# Offline war in the online arena: Showing that civil war leads to conflict over national history using 250,000 Wikipedia texts

Bertel Teilfeldt Hansen*

October 30, 2018

**Abstract**

[ABSTRACT COMING HERE]

VERY EARLY AND UNFINISHED DRAFT - DO NOT CITE OR CIRCULATE]

*University of Copenhagen

## Introduction

John Stuart Mill believed that the most essential building-block of a nation was a sense of shared history among its people:

"(...) the strongest of all [causes of a feeling of nationality] is identity of political antecedents; the possession of a national history, and consequent community of recollections; collective pride and humiliation, pleasure and regret, connected with the same incidents in the past.

Mill (1910, p. 360).

With the benefit of 50 years of postmodernist thought we know that the construction of such a "community of recollections" is not merely a matter of collecting and stacking historical facts abound in the world. Rather, it is a continuous contest, a constant and often conflictual process of remembering and forgetting, ebbing and flowing over time (Wilmer, 2004). But what causes these ebbs and flows in the construction of a national narrative? What makes people fight over their history? Knowledge of this process is crucial to our understanding of the formation and destruction of nations and nation-states (Cinar, 2001). However, the topic has so far escaped systematic large-N inquiry, likely due to the difficulty of reliably measuring the fight over history at a large scale.

Using the first global dataset on the social construction of history, this paper shows that violent conflict, specifically civil war, is a major catalyst of the contestation over national history. When a nation is at war with itself, its national narrative is as much a battleground as its streets and cities.

The logic behind this is straightforward. When groups in society engage in violent conflict with each other they often have radically different conceptions of the identity of that society; radically different answers to questions like "who are we?" and "who were we?". Indeed they may not even agree that they are part of the same "we" at all. Sometimes such identity cleavages are the driving force behind the conflict (Wimmer, Cederman and Min, 2009), at other times they develop and sharpen over the course of the fighting (Fearon and Laitin, 2000). In most instances both dynamics are likely to occur.

Thus, when armies clash on the field, opposing representations of the nation and its history clash as well. Not just the combatants themselves, but also the sections of society that they fight for. A pertinent example can be found in the controversy over the mass killings in the infamous WWII *Jasenovac* concentration camp in Ustaša Crotia, on which Croatian and Serbian scholars and politicians – including later Croatian president, Franjo Tudman – published a number of wildly different accounts and casualty estimates in the beginning of the Yugoslav wars. The representation of the atrocities committed by the fascist Ustaša in Jasenovac and elsewhere were widely used in Serbian propaganda as proof of the "genocidal national character" of Croats (Kolstø, 2011). As one of Franke

Wilmer's interviewees from Belgrade put it: "history is a battlefield" (Wilmer, 2004, p. 135).

Other such single cases and anecdotes can be marshalled, but is this a general phenomenon? Can we generally expect to see national history become a second battlefield during civil war? To find out, this paper develops a global dataset on the social construction of history using 250,000 Wikipedia texts.

Wikipedia is a veritable treasure trove of data on the social construction of the "truth" of just about anything. It is without comparison the largest and most popular source of information that has ever existed. wikipedia.org is the fifth most visited site on the internet, and the english Wikipedia alone consists of over five million articles, which with revision history and markup amount to tens of terabytes. Crucially for this study, Wikipedia allows anyone to edit an article, and due to its unique revision history system every edit is recorded and available for download. Thus the contest over how a nation's history should be written can be observed in minute detail at a massive scale.

This paper uses MediaWiki's open API to download and parse full revision histories for the national history pages of every state in the Correlates of War universe. For each page, I use a revert-based method inspired by Yasseri et al. (2012) to measure conflict between Wiki users over the content of the article. These time series of conflict frequencies are aggregated in order to build a full country-year panel on to which data on civil war along with various controls are merged.

A standard two-way fixed effects model is then applied to this panel, yielding the main result of the paper: When a civil war is ongoing in a country, conflict over its history increases on average by between 66 and 84 %. Using a support vector machine and 10,000 hand-coded sentences to identify calendar years in the texts, I find that this war-induced increase in conflict is just as much about older history as it is about the present. This indicates that civil war does not just cause people to fight over its contemporary origins, but also about "ancient" history.

These findings touch on a number of literatures. The link between civil war and the identity of groups and nations has been the subject of intense quantitative study at least since the seminal works of Fearon and Laitin (2003) and Collier and Hoeffler (2004). Centered on scholarship from the EPR project (Vogt et al., 2015), something like a consensus has emerged maintaining that identity politics – specifically ethnic identity politics – is an important predictor of civil war (see e.g. Wimmer, Cederman and Min (2009); Cederman, Wimmer and Min (2010); Cederman, Weidmann and Gleditsch (2011); Wucherpfennig, Hunziker and Cederman (2016). Though essential to understanding how much any correlation between war and group identity is due to identity affecting war, the opposite direction of the causal arrow remains relatively unstudied in the quantitative conflict literature. There exists a body of qualitative evidence suggesting that group identity and conflict are mutually endogenous (Wood, 2008; Kalyvas, 2008) along with some single-country surveys (Dyrstad, 2012; Kubota, 2017), but war's impact on culture and identity is generally not well understood (Blattman and Miguel, 2010).

