</sup><https://covid-data.wmflabs.org/>



of identity reverts (by humans or bots) by the number of non-bot edits on this given day. Correspondingly, revert rate relates the amount of reverts to the amount of human contribution.

*Daily editors by activity level.* We measure daily active editors in a Wikipedia by counting the number of registered, non-bot users which perform revisions in the article namespace. To detect effects across the editor population, we collect data for multiple activity levels, keeping count of how many editors perform 1 to 4, 5 to 24, 25 to 99, or more than 99 daily edits. In contrast to other metrics, we do not compute the number of daily editors from the Wikimedia history dumps, but retrieve it via the Wikimedia REST API<sup>4</sup> instead.

**Changepoint detection.** We adopt the approach by Horta Ribeiro et al. [24] to automatically detect *mobility* and *normality* changepoints via Google and Apple mobility reports.<sup>5</sup> These reports capture population-wide movement patterns based on cellphone location signals and specify, on a daily basis, the percentage of time spent in variety of locations (e.g., residential areas, workplaces, or retail). Government-mandated lockdowns and self-motivated social distancing measures manifest themselves as sharp changes in these mobility time series. To detect changepoints in mobility, the approach consists of a simple binary segmentation algorithm [42]. For Wikipedias of languages widely spoken across many countries (e.g. English, German, etc), we determine a changepoint by aggregating mobility reports for the countries in which the language is official with weights proportional to the population of each of these countries. Notice that the link between Wikipedia and language editions is merely approximate—in particular for English, which is accessed from all over the world. We use the changepoints at which mobility drops as heuristics for dates when people started spending substantially more time in their homes and term them *mobility changepoints*. To detect *normality changepoints*, we compute the point in time for which the future average mobility remains within a 10% band around baseline levels before the initial mobility changepoint (defined as pre-pandemic mobility levels by Google and Apple). For languages spoken across multiple countries, we maintain the same aggregation scheme as before. Compared to choosing specific dates, this changepoint detection approach leads to more comparable treatments across different regions. Supplementary Table 3 summarizes the detected changepoints for the investigated Wikipedias, which we also make available in our code repository.

**Difference-in-differences setup.** To compare values of metrics during the COVID-19 pandemic with reference values from previous years and pre-pandemic periods, we employ a difference-in-differences regression (DiD). DiD allows us to quantify changes in these metrics in multiple Wikipedia language editions around times of region-specific mobility changepoints in early spring, while controlling for (long-term) temporal trends.

Our basic DiD equation models a dependent variable’s value ( $V$ ) as a function of the independent variables year ( $Y$ ), period ( $P$ ), and Wikipedia language ( $L$ ), as well as their interactions. Year is a binary variable which differentiates between pre-pandemic (2018 and 2019) and pandemic years (2020), whereas period encodes the treatment period via a binary variable, in our case represented by the pre- and post-phases of the region-specific mobility changepoints. Lastly, we model our 12 Wikipedia language versions with a categorical variable to control for language-specific effects. To account for outliers and normalize regression results across various-sized Wikipedias, we use logarithmic scales for  $V$ . Literature often refers to our setup, which uses three independent variables, as “triple-difference” or “difference-in-difference-in-difference” estimators [20, 31]. Mathematically, our DiD setup is:

$$V = \beta_0 + \beta_1^\top \mathbf{L} + \beta_2 Y + \beta_3 P + \beta_4^\top (YL) + \beta_5^\top (PL) + \beta_6 (YP) + \beta_7^\top (YPL) + \varepsilon \quad (1)$$

We depict the 12 Wikipedia language versions as a vector of 11 binary indicators ( $\mathbf{L}$ ). Scalar coefficients ( $\beta_0$ ,  $\beta_2$ ,  $\beta_3$ ,  $\beta_6$ ) describe effects for the reference language (i.e., baseline). Coefficient vectors ( $\beta_1$ ,  $\beta_4$ ,  $\beta_5$ ,  $\beta_7$ , printed in bold) collect language-specific effects of non-baseline Wikipedias. Lastly,  $\varepsilon$  is the normally distributed residual. Given this mathematical formulation, the coefficient  $\beta_7$  captures the change in  $V$  post mobility changepoint relative to the baseline Wikipedia, after accounting for differences stemming from year or period alone ( $\beta_4$  and  $\beta_5$ , resp.). We therefore compute the effect of interest for all Wikipedias via summation of  $\beta_6$  and  $\beta_7$ . For each Wikipedia, we term this effective change in  $V$  as  $\delta_m$ , where  $m$  stands for the metric representing the dependent variable.

*Interpretation of DiD coefficients.* We now elaborate in more detail on how to interpret the coefficients of our DiD model. We model the categorical language variable via vector  $\mathbf{L}$  containing 11 binary indicator variables for the 12 Wikipedias. As is customary, the regression utilizes a “reference Wikipedia” baseline, which is represented by the intercept of the model given  $Y = 0$  and  $P = 0$ . In our setup, we arbitrarily choose Danish as the baseline. Consequently,  $\beta_1$  describes the respective difference between the baseline Wikipedia and the 11 non-baseline Wikipedias using indicator variables. Thus, adding  $\beta_0$  and  $\beta_1$  yields the intercept of each language’s sub-model.

The binary year variable ( $Y$ ) indicates whether a data point lies in 2020 ( $= 1$ ) or in the previous two pre-pandemic years ( $= 0$ ), regardless of period. As Danish represents the arbitrary baseline, the corresponding coefficient  $\beta_2$  is a scalar which describes the overall change between the pre-pandemic years (2018 and 2019) and 2020 for Danish. For

<sup>4</sup>[https://wikimedia.org/api/rest\\_v1/](https://wikimedia.org/api/rest_v1/)

<sup>5</sup><https://www.apple.com/covid19/mobility>

non-baseline Wikipedias, the interaction  $YL$  models the language-specific effects for the change in years relative to the baseline Wikipedia and is quantified by the corresponding coefficient vector  $\beta_4$ . Therefore, the summation of  $\beta_2$  and  $\beta_4$  is equal to the effective overall difference of 2020 to the previous two years for all Wikipedias.

We model seasonal differences between pre- and post-changepoint windows via the binary period indicator ( $P$ ). The corresponding scalar coefficient ( $\beta_3$ ) measures the difference between before and after the mobility changepoint over all years for the baseline. Consequently,  $PL$  and coefficient vector  $\beta_5$  describe the period effect for non-baseline Wikipedias in relation to the baseline. Calculating the sum of  $\beta_3$  and  $\beta_5$  then gives the total pre- and post-changepoint effects.

Lastly, the interaction between year and period ( $YP$ ) enables our model to capture the change in  $V$  for the baseline Wikipedia via  $\beta_6$ , after accounting for change in  $Y$  (via  $\beta_2$ ) and  $P$  (via  $\beta_3$ ) alone. To measure this effective change for all Wikipedias, we employ the coefficient vector  $\beta_7$  of the three-way interaction  $YPL$ . While  $\beta_6$  describes the baseline’s effect,  $\beta_7$  contains the aforementioned change relative to the baseline Wikipedia. Therefore, the sum of  $\beta_6$  and  $\beta_7$  captures the effective change in  $V$  for all Wikipedias. For a single Wikipedia  $l$  and metric  $m$ , we name this effect of interest  $\delta_m$ . Correspondingly,  $\delta_m$  describes language-specific post-changepoint effects in 2020, as it excludes differences that are due to year or period alone.

**Quantifying changes in volunteer contribution.** Wikipedia is a dynamic ecosystem, in which edit behavior and the amount of volunteer contribution can change rapidly—especially in times of turmoil. To track these changes and detect short-, medium-, and long-term effects of mobility restrictions on volunteer contributions, we fit our statistical model on different data-points obtained from the same longitudinal dataset. This methodology, pioneered by Gelman and Huang [17], allows us to observe trends rather than mere point estimates.

We compute our DiD analysis for a sequence of post-changepoint windows, always retaining the Wikipedias’ pre-changepoint periods. For each language version, we choose a fixed 30-day period before the respective mobility changepoint as the pre-changepoint baseline. As post-changepoint analysis intervals, we then extract a sequence of 120 overlapping left-aligned seven-day-windows starting with the changepoints. Mathematically, we set the treatment period to days  $\{n, n + 1, \dots, n + 6\}, \forall n \in \{0, 1, \dots, 119\}$ . For each post-changepoint window  $n$ , we perform a separate DiD analysis across all languages using the retained baseline periods. By doing so, each DiD analysis compares the week starting at day  $n$  after the language-specific changepoint to the baseline periods. In this default setup, each of the 12 Wikipedias is represented by 37 data points for every year in the DiD regression (2018, 2019, and 2020), yielding a total of 1332 data points ( $= (30 \text{ pre-changepoint days} + 7 \text{ post-changepoint days}) \times 3 \text{ years} \times 12 \text{ Wikipedias}$ ) for each of the 120 experiments. For each Wikipedia, we conservatively detect outliers via the Median Absolute Deviation (MAD) approach [30] with a threshold of  $5 * \text{MAD}$  from the monthly median and replace such outliers by the monthly median. We then build a time series of the 120 DiD results using  $\delta_m$  and approximate the 95% two-sided confidence intervals (CI) as two standard errors. As robustness checks, we compute variations of our DiD experiments with wider window size (14 days) and slightly varied mobility changepoint dates ( $\pm 7$  days) as described in Supplementary Information (Supplementary Figs. 8, 9, 10, 11, 12, 13, 14, 15, and 16). These results corroborate the findings reported under Results.

## Data availability

The openly accessible MediaWiki history dataset dumps are available at [https://dumps.wikimedia.org/other/mediawiki\\_history/readme.html](https://dumps.wikimedia.org/other/mediawiki_history/readme.html). We further provide preprocessed data and results relevant to the manuscript in the code repository at <https://github.com/ruptho/wiki-volunteers-covid>. Any other supplementary data is available upon request from the corresponding author.

## Code availability

The code repository for this paper can be found at <https://github.com/ruptho/wiki-volunteers-covid>.

## Author contributions

T.R. retrieved the dataset, processed the data, and performed the experiments. T.R. and M.H.R. wrote the code. T.R., M.H.R., and T.S. analyzed the data. T.R., M.H.R., T.S., F.L., M.S., and D.H. conceived and designed the experiments, developed the arguments, and wrote the paper.

## Competing interests

The authors declare no competing interests.

## References

- [1] AL TAMIME, R., GIORDANO, R., AND HALL, W. Observing burstiness in Wikipedia articles during new disease outbreaks. In *Proceedings of the 10th ACM Conference on Web Science* (2018), pp. 117–126.
- [2] ALEXA. The top 500 sites on the Web. <https://www.alexa.com/topsites>, 2021. Accessed 2021-01-07.
- [3] ATCHISON, C. J., BOWMAN, L., VRINTEN, C., REDD, R., PRISTERA, P., EATON, J. W., AND WARD, H. Perceptions and behavioural responses of the general public during the COVID-19 pandemic: A cross-sectional survey of UK adults. *medRxiv* (2020).
- [4] BLOOM, N., BUNN, P., MIZEN, P., SMETANKA, P., AND THWAITES, G. The impact of COVID-19 on productivity. Working Paper 28233, National Bureau of Economic Research, December 2020.
- [5] BUTLER, B., JOYCE, E., AND PIKE, J. 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* (2008), pp. 1101–1110.
- [6] CHANG, S., PIERSON, E., KOH, P. W., GERARDIN, J., REDBIRD, B., GRUSKY, D., AND LESKOVEC, J. Mobility network models of COVID-19 explain inequities and inform reopening. *Nature* (2020), 1–6.
- [7] CHETTY, R., FRIEDMAN, J. N., HENDREN, N., STEPNER, M., AND TEAM, T. O. I. The economic impacts of COVID-19: Evidence from a new public database built using private sector data. Working Paper 27431, National Bureau of Economic Research, June 2020.
- [8] CIAMPAGLIA, G. L., AND TARABORELLI, D. MoodBar: Increasing new user retention in Wikipedia through lightweight socialization. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (2015), pp. 734–742.
- [9] COLAVIZZA, G. COVID-19 research in Wikipedia. *bioRxiv* (2020).
- [10] COLLIER, B., AND BEAR, J. Conflict, criticism, or confidence: An empirical examination of the gender gap in Wikipedia contributions. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (New York, NY, USA, 2012), CSCW '12, Association for Computing Machinery, p. 383–392.
- [11] DESVARS-LARRIVE, A., DERIC, E., HAUG, N., NIEDERKROTENTHALER, T., CHEN, J., DI NATALE, A., LASSER, J., GLIGA, D. S., ROUX, A., CHAKRABORTY, A., ET AL. A structured open dataset of government interventions in response to COVID-19. *medRxiv* (2020).
- [12] FELDMANN, A., GASSER, O., LICHTBLAU, F., PUJOL, E., POESE, I., DIETZEL, C., WAGNER, D., WICHTL-HUBER, M., TAPIADOR, J., VALLINA-RODRIGUEZ, N., ET AL. The lockdown effect: Implications of the COVID-19 pandemic on internet traffic. In *Proceedings of the ACM Internet Measurement Conference* (2020), pp. 1–18.
- [13] FLAXMAN, S., MISHRA, S., GANDY, A., UNWIN, H. J. T., MELLAN, T. A., COUPLAND, H., W-TAKER, C., ZHU, H., BERAH, T., EATON, J. W., ET AL. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* 584, 7820 (2020), 257–261.
- [14] FOLKE, C., CARPENTER, S. R., WALKER, B., SCHEFFER, M., CHAPIN, T., AND ROCKSTRÖM, J. Resilience thinking: integrating resilience, adaptability and transformability. *Ecology and society* 15, 4 (2010).
- [15] GALLOTTI, R., VALLE, F., CASTALDO, N., SACCO, P., AND DE DOMENICO, M. Assessing the risks of infodemics in response to COVID-19 epidemics. *arXiv preprint arXiv:2004.03997* (2020).
- [16] GARCIA, D., AND RIMÉ, B. Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological science* 30, 4 (2019), 617–628.