The paper also relates to the vast literature on the origins of states and nations. Here, war and violent struggle is seen not only as "making the state" (Tilly, 1990), but also as a key ingredient in the myths and narratives that make up the history of the nation (Renan, 1882; Smith, 1986). Indeed, according to Michael Howard, the most important incidents in group memory are usually of violent conflict, and thus "(...) no Nation, in the true sense of the word, could be born without war" (Howard, 1979, p. 102).

Finally, this paper's inherently positivist and quantitative data science approach to an essentially constructivist question contributes to a broader discussion in the social sciences about the recent spread of methods developed in statistics and computer science to areas traditionally dominated by qualitative research (Bennett, 2015; Barkin and Sjoberg, 2017; Grimmer and Stewart, 2013). Depending on one's requirement of the "big" in "big data", the data collected here can be seen as an example of how big data can help social scientists measure and investigate difficult-to-observe and difficult-to-quantify phenomena, such as the social construction of history (Monroe et al., 2015).

The remainder of the paper is structured as follows. First, I provide a brief introduction to the literature on the social construction of history and sketch out the argument for civil war as a catalyst of conflict over it. Second, I describe how the dataset used to test this argument is constructed from the MediaWiki API before going through a number of validation exercises to gauge how sensible the measure of conflict derived from it is. I then use a standard two-way fixed effects model along with a support vector machine to show that civil war substantially increases the amount of conflict over history and that the history being contested is not merely the contemporary origins of the war, but just as much the distant past. The final section concludes and offers implications for theory and policy.

## The social construction of history

[LITERATURE REVIEW COMING HERE]

## The argument: Jasenovac all over again

In 1989, less than two years before becoming the first president of Croatia, Franjo Tudman published the book, *Horrors of War: Historical Reality and Philosophy. Horrors of War* was written during Tudman's time in prison in the early eighties and painted a markedly different picture of the actions of the Ustaša than the official Titoist account. Most notably, it downgraded the number of victims in the Jasenovac concentration camp from somewhere around 700,000 to only 30,000-40,0000 and claimed that Serbs were not the primary targets of the killings.

While the concentration camp death estimates had been the subject of some debate after the war, this was mostly shut down by the communist party in the fifties. With Tito dead and nationalist sentiment simmering in the republics in the late eighties and early nineties, however, the book spawned a mass of more or less scholarly literature on Jasenovac and the other Ustaša camps, where ethnicity of the author came to be a near-perfect predictor of histographical position. Serbian writers claimed huge death tolls and bestial torture, and Croats emphasized the labour aspect of the camps and generally agreed with Tudman's low casualty estimates. During the Yugoslav Wars both sides frequently asserted that their counterpart's allegedly false accounts of history were deliberate attempts to justify their violent actions.[1]

The contested history of Jasenovac illustrates the dynamics at the heart of the argument presented here: Civil war and its precursors tend to make conflict participants fight over their shared history. This discursive struggle may in itself also contribute to the violence.

The proposed mechanism has both strategic and non-strategic elements.[2] The strategic part is based on the simple assumption that the elites of a group may be divided into extremists and moderates with respect to their views on the nature of some out-group. When these two types of elites vie for power the extremist will often stand to gain in in-group popularity from conflict with the out-group, since that will validate their own extremist views and discredit the moderates. Therefore they will have an incentive to portray the out-group as vicious and dangerous in order to spur or sustain conflict. As in the Jasenovac case, a major tool in such rhetoric is manipulation of the shared history of the two groups; particularly history of any violent conflict between them. Such arguments have a long pedigree in social science, both in constructivist and non-constructivist garb (see e.g. Brass (1997); Glaeser (2005); Simmel (2010)).

For the extremist elites to be successful in causing or continuing conflict, of course, their propaganda needs to affect the behaviour of some or all of the in-group non-elites, who are typically the ones to do the actual fighting. This can come about in two ways. Either the

[1] For a more thorough treatment of the Jasenovac debate, see Kolstø (2011).