- [17] GELMAN, A., AND HUANG, Z. Estimating incumbency advantage and its variation, as an example of a before–after study. *Journal of the American Statistical Association* 103, 482 (2008), 437–446.
- [18] GILES, J. Internet encyclopaedias go head to head, 2005. Nature Publishing Group.
- [19] GOZZI, N., TIZZANI, M., STARNINI, M., CIULLA, F., PAOLOTTI, D., PANISSON, A., AND PERRA, N. Collective response to the media coverage of COVID-19 pandemic on Reddit and Wikipedia. *arXiv preprint arXiv:2006.06446* (2020).
- [20] GRUBER, J. The incidence of mandated maternity benefits. *The American economic review* (1994), 622–641.
- [21] HALFAKER, A., GEIGER, R. S., MORGAN, J. T., AND RIEDL, J. The rise and decline of an open collaboration system: How Wikipedia’s reaction to popularity is causing its decline. *American Behavioral Scientist* 57, 5 (2013), 664–688.
- [22] HALFAKER, A., KITTUR, A., AND RIEDL, J. 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* (2011), pp. 163–172.
- [23] HILL, B. M., AND SHAW, A. The Wikipedia gender gap revisited: Characterizing survey response bias with propensity score estimation. *PLoS one* 8, 6 (2013), e65782.
- [24] HORTA RIBEIRO, M., GLIGORIĆ, K., PEYRARD, M., LEMMERICH, F., STROHMAIER, M., AND WEST, R. Sudden attention shifts on Wikipedia during the COVID-19 crisis. *arXiv preprint arXiv:2005.08505* (2020).
- [25] JAY, J., BOR, J., NSOESIE, E. O., LIPSON, S. K., JONES, D. K., GALEA, S., AND RAIFMAN, J. Neighbourhood income and physical distancing during the COVID-19 pandemic in the United States. *Nature human behaviour* (2020), 1–9.
- [26] KEEGAN, B., AND FIESLER, C. The evolution and consequences of peer producing Wikipedia’s rules. *Proceedings of the International AAAI Conference on Web and Social Media* 11, 1 (May 2017).
- [27] KEEGAN, B. C., AND TAN, C. A quantitative portrait of Wikipedia’s high-tempo collaborations during the 2020 coronavirus pandemic. *arXiv preprint arXiv:2006.08899* (2020).
- [28] KITTUR, A., SUH, B., PENDLETON, B. A., AND CHI, E. H. He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2007), pp. 453–462.
- [29] LEMMERICH, F., SÁEZ-TRUMPER, D., WEST, R., AND ZIA, L. Why the world reads Wikipedia: Beyond english speakers. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (2019), pp. 618–626.
- [30] LEYS, C., LEY, C., KLEIN, O., BERNARD, P., AND LICATA, L. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology* 49, 4 (2013), 764 – 766.
- [31] MIAN, A., AND SUFI, A. House prices, home equity-based borrowing, and the US household leverage crisis. *American Economic Review* 101, 5 (2011), 2132–56.
- [32] MIQUEL-RIBÉ, M., AND LANIADO, D. Cultural identities in Wikipedias. In *Proceedings of the 7th 2016 International Conference on Social Media & Society* (2016), pp. 1–10.
- [33] RANSBOTHAM, S., AND KANE, G. Membership turnover and collaboration success in online communities: Explaining rises and falls from grace in Wikipedia. *MIS Quarterly* 35 (09 2011), 613–628.
- [34] ROBERT JR, L. P., AND ROMERO, D. M. The influence of diversity and experience on the effects of crowd size. *Journal of the Association for Information Science and Technology* 68, 2 (2017), 321–332.
- [35] SCHEFFER, M., WESTLEY, F., AND BROCK, W. Slow response of societies to new problems: causes and costs. *Ecosystems* 6, 5 (2003), 493–502.
- [36] SCHEFFER, M., AND WESTLEY, F. R. The evolutionary basis of rigidity: locks in cells, minds, and society. *Ecology and Society* 12, 2 (2007).

- [37] SHAW, A., AND HARGITTAI, E. The pipeline of online participation inequalities: The case of Wikipedia editing. *Journal of communication* 68, 1 (2018), 143–168.
- [38] SUH, B., CONVERTINO, G., CHI, E. H., AND PIROLI, P. The singularity is not near: Slowing growth of Wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration* (New York, NY, USA, 2009), WikiSym '09, Association for Computing Machinery.
- [39] SUH, J., HORVITZ, E., WHITE, R. W., AND ALTHOFF, T. Population-scale study of human needs during the COVID-19 pandemic: Analysis and implications. *arXiv preprint arXiv:2008.07045* (2020).
- [40] SULTANA, A., TASNIM, S., BHATTACHARYA, S., HOSSAIN, M. M., AND PUROHIT, N. Digital screen time during COVID-19 pandemic: A public health concern. [osf.io/preprints/socarxiv/e8sg7](https://osf.io/preprints/socarxiv/e8sg7), Sep 2020.
- [41] SUMI, R., YASSERI, T., ET AL. Edit wars in wikipedia. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing* (2011), IEEE, pp. 724–727.
- [42] TRUONG, C., OUDRE, L., AND VAYATIS, N. Selective review of offline change point detection methods. *Signal Processing* 167 (2020), 107299.
- [43] TSVETKOVA, M., GARCÍA-GAVILANES, R., FLORIDI, L., AND YASSERI, T. Even good bots fight: The case of wikipedia. *PloS one* 12, 2 (2017), e0171774.
- [44] VRANDEČIĆ, D., AND KRÖTZSCH, M. Wikidata: a free collaborative knowledgebase. *Communications of the ACM* 57, 10 (2014), 78–85.
- [45] WIKIMEDIA FOUNDATION. Responding to COVID-19: How we can help in this time of uncertainty, Oct 2020.
- [46] WORLD HEALTH ORGANIZATION. The world health organization and wikimedia foundation expand access to trusted information about COVID-19 on Wikipedia, Oct 2020.
- [47] XIONG, W., AND LAGERSTRÖM, R. Threat modeling—a systematic literature review. *Computers & security* 84 (2019), 53–69.
- [48] XU, B., AND LI, D. An empirical study of the motivations for content contribution and community participation in Wikipedia. *Information & management* 52, 3 (2015), 275–286.
- [49] YASSERI, T., SUMI, R., AND KERTÉSZ, J. Circadian patterns of wikipedia editorial activity: A demographic analysis. *PloS one* 7, 1 (2012), e30091.
- [50] YASSERI, T., SUMI, R., RUNG, A., KORNAI, A., AND KERTÉSZ, J. Dynamics of conflicts in Wikipedia. *PloS one* 7, 6 (2012), e38869.
- [51] ZHANG, A. F., LIVNEH, D., BUDAK, C., ROBERT, L., AND ROMERO, D. Shocking the crowd: The effect of censorship shocks on chinese Wikipedia. *Proceedings of the International AAAI Conference on Web and Social Media* 11, 1 (May 2017).
- [52] ZHANG, A. F., WANG, R., BLOHM, E., BUDAK, C., ROBERT JR, L. P., AND ROMERO, D. M. Participation of new editors after times of shock on Wikipedia. In *Proceedings of the International AAAI Conference on Web and Social Media* (2019), vol. 13, pp. 560–571.
- [53] ZHENG, L., ALBANO, C. M., VORA, N. M., MAI, F., AND NICKERSON, J. V. The roles bots play in Wikipedia. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–20.
- [54] ZHENG, Q., JONES, F. K., LEAVITT, S. V., UNG, L., LABRIQUE, A. B., PETERS, D. H., LEE, E. C., AND AZMAN, A. S. HIT-COVID, a global database tracking public health interventions to COVID-19. *Scientific data* 7, 1 (2020), 1–8.

# Supplementary Information

## Wikipedia Statistics

This supplementary section summarizes Wikipedia statistics relevant to this study. First, we visualize the total, monthly, weekly, as well as daily edits to Wikipedia in the year before the COVID-19 pandemic began (2019) in Supplementary Table 1. Secondly, Supplementary Table 2 shows statistics describing the yearly growth of Wikipedia between 2015 and 2020 until October of each year. Finally, Supplementary Table 3 lists all Wikipedias we studied, alongside their automatically detected mobility changepoints (Methods).

Table 1: **Edits to Wikipedia in 2019.** Overall, Wikipedia language editions were subject to 193 412 million visits and 319.82 million content edits by non-bot users in 2019.<sup>1,2</sup> This translates to about 530 million visits and 876 thousand edits per day.

<sup>1</sup>Total page views, via Wikimedia Statistics (<https://w.wiki/vkf>)

<sup>2</sup>User edits, via Wikimedia Statistics (<https://w.wiki/vkm>)

In 2019	Views (in Millions)	Edits (in Millions)
<b>Total</b>	193 412.78	319.82
<b>Monthly</b>	16 117.73	26.65
<b>Weekly</b>	3 719.48	6.15
<b>Daily</b>	529.9	.876

Table 2: **Growth between 2019 and 2020 in the English Wikipedia nearly doubles the relative increase between 2015 and 2019.** Edit growth in Wikipedia stagnated in recent years. From 2019 to 2020, edits grew by 8.41% (2.3 million edits), that is nearly double the growth between 2015 and 2019 (4.48%, or about 1.2 million edits).<sup>1</sup> In 2018 and 2019, human edits even declined in comparison to previous years. Therefore, the developments in 2020 mark a clear difference to the downward trend of edits in Wikipedia in the last 5 years.

<sup>1</sup>Edits by anonymous or registered users, via Wikimedia Statistics (<https://w.wiki/u8R>)

Year	Non-Bot Edits (until September 31 <sup>st</sup> )	Difference to Previous Year		Difference to 2019	
		Edits	Percent	Edits	Percent
2015	25 943 729	–	–	–1 163 251	–4.48
2016	27 095 495	1 151 766	4.44	–11 485	–0.04
2017	27 406 797	311 302	1.15	299 817	1.09
2018	27 374 765	–32 032	–0.12	267 785	0.98
2019	27 106 980	–267 785	–0.98	–	–
2020	29 386 750	2 279 770	8.41	2 279 770	7.76

## Number of editors by daily activity level during COVID-19 mobility restrictions

On Wikipedia, human users contribute edits with varying daily intensity. Wikipedia categorizes editors into five groups, according to their daily activity: 1 to 4, 5 to 24, 25 to 99, and more than 99 daily edits. We retrieve the number of registered editors (and their activity level) via the Wikimedia REST API and apply DiD analysis to detect significant changes across the editor population (Methods).