[2] The following borrows heavily from the arguments and discussion of case evidence in Fearon and Laitin (2000).

non-elites are simply swayed by the rhetoric and duped into engaging in costly conflict with their out-group counterparts, or they act strategically and use the elite talk of communal violence and ancient enemies as a cover for pillaging and settling local grudges (Fearon and Laitin, 2000)[3]. In the context of the Yugoslav Wars, Woodward finds that many of the irregulars fighting were motivated not by ethnic hatred, but by the opportunity for personal enrichment in a period of severe economic decline (Woodward, 1995). Whether convinced by elite framing or exploiting it for their own ends, non-elites will have a strong incentive to further propagate this framing, since it justifies their violence – the more people that believe one's actions are justified, the less costly those actions tend to become. Thus non-elites are expected to participate in the (re)construction of history.

The previous two paragraphs have highlighted elite machinations as the driver of the contest over history, but this is complemented by a more organic process in which group identities are recast in a more antagonistic image through participating in violent conflict with other groups. When violent conflict – for whatever reason – occurs between two groups, their way of thinking about both themselves and each other is likely to change. Members of group A will become more inclined to view themselves as "those-who-fight-group-B" and to view members of group B as "those-who-fight-us" (and vice versa). This will make group members more prone to believe and disseminate historical representations of the other group as particularly murderous and of their actions as unjustified, which in turn can lead to more violence.

The proposed explanation for why civil war should breed conflict over history is thus a mix of different processes, some strategic, some not. The actual mix will probably vary from case to case, with some being driven mainly by scheming elites and others mainly by continuously hardening identities arising out of war itself. The central prediction is the same, however: Civil war leads to increased fighting over the shared history of its participants.

## Data

The connection between war and the social construction of history has escaped systematic inquiry not because the idea of such a connection is new, but because the data needed has been unavailable. Indeed the social construction of most things are tricky to measure reliably. Here, I exploit the online, collaborative encyclopaedia, Wikipedia, which I argue provides an excellent and unprecedented data source on the social construction of "facts".

[3]See Yanagizawa-Drott (2014) for an alternative explanation based on signals about future government punishments or rewards.

### Wikipedia and the social construction of facts

As noted above, Wikipedia is the largest information source ever created (not counting the internet itself), with the English Wiki alone consisting of about 5.7 million articles and an article history stretching into the tens of terabytes. According to Alexa.com, Wikipedia.org is currently the fifth most visited site on the net, beaten only by such giants as Google, Youtube, Facebook, and the Chinese search provider Baidu. The extensive use of Wiki data by applications like Siri and Amazon Alexa and indeed Google's own infoboxes is a testament to the pervasiveness of the site.

But Wikipedia is not just huge and hugely popular, it is also a huge social experiment. Its content is generated entirely by a massive number of unpaid users (called "editors") who collaboratively and incrementally put together each article. With very few exceptions, anyone can edit any article on Wikipedia.

While many sites rely on user-generated content and some on collaboration, Wikipedia is unique because the content its editors collaborate on generating is supposed to be *the truth*. Wikipedia is not a discussion forum or a soapbox, but an encyclopaedia. When users edit and argue over content on wiki, they are not arguing over tastes or opinions, but over the true representation of a topic. The final content of an article is thus the result of a discursive struggle over the truth. This means that what goes on on Wikipedia *is* the social construction of facts.

When people engage in a discursive struggle over facts in the offline world, their individual representations are usually not recorded systematically, if at all. On Wikipedia, however, every edit made is timestamped and stored in its entirety, such that any editor's version of the truth can be called forth and inspected. All these versions along with a host of metadata is freely available on MediaWiki's public API, and this allows one to observe the struggle over facts – of national histories or just about anything else – in minute detail and at a massive scale.

Despite all these advantages, Wikipedia also comes with some limitations as a data source on social construction. First, it is difficult to know who the people doing the constructing are. Available survey evidence indicates that editors are mostly male (Wikimedia, 2012), but representativeness of such surveys is hard to gauge. Futher, we generally do not know the geographical distribution of editors. When users edit without being logged in to an account, they are identified by their IP address, which enables one to geolocate them (with some caveats, such as the use of proxy servers). In the dataset I generate below, however, less than a third of editors can be geolocated in this way, and it is highly doubtful how representative of the general editor population they are, since editor dedication and seriousness is likely correlated with the tendency to use a named account.

Second, it is not straightforward to compare content across languages. While the general revert-based measure I employ can be used regardless of language, the later validation techniques and the supervised learning of calendar years cannot; at least not without very extensive, language-specific modifications. Perhaps more importantly, any sort of qualitative understanding of text samples from the more than 300 different languages that have Wikipedias would be highly impractical. Therefore, I focus on the English Wikipedia.

Relying exclusively on the English Wikipedia probably means missing some of the fight over history, since a non-negligible part may be occurring on national Wikis. Examples of this include Croatian fascist fighting intensely with moderate editors over both contemporary issues and WWII history on the Croatian Wiki (Sampson, 2013), and Russian state-run media manipulating the Russian Wiki article on the crash of Malaysian Airlines Flight 17 by inserting the official government narrative and claiming that the Ukrainians shot it down (Lowensohn, 2014).

However, the English Wikipedia is by far the largest Wiki (particularly when discounting purely bot-made content), has the overwhelmingly largest and most active user-base, and has between seven and eight billion views per month – more than seven times that of the closest runner-ups, the Spanish and the German Wikis (Wikimedia, 2018). Often, local topics will have better coverage and more views on the English Wiki than on the national ones, which makes it likely that purely national fights take place on the English Wiki as well, if not only there. Further, the content on the English Wikipedia is what the world sees, so that any actor who cares about the representation of their national history in the eyes of the world will want to push their version here. A closer reading of discussions on the forums association with each article (the "talk pages") also makes it clear that struggles over history on the English Wiki frequently occur between locals.[4]

### Building a dataset of history

The dataset of national histories created for this paper takes its outset in the widely used Correlates of War (COW) universe of states (Correlates of War Project, 2017). For each of the 195 states that existed in COW in 2016, I manually identify a corresponding history article on the English Wikipedia (hereafter just "Wikipedia"). While Wikipedia has no mandatory articles or topics that must be covered in their entirety, all states in the current COW universe happen to have an article dedicated specifically to their history, in addition to having a general article about them. Nearly all of these history articles were started in the first or second year after Wikipedia's founding in January 2001.

Iterating through this list of pages, I then query the MediaWiki API for all the different versions (called "revisions") of each history article there have been on Wikipedia and stack the resulting texts into a panel according to timestamp and country. As of 31

[4]See for instance the very intense debate over the proper name and origins of the Republic of Macedonia/Republic of North Macedonia/FYROM on the talk page of Macedonia's history article.

December 2017, this yields a total of 247,254 raw texts, which combines to about 10-13 gigabytes. These text data are supplemented with revision-level metadata such as editor names, IDs, comments, tags etc. along with derived variables like geolocations of editors and bibliographic references used in the text. For later validation, I also make a similar panel of the talk pages of each history article.[5]

The next step is to create a workable measure of conflict over history to correlate with real-world civil war. For this, I rely to a large extent on the work of Yasseri et al. (2012), who are interested in measuring so-called "edit wars" on Wikipedia. An edit war occurs when editors repeatedly override each other's contributions to a page by taking it in turns to revert the article text back to their respective preferred version. Edit warring is highly destructive for the collaborative writing effort and a rich system of guidelines and counter-measures has evolved to prevent it, the latter including locking articles for non-registered editors, tagging controversial articles, forming and running arbitration committees, and temporal or indefinite banning of continuously violating editors. Continuous mutual reverts are taken quite seriously on Wikipedia and is seen as a sign of real conflict between parties over the correct description of a topic or event.

Yasseri et al exploit this fact in their measure of Wiki conflict, which at its heart consists of a simple identification of reverts: Let $\ldots, i-1, i, i+1, \ldots, j-1, j, j+1, \ldots$ be successive revisions of an article, and let there then be a revert between the editor of $j$ and $i$ if revision $j$ is identical to revision $i-1$. I borrow this scheme for the basis of my measure, and since it necessitates that each revision text has to be compared to all the other revisions of an article, and since each article can have several thousand revisions each with thousands of words, I also follow Yasseri et al in first calculating the MD5 hash of the texts and then using these hashes in the comparison instead of the full strings (Rivest, 1992).[6]

Now, one could stop there and have some kind of measure of conflict. However, not all reverts are about the kind of substantive and sincere disagreements that the theoretical argument connects with civil war. A non-negligible portion of the reverts made on Wikipedia are against vandals who aim not to push some actual agenda, but merely to troll and disrupt serious users. When a user vandalizes an article, bots or human editors eventually (generally very quickly) revert it back to its last clean version and usually tag the vandal's revision with one of several tags indicating vandalism. To prevent the conflict measure from picking up such unsubstantive disagreements, I exclude all reverts made to revisions that are tagged with one or more vandalism-related tags. I also exclude all reverts made to revisions that add profanity to an article's text.[7]

[5] A more technical and in-depth explanation of the process of acquiring and preparing the data is available in the appendix along with the Python scripts used and the full list of URLs for the history articles.

[6] See e.g. Preneel (2010) for an overview of hashing.

[7] A full list of these tags and profanities is located in the appendix.