We again remove outliers before performing DiD analysis (see Methods) and visualize the results for all daily activity levels in Supplementary Figures 1 (1 to 4 edits), 2 (5 to 24 edits), 3 (25 to 99 edits), and 4 (more than 99 edits). The results corroborate our previous newcomer and edit volume findings, as the number of editors increases significantly after mobility changepoints. Our findings signal an increase in contribution across all activity levels for the editor population, particularly in large and medium Wikipedias, while results for small Wikipedias remain consistent with pre-pandemic baselines.

## Revert rate during COVID-19 mobility restrictions

The supplementary explanations in this section extend the revert rate analysis carried out in Discussion. We plot the rolling 7-day average revert rate in Supplementary Figure 5a as well as logarithmic effects for  $\delta_{rr}$  captured by DiD analysis with revert rate as the dependent variable in Supplementary Figure 5b for large (top), medium (middle), and small (bottom) Wikipedias. We find recession of revert rates for most Wikipedias during the initial weeks

Table 3: **Wikipedia language versions.** The 12 Wikipedia language editions relevant to this study, ordered by the total number of edits (bot or non-bot) in 2019, including mobility and normality changepoint dates (Methods).

Language	Changepoints (2020)		Wikipedia Version	
	Mobility	Normality	Code	Edits in 2019 (Millions)
English	03/16	05/21	en	40.56
French	03/16	07/02	fr	7.45
German	03/16	07/10	de	7.33
Italian	03/11	06/26	it	5.80
Japanese	03/31	06/14	ja	3.84
Swedish	03/11	06/05	sv	2.73
Dutch	03/16	05/29	nl	1.78
Korean	02/25	04/15	ko	1.61
Serbian	03/16	05/02	sr	1.29
Norwegian	03/11	06/04	no	0.71
Finnish	03/16	05/21	fi	0.65
Danish	03/11	06/05	da	0.31

of mobility restrictions, possibly indicating a reduction in negative contributions that need to be reverted (e.g., vandalism). Coefficient values for  $\delta_{rr}$  support this sentiment for large and particular medium or small Wikipedias. Nevertheless, it must be mentioned that revert rate can not be interpreted so simply, as specific bots periodically refactor revisions (e.g., monthly, quarterly) or some editor groups conduct article maintenance in coordinated events. Such difficult-to-predict patterns might be especially notable in Wikipedias with a generally lower amount of reverts, which is often the case in smaller Wikipedias. However, even these spontaneous patterns that would normally drive up revert rates appear to be mostly muted during the COVID-crisis.

Revert rates in the large English, German, and French Wikipedias drop after mobility restrictions come into effect in March 2020 (Supplementary Fig. 5a, top). For English, we observe an average revert rate of 0.09 in the month before mobility restrictions take effect and 0.08 in the month after. Revert rate for both German and French averages approximately 0.06 before the changepoints, but reaches respective minima of 0.043 and 0.045 in the subsequent weeks. DiD analysis and corresponding  $\delta_{rr}$  (Supplementary Fig. 5b, top) confirm significant relative decreases by measuring respective logarithmic effects of  $-0.327$  and  $-0.329$  for the French and German Wikipedia, signaling a 28%-decline for both Wikipedias ( $e^{-0.327} \approx e^{-0.329} \approx 72\%$  of previous levels). Italian shows a similar drop in revert rate ( $-23\%$ ). Relative decrease for English is considerably lower ( $-11\%$ ), but is deemed significant by our DiD analysis. Altogether, we find significant decreases in revert rate for all large Wikipedias.

Most medium and small Wikipedias seem to not exhibit considerable negative effects for revert rates (Supplementary Fig. 5a, middle and bottom). However, instead of explicitly showing visible dips in the revert rate graphs, the general level seems to be subdued during the investigate periods in 2020, especially close to the mobility changepoints. As an exception, Korean is the only language version that increases its revert rate during times of mobility restrictions, from 0.06 before the changepoint to a maximum of 0.075 in the month thereafter. DiD analysis reveals significant relative decreases in  $\delta_{rr}$  (Supplementary Fig. 5b, middle and bottom) for the medium Japanese ( $-42\%$ ), Dutch ( $-28\%$ ), and Swedish ( $-28\%$ ) Wikipedias within two months post changepoint, as well as for the smaller Norwegian ( $-63\%$ ), Serbian ( $-56\%$ ), and Danish ( $-48\%$ ) Wikipedias. We explain some of these significant effects by the generally higher revert rate in the same period in previous years. Although our DiD analysis uncovers these significant short-term declines in revert rates for medium and small Wikipedias, results must be taken with a grain of salt due to the aforementioned nature of reverts in smaller Wikipedias.

## Edit volume in articles not related to COVID-19

We investigate the impact of articles strongly related to COVID-19 (Methods) on the edit volume on Wikipedia. Supplementary Table 4 lists information about the total percentage of edits to COVID-19 articles, as well as the percentage of edited articles that are related to COVID-19. Supplementary Figure 6 shows the percentage of edits going towards articles strongly related to COVID-19 articles, as well as the overall percentage of edited articles that were strongly related to COVID-19. It appears that edits to COVID-19 articles make up an insignificant account of daily activity in most Wikipedias (mostly  $< 4\%$ ), whereas some Wikipedias, for example German (at one point 15% daily COVID-19 edits), have a somewhat higher but short-lived affinity for COVID-19 topics.

To quantify the effect of COVID-19 articles on overall edit volume, we perform the same DiD analysis for






Figure 1: **Editors with 1 to 4 Edits during COVID-19 mobility restrictions.** **a**, We depict the 7-day rolling average of active registered editors with an activity level of 1 to 4 daily edits in the context of COVID-19 mobility restrictions, delineated via mobility and normality changepoints. **b**, We show the relative change in active registered editors ( $ae$ ) with 1 to 4 edits per day as retrieved from the DiD via  $\delta_{ae}$  (95% confidence intervals as two standard deviations). The number of editors with 1 to 4 daily edits increases significantly after mobility restrictions come into effect, for all but a few medium and small Wikipedias.

edit volume as in Results, this time specifically excluding edits to articles that are strongly related to COVID-19. Supplementary Figure 7 visualizes the performed DiD analysis for edit volume, excluding edits to COVID-19 articles. We find that excluding edits to COVID-19 articles does not significantly alter the results reported beforehand.