Bots play an important role on Wikipedia and carry out a wide range of tasks,[8] including instant reverting of obvious vandalism. The anti-vandalism bots contribute a large amount of reverts (particularly the different iterations of ClueBot), but again without constituting the type of contesting of history that a valid conflict measure should capture. Therefore, I query the MediaWiki API for a list of all user profiles registered as bots and exclude reverts made from any of these. Human editors occasionally make reverts against bot-made revisions, and in these instances I include the revert, but count it as being against the first human revision between the bot revision and the revert in question (if one exists). As bots generally just revert or make tiny changes, they are difficult to disagree with about historical content, and so humans making reverts against them are more likely to be in disagreement with the following man-made, and presumable substantive, revision. Finally, I exclude all reverts where the reverter and the revertee is the same editor (you cannot be in an edit war with yourself).

For a revert to be part of an actual conflict between editors, there has to be reciprocity. One side attacking the other's representation of history without any rebuttals does not constitute a contest, but merely a one-sided correction. Thus, in the main version of the measure, I further exclude reverts that are never reciprocated, such that for a revert to be counted, the revertee has to either already have made a revert against the reverter or do so before the time series ends. This criteria removes a large number of the otherwise ubiquitous reverts against anonymous editors without a logged-in profile, who do the majority of the trolling, but are often too subtle to be caught by the procedure relying on tags and profanities alone.

With these rules in place, we arrive at a variable that is able to reasonably distinguish revisions with substantive conflict from regular revisions, vandalism, and bot activity. According to this measure, for example, about 12 % of the revisions to the history article for Croatia spark conflict, whereas less then 1 % of the revisions to the history article for Denmark do so. If one includes all reverts, both numbers increase markedly – to 22 and 13 % for Croatia and Denmark, respectively – but, more importantly, the relative difference between them decrease, implying that general reverts are not as useful in detecting real conflict.

In order to carry out more systematic validation checks and correlate the measure with civil war, I aggregate each country's time series to the yearly level by summing the number of reverts that meet the above criteria. In an alternative version of the measure I instead sum the number of editor pairs that make such reverts against each other. I prefer the former version, however, since it implies a natural weighting of serious conflicts between editor pairs rather than treating all conflicts the same. Results using both measures are similar, if slightly stronger with the alternative measure (see additional tables in the appendix).

[8] Currently, 2,192 different bot tasks have been approved (Wikipedia, 2018).

Below I present descriptive statistics for the 3,205 country years in the 2001-2017 period for which a given country has both existed and has had a Wiki page about its history.

Table 1: Summary statistics

| Statistic | N | Mean | St. Dev. | Min | Max |
|---|---|---|---|---|---|
| Number of revisions | 3,205 | 76.8 | 128.2 | 0 | 1,706 |
| Median article size (bytes) | 3,205 | 38,303.6 | 42,268.9 | 136.0 | 291,484.0 |
| Median added text size (bytes) | 3,205 | 37.9 | 285.1 | 0.0 | 6,647.5 |
| Median removed text size (bytes) | 3,205 | 8.2 | 44.6 | 0.0 | 2,232.5 |
| Reverts | 3,205 | 11.6 | 25.8 | 0 | 394 |
| Reverts (no bots, no vandalism) | 3,205 | 9.9 | 23.5 | 0 | 378 |
| Reverts (mutual) | 3,205 | 1.1 | 4.8 | 0 | 126 |
| Reverts (mutual, no bots, no vandalism) | 3,205 | 1.0 | 4.7 | 0 | 126 |
| Reverts (mutual, editor pairs) | 3,205 | 0.3 | 1.2 | 0 | 22 |
| Reverts (mutual, editor pairs, no bots, no vandalism) | 3,205 | 0.3 | 1.2 | 0 | 22 |
| Revisions with profanities* | 3,205 | 0.5 | 1.9 | 0 | 36 |

*Full list of profanities available upen request

The average country year has 77 revisions made to its national history page, but the level of activity varies immensely; from no yearly revisions to 1,700. The variables counting bytes of text added and removed imply that these revision have tended to grow the pages in size quite substantially. For the preferred version of the revert measure ("mutual, no bots, no vandalism"), about 1 % of all revisions lead to conflict. Again, this exhibits quite a lot of variation, with the most conflictual country year (Croatia in 2006) having 126 reverts and many country years having none at all. Approximately 1 in 200 revisions includes profanities.

## Validation

As is always the case when developing a measure for a phenomenon that has not previously been measured, it is hard to know how a valid measure would be distributed. As a first cut, though, we can gain some insight by examining cross-country patterns and seeing if they line up reasonably well with expectations based on general information about the world.

To get a valid comparison between countries one must first take out some of the variation in reverts that is due purely to structural factors. More populous countries can, on average, be expected to have more conflict then less populous ones, since they will have more editors writing on the material. The same goes for richer countries, who will have a larger share of people with access to computers and the necessary education and outlook to want to edit national history pages. Further, since I deal only with the English Wikipedia, English-speaking countries will probably also tend to have more editors. Finally, the longer a country has existed, the more history there is to write about (arguably), and so the more revisions we can expect.

Below I present a table of the twenty most conflictual national histories according to the preferred revert measure. The ranking is generated by first summing all the reverts by country and then adding median population and median GDP per capita from the World Development Indicators (World Bank, 2017), a dummy for English as the most widespread language from GeoNames (GeoNames, 2017), and a count for years existed calculated from COW. These variables are used to predict the total number of reverts using standard OLS, and the ranking is subsequently computed from the residuals of this model. These residuals can be interpreted as excess (or deficit) conflict considering the structural conditions of a country.

Table 2: Top 20 countries with most excess reverts
(given size, income, language, and years as an independent country)

| Rank | Country | Residual | Total mutual reverts (no bots, no vandalism) | Total reverts |
|---|---|---|---|---|
| 1 | Croatia | 185.12 | 206 | 370 |
| 2 | Bosnia and Herzegovina | 114.25 | 129 | 238 |
| 3 | Japan | 111.05 | 157 | 1634 |
| 4 | Pakistan | 99.47 | 129 | 521 |
| 5 | Israel | 90.42 | 118 | 741 |
| 6 | India | 78.48 | 118 | 1203 |
| 7 | Moldova | 65.50 | 75 | 122 |
| 8 | Italy | 65.50 | 106 | 657 |
| 9 | China | 61.08 | 109 | 1436 |
| 10 | Azerbaijan | 54.47 | 75 | 240 |
| 11 | Hungary | 51.52 | 77 | 237 |
| 12 | United States of America | 49.72 | 99 | 2278 |
| 13 | Russia | 43.24 | 83 | 746 |
| 14 | Afghanistan | 39.05 | 55 | 359 |
| 15 | Albania | 35.49 | 48 | 184 |
| 16 | Iran | 31.75 | 65 | 417 |
| 17 | Greece | 27.61 | 57 | 806 |
| 18 | Armenia | 25.17 | 37 | 208 |
| 19 | Kosovo | 25.11 | 37 | 109 |
| 20 | Bulgaria | 22.42 | 43 | 151 |

The ordering is fairly consistent with what one would expect given differences in historical controversies between countries, and plausible, if post hoc, explanations can be offered for most positions. Not entirely surprisingly, the history of Croatia is by far the most conflictual, followed by the history of Bosnia and Herzegovina. Japan is third, likely due to controversy over the empire's war crimes during WWII, which have also spawned serious offline controversy and government revisionism in history textbooks (Guex, 2015). In fourth place, Pakistan has had its own textbook controversy with alleged Indophobia and Islamification of the country's roots, and one scholar has placed its history textbooks "(...) among the best available sources for assessing the nexus between power and bigotry in creative imaginings of a national past." (Jalal, 1995, p. 78). Further, in the light of the country's highly contested origins and borders, Israel being in the top five is well in line with expectations.

Though reasonable stories can be told for the remaining countries on the top 20 as well, there are some unexpected omissions that bear mentioning. Considering the extreme controversy surrounding Croatia's history page, for example, it is surprising that Serbia is not also up there, just as the high rank of Armenia raises questions about Turkey's placement.

This is mainly due to idiosyncrasies in how some Wiki pages have been created and how their discussions have developed. Turkey's history page was split in two in 2010, and much of the controversy happens on the sister page dedicated exclusively to the modern Republic of Turkey founded in 1922 (despite the fact that the Armenian Genocide occurred in 1915). Similarly, much of the contest over Serbian history occurs on the general article for Yugoslavia. None of these pages figure in the data used here.[9]

This highlights both some of the sources of noise in the Wiki data and the limitations of cross-country comparisons. Since different country histories and online contest over them are affected by different idiosyncrasies, there is likely more useful information in within-country comparisons. The regression-based analysis below therefore limits itself to within-country variation by always including country-specific fixed effects.

Even so, seeing Israel and the Balkans at the top of the list rather than, say, Scandinavia and some tiny island countries is reassuring. The map below displays the residuals for the entire world, and is again reasonably in line with expectations, with the Balkans being the most conflictual location and Switzerland being the least. Regions like South America and Africa generally seem less conflictual than one would anticipated, however; perhaps due to a smaller number of editors active on the English Wikipedia.

[9] In principle, they could be added, of course, but this would raise some new and very thorny issues on how to decide which present-day countries to assign conflict from articles on Yugoslavia and other past entities, and which of the huge amount of relevant history pages about specific events and topics to add to the data.