## Robustness checks for edit volume, newcomers, and revert rate

As robustness checks, we perform variations of our DiD experiments for edit volume, newcomers, and revert rate. Supplementary Figures 8, 9, and 10 visualize DiD with a 14-day post-changepoint period. Supplementary Figures 11, 12, and 13 depict DiD with 7-day post-changepoint periods, but move mobility changepoints to seven days before the actual dates. Similarly, Supplementary Figures 14, 15, and 16 move mobility changepoints to seven days after the actual changepoints. Our DiD robustness checks show that longer post-changepoint periods or modified changepoint dates do not significantly influence results and prove the robustness of our methodology.




Figure 2: **Editors with 5 to 24 Edits during COVID-19 mobility restrictions.** **a**, We depict the 7-day rolling average of active registered editors with an activity level of 5 to 24 daily edits in the context of COVID-19 mobility restrictions, delineated via mobility and normality changepoints. **b**, We show the relative change in active registered editors ( $ae$ ) with 5 to 24 edits per day as retrieved from the DiD via  $\delta_{ae}$  (95% confidence intervals as two standard deviations). In all large and most medium Wikipedias the number of editors with 5 to 24 daily edits significantly increases over longer periods of time, while smaller Wikipedias do not consistently increase their editor numbers for this activity level.

Table 4: **COVID-19 edits and edited COVID-19 articles.** We list the total percentage of non-bot edits to COVID-19 articles as well as the percentage of edited articles that were related to COVID-19 for the total time span between January 1<sup>st</sup> and September 31<sup>st</sup> 2020 (second and third column). Furthermore, we show the maximum daily percentage for these two metrics during this period (third and fourth column).

Code	% of Edits to COVID-19 Articles	% of Edited Articles that are COVID-19 Articles	Max. Daily % of Edits to COVID-19 Articles	Max. Daily % of Edited Articles that are COVID-19 Articles
de	2.40	0.44	17.0	1.28
fr	0.73	0.24	3.44	1.00
it	0.29	0.12	1.78	0.72
sr	0.02	0.02	1.27	0.65
no	0.12	0.07	2.19	1.10
ko	0.04	0.03	1.15	0.26
da	0.33	0.16	6.35	1.55
sv	0.25	0.07	3.99	0.49
ja	0.22	0.11	1.49	0.46
nl	0.64	0.21	4.23	1.49
fi	0.75	0.37	4.12	1.43
en	1.03	0.33	4.44	0.82




Figure 3: **Editors with 25 to 99 Edits during COVID-19 mobility restrictions.** **a**, We depict the 7-day rolling average of active registered editors with an activity level of 25 to 99 daily edits in the context of COVID-19 mobility restrictions, delineated via mobility and normality changepoints. **b**, We show the relative change in active registered editors ( $ae$ ) with 25 to 99 edits per day as retrieved from the DiD via  $\delta_{ae}$  (95% confidence intervals as two standard deviations). We observe significant increases for the number of registered editors who perform 25 to 99 daily edits during COVID-19 mobility restrictions in large Wikipeidias. Results for medium and small Wikipeidias are mostly inconsistent.




Figure 4: **Editors with more than 99 Edits during COVID-19 mobility restrictions.** **a**, We depict the 7-day rolling average of active registered editors with an activity level of 25 to 99 daily edits in the context of COVID-19 mobility restrictions, delineated via mobility and normality changepoints. **b**, We show the relative change in active registered editors ( $ae$ ) with 25 to 99 edits per day as retrieved from the DiD via  $\delta_{ae}$  (95% confidence intervals as two standard deviations). For all large Wikipedias, besides German, we find increased counts for editors with more than 99 edits per day. Additionally, some medium Wikipedias exhibit significantly more of these high-intensity editors, while smaller Wikipedias show no strong significant trends.




Figure 5: **Revert rate during COVID-19 mobility restrictions.** We depict results for revert rate in large (top), medium (middle), and small (bottom) Wikipedias during COVID-19 mobility restrictions, which in the figure we delineate using mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We plot rolling 7-day average revert rate in 2018, 2019, and 2020 up until October. Even though our analysis showed that newcomers and edit volume increased in 2020, especially after the mobility changepoint, we do not observe increases for revert rate in any Wikipedias. **b**, We calculate the relative change in revert rate ( $rr$ ) to pre-changepoint periods from the DiD via  $\delta_{rr}$  (95% confidence intervals as two standard deviations), and plot  $\delta_{rr}$  for 120 left-aligned seven-day-windows (see Methods), beginning with the respective mobility changepoint. We detect no significant increase after mobility restrictions come into effect in virtually all Wikipedias, with the exception of Korean, and even find decreased revert rates for most large and medium Wikipedias.




Figure 6: **COVID-19 edits and edited COVID-19 articles.** We visualize the percentage of COVID-19 article edits per day and daily edited COVID-19 articles for large (top), medium (middle), and small (bottom) Wikipedias. **a.** We show the daily percentage of non-bot edits to COVID-19 articles between January 1<sup>st</sup> and September 31<sup>st</sup> 2020. **b.** We visualize the daily percentage of edited articles that were related to COVID-19 in 2020 until October.




Figure 7: **Edit volume in articles not related to COVID-19 during mobility restrictions.** We show edit volume in non-COVID-19 articles in large (top), medium (middle), and small (bottom) Wikipedias during the mobility restrictions, which we delineate using mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We show rolling 7-day average edit volume generated by human editors for 2018, 2019, and 2020 until October. **b**, We depict relative change in edit volume ( $ev$ ) as retrieved from DiD via  $\delta_{ev}$  (95% confidence interval as two standard deviations) and plot  $\delta_{ev}$  for 120 left-aligned seven-day-windows. Edits to articles closely related to COVID-19 mostly only make up a small fraction of daily edits (Supplementary Table 4 and Supplementary Fig. 6). Accordingly, we observe that findings for edit volume excluding COVID-19 edits barely differ from those for overall edit volume depicted in Figure 2. Although a minimal visual effect is observable for a few select Wikipedias (e.g., German), it does not affect significance of our results.