Figure 1: Excess reverts across the world

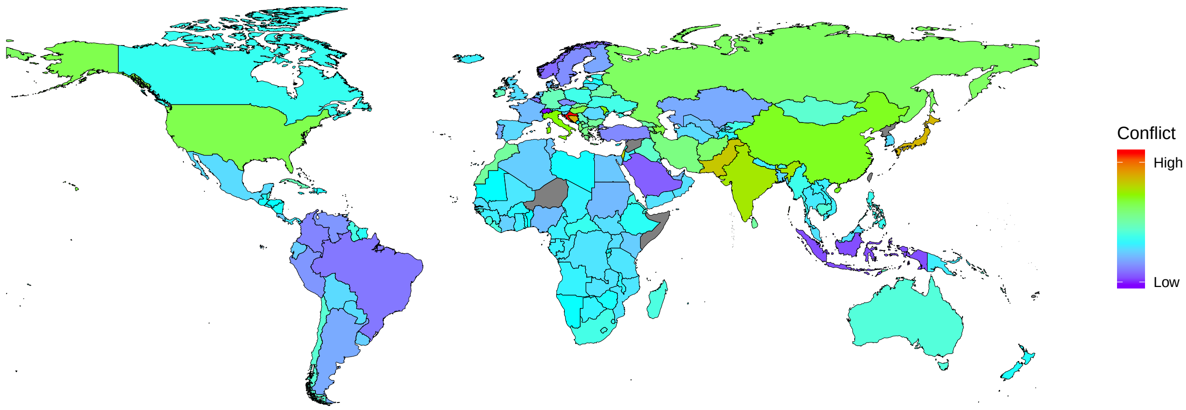


Though the overall patterns thus appear to make some sense, a better understanding of the validity of the measure can be had by comparing the revert variable with other plausible and time-variant indicators of conflict on Wikipedia. The most natural candidate here is talk page sentiment, meaning the tone of the discussion on the forums attached to each history page. If the revert measure is in fact capturing conflict between editors, the tone of the discussion should on average be more hostile in country years with many mutual reverts than in country years with few. People fighting over history are likely to get angry with each other.

The most common way to capture text sentiment is by applying an unsupervised dictionary approach. A dictionary in this sense is simply a long, prefabricated list of words each associated with a score that aims to capture whether a given word is negative or positive and sometimes the degree of negativity or positivity (Taboada et al., 2011). In the widely used sentiment lexicon from Bing Liu and collaborators, for example, the words "accomplished" and "accurate" are labelled as positive, whereas the words "abominably" and "abysmal" are labelled as negative (Hu and Liu, 2004).

To my knowledge, none of the publicly available dictionaries were created for analysing Wiki talk pages (the Bing dictionary was developed for customer reviews of online merchandise), and as is often the case when transplanting a learning procedure to a different environment, they tend to perform poorly here. This is mainly due to the nature of the content on the national history pages. Many completely peaceful and constructive Wiki discussions about national history contains words like "war", "death", "killings" and the like, because such words are simply relevant descriptors of historical events. An unsupervised approach from a different context, however, will frequently label such words as negative, obscuring the actual sentiment of the discussion.

To remedy this, I develop a small, but highly specific dictionary of words that as unambiguously as possible identify conflict on Wikipedia talk pages. These words are based on a close reading of numerous talk pages along with the guidelines for discussion produced by the Wiki community, and most of them are highly Wiki-specific abbreviations or links that editors use to indicate that their opponents are in violation of the rules of conduct. Examples include "NPOV" ("neutral-point-of-view"; a reference to the guiding principle of writing on Wikipedia), "WP:CIVIL" (a reference to a specific Wiki subsite containing guidelines for civil conduct), "WP:NOTPROPAGANDA" (a reference to a similar subsite about propaganda on Wiki), "Wikipedia:Conflict_of_interest" (a reference to a subsite about rules on conflict of interest), and "sockpuppet" (a common internet term for alternative user accounts exploited to break rules without having one's main account sanctioned).[10] To these Wiki-specific words I add the same list of profanities employed earlier.

With the dictionary of conflict words in place, I identify all the words added to a talk page by a given revision and then simply count the number of these that appear in the dictionary. Manual reading of a sample of talk page revisions with negative and non-negative scores reveals that this procedure yields a fairly accurate identification of conflict, though with somewhat of a tendency to miss the more civil and understated occurrences. To correlate it with reverts, I finally aggregate the talk page score to the country-year level by summing.

Figure 2 plots within-country and within-year slope estimates for the associaton between reverts (mutual and without bots and vandals) and negative talk page sentiment with 95 % confidence intervals.

[10] The full dictionary is located in the appendix.