Figure 8: **Edit volume during COVID-19 mobility restrictions (14-day windows).** We show results of our 14-day window robustness experiment for edit volume in large (top), medium (middle), and small (bottom) Wikipedias during COVID-19 mobility restrictions, delineated using mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We show rolling 14-day average daily edit volume generated by human editors for 2018, 2019, and 2020 until October. **b**, We depict relative change in edit volume ( $ev$ ) as retrieved from DiD via  $\delta_{ev}$  (95% confidence interval as two standard deviations) and plot  $\delta_{ev}$  for 120 left-aligned fourteen-day-windows. Edit volume results for 14-day windows represent the same trends and similar significant effects as previous experiments (Figure 2), only smoothing the 7-day-window results more.




Figure 9: **Newcomers during COVID-19 mobility restrictions (14-day windows).** We show results of our 14-day window robustness experiment for newcomers in large (top), medium (middle), and small (bottom) Wikipedias during COVID-19 mobility restrictions, delineated using mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We show rolling 14-day average newcomer counts for 2018, 2019, and 2020 until October. **b**, We depict relative change in newcomers ( $nc$ ) as retrieved from DiD via  $\delta_{nc}$  (95% confidence interval as two standard deviations) and plot  $\delta_{nc}$  for 120 left-aligned fourteen-day-windows. Newcomer results for 14-day windows represent the same trends and similar significant effects as previous experiments (Figure 3), only smoothening the 7-day-window results more.

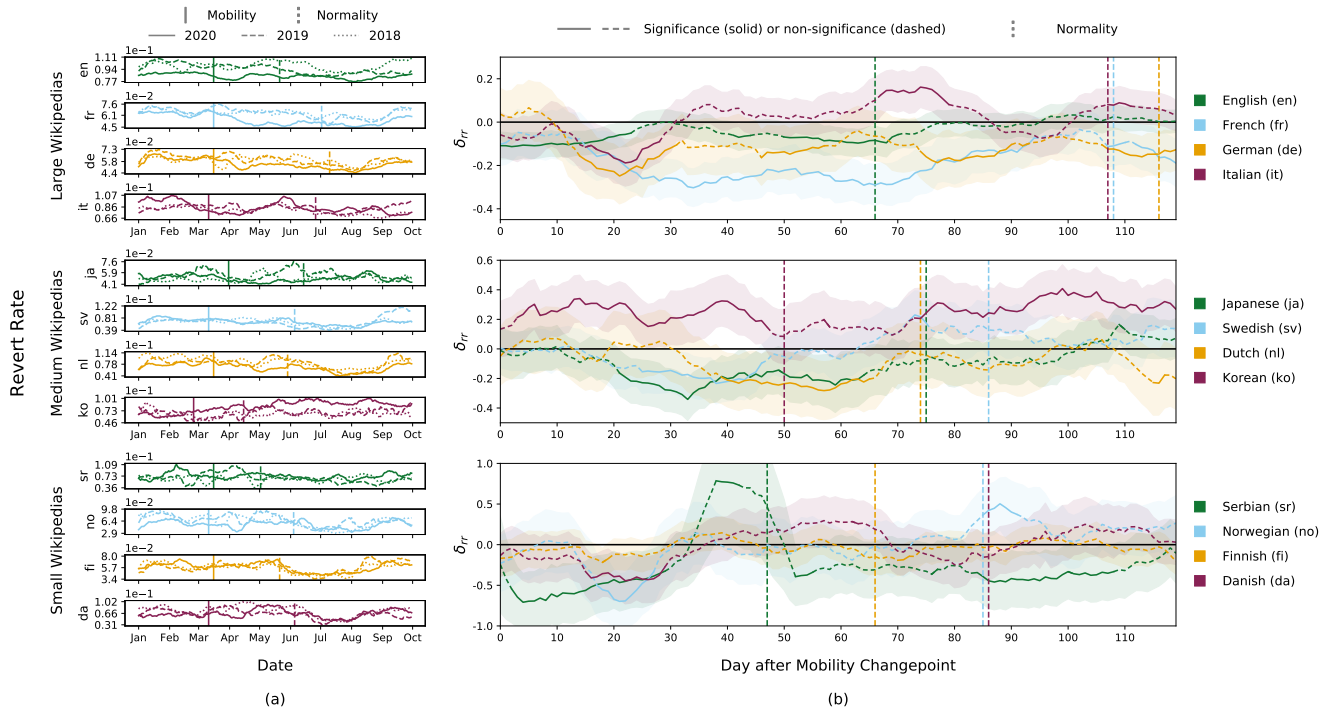


Figure 10: **Revert rate during COVID-19 mobility restrictions (14-day windows).** We show results of our 14-day window robustness experiment for revert rate in large (top), medium (middle), and small (bottom) Wikipedias during COVID-19 mobility restrictions, delineated using mobility (when restrictions become effective) and normality (when restrictions are lifted) changepoints. **a**, We show rolling 14-day average revert rate for 2018, 2019, and 2020 until October. **b**, We depict relative change in revert rate ( $rr$ ) as retrieved from DiD via  $\delta_{rr}$  (95% confidence interval as two standard deviations) and plot  $\delta_{rr}$  for 120 left-aligned fourteen-day-windows. Revert rate results for 14-day windows represent the same trends and similar significant effects as previous experiments (Supplementary Fig. 5), only smoothing the 7-day-window results more.




Figure 11: **Edit volume during COVID-19 mobility restrictions (beginning 7 days earlier)**. We vary our DiD for edit volume by changing the mobility (when restrictions take effect) and normality changepoint (when restrictions are lifted) to 7 days earlier. **a**, We show rolling 7-day average edit volume for 2018, 2019, and 2020 until October. **b**, We depict relative change in edit volume ( $ev$ ) as retrieved from DiD via  $\delta_{ev}$  (95% confidence interval as two standard deviations) and plot  $\delta_{ev}$  for 120 left-aligned seven-day-windows. As mobility changepoints generally mark dates of decreased activity, moving the changepoint before these declines lead to the first few days representing a more negative trend and values for later days are slightly lower than in the original experiment (Figure 2), which is an expected effect.




Figure 12: **Newcomers during COVID-19 mobility restrictions (beginning 7 days earlier)**. We vary our DiD for newcomers by changing the mobility (when restrictions take effect) and normality changepoint (when restrictions are lifted) to 7 days earlier. **a**, We show rolling 7-day average newcomers for 2018, 2019, and 2020 until October. **b**, We depict relative change in newcomers ( $nc$ ) as retrieved from DiD via  $\delta_{nc}$  (95% confidence interval as two standard deviations) and plot  $\delta_{nc}$  for 120 left-aligned seven-day-windows. As mobility changepoints generally mark dates of decreased activity, moving the changepoint before these declines lead to the first few days representing a more negative trend, before than recovering and increasing to slightly lower values than in the original experiment (Figure 3), which is an expected effect.




Figure 13: **Revert rate during COVID-19 mobility restrictions (beginning 7 days earlier).** We vary our DiD for revert rate by changing the mobility (when restrictions take effect) and normality changepoint (when restrictions are lifted) to 7 days earlier. **a**, We show rolling 7-day average revert rate for 2018, 2019, and 2020 until October. **b**, We depict relative change in revert rate ( $rr$ ) as retrieved from DiD via  $\delta_{rr}$  (95% confidence interval as two standard deviations) and plot  $\delta_{rr}$  for 120 left-aligned seven-day-windows. We find no strong differences to the original experiment (Supplementary Fig. 5), as revert rates are relatively stable close to the mobility changepoint.




Figure 14: **Edit volume during COVID-19 mobility restrictions (beginning 7 days later).** We vary our DiD for edit volume by changing the mobility (when restrictions take effect) and normality changepoint (when restrictions are lifted) to 7 days later. **a**, We show rolling 7-day average edit volume for 2018, 2019, and 2020 until October. **b**, We depict relative change in edit volume ( $ev$ ) as retrieved from DiD via  $\delta_{ev}$  (95% confidence interval as two standard deviations) and plot  $\delta_{ev}$  for 120 left-aligned seven-day-windows. As mobility changepoints generally mark dates of decreased activity, moving the changepoint past these declines leads to these lower values now being counted towards the 30-day baseline period, generating overall higher post-changepoint values than in the original experiment (Figure 2).



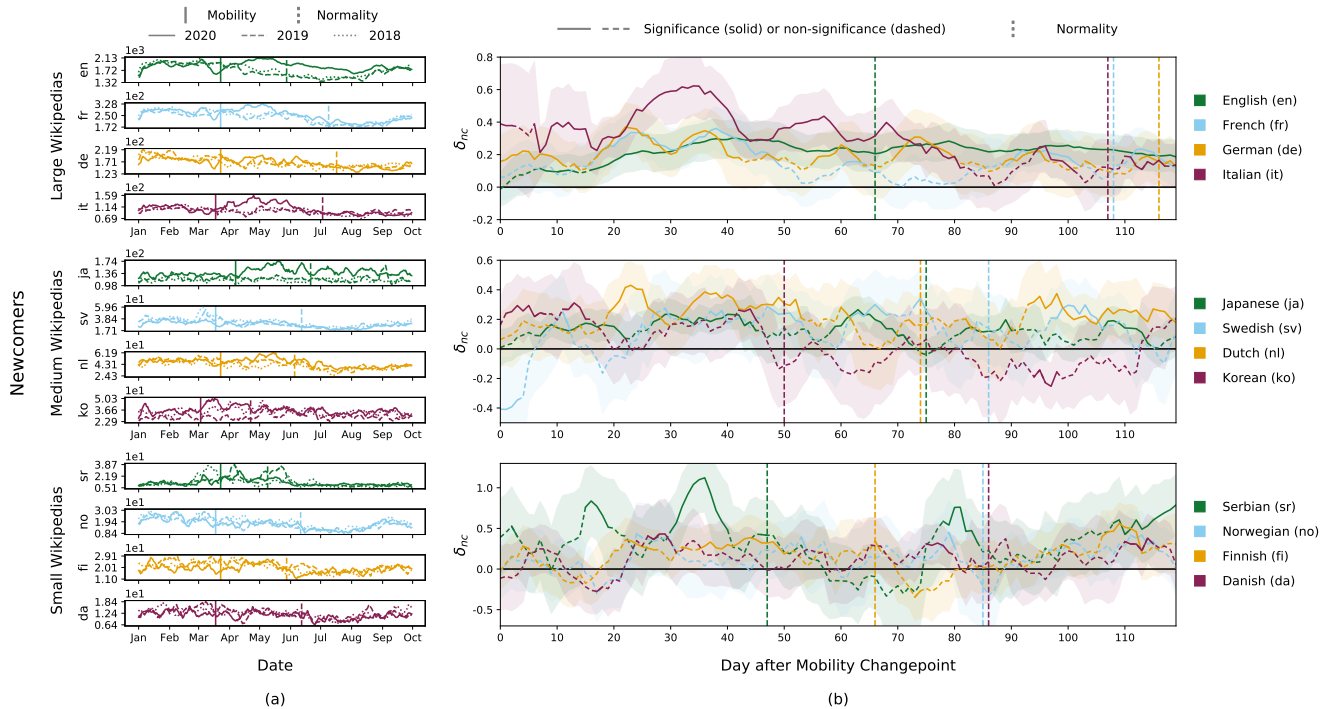


Figure 15: **Newcomers during COVID-19 mobility restrictions (beginning 7 days later).** We vary our DiD for newcomers by changing the mobility (when restrictions take effect) and normality changepoint (when restrictions are lifted) to 7 days later. **a**, We show rolling 7-day average newcomers for 2018, 2019, and 2020 until October. **b**, We depict relative change in newcomers ( $nc$ ) as retrieved from DiD via  $\delta_{nc}$  (95% confidence interval as two standard deviations) and plot  $\delta_{nc}$  for 120 left-aligned seven-day-windows. As mobility changepoints generally mark dates of decreased activity, moving the changepoint past these declines leads to these lower values now being counted towards the 30-day baseline period, generating overall higher post-changepoint values than in the original experiment (Figure 3).




Figure 16: **Revert rate during COVID-19 mobility restrictions (beginning 7 days later).** We vary our DiD for revert rate by moving the mobility (when restrictions take effect) and normality changepoint (when restrictions are lifted) to 7 days later. **a**, We show rolling 7-day average revert rate for 2018, 2019, and 2020 until October. **b**, We depict relative change in revert rate ( $rr$ ) as retrieved from DiD via  $\delta_{rr}$  (95% confidence interval as two standard deviations) and plot  $\delta_{rr}$  for 120 left-aligned seven-day-windows. We find no strong differences to the original experiment (Supplementary Fig. 5), as revert rates are relatively stable close to the mobility changepoint.