Figure 2: Reverts and negative talk page sentiment
(with country and year FE)

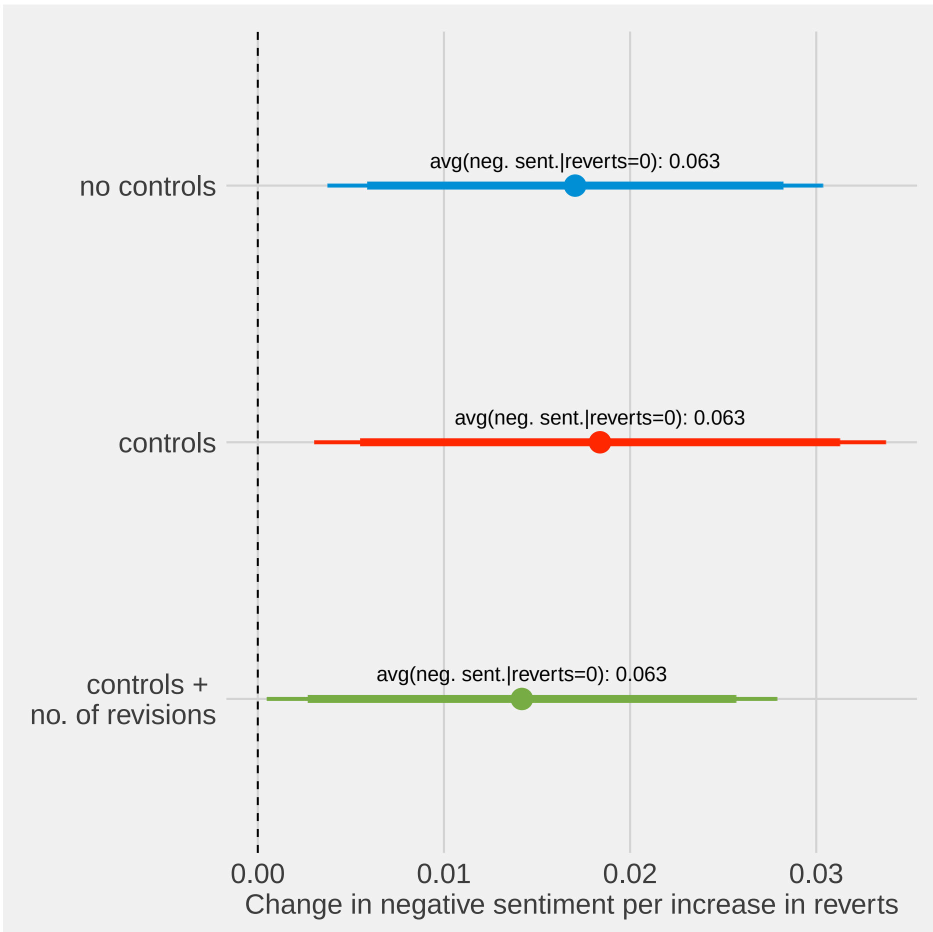


controls: log(GDP per capita), log(population), GDP growth, and Polity2

The plot shows that increases in revert frequency is quite predictive of increases in talk page negativity. An increase of one standard deviation in reverts is associated with a change of 0.16 standard deviations in negativity (results are similar without fixed effects), and including yearly variables for GDP per capita, population size, GDP growth, and democracy level (Polity2) in the model makes practically no difference. Even when adding the number of revisions in a country year, the association stays strong.[11] This means that for a given level of user activity on a history article, the more reverts there are the more negative sentiment one can expect to see on the talk page. This pattern is highly consistent with the revert measure actually capturing true and substantive conflict over history.

In summary, there is considerable evidence to suggest that the preferred revert measure is a serviceable proxy for conflict over history on Wikipedia.

## Results

Before correlating such conflict with civil war, it would be informative to get an idea of what the conflicts are about. We know now that conflict is associated with negative sentiment, but what kinds of themes of history are they associated with?

To answer this question, one needs to examine the revision text itself and determine whether conflictual revisions generally deal with different topics than non-conflictural revisions. The ideal way to do this would likely be a human reading of revisions, but this is clearly infeasible given the sheer size of the text material. To investigate topics, I instead turn to topic models.

Broadly, topic model work by taking as input a set of documents and some pre-specified parameters (often the number of topics to be found) and outputting a probability distribution over topics for each document, where a topic is in itself a probability distribution over the collection of words in the documents. Topic models are generative models, meaning that they assume some probabilistic process that generates documents and then use the observed data to find the most likely values for the parameters of this process. The data generating process is thought of as a series of draws of words, such that for word position $i$ in document $j$, one first draws a topic $k_{ij}$ from the topic distribution for $j$, $p(k|j)$, and then draws a word $w_{ij}$ from the word distribution of $k$, $p(w|k)$. The key parameters to be estimated from the data are thus the probabilities associated with each word in each topic and the probabilities associated with each topic in each document.[12] The Structural Topic

[11] If the models were attempting to identify the effect of reverts on sentiment, one would probably not want to include the number of revisions, since this number is bound to be "post-treatment" or at least simultaneously given with respect to conflict – more conflict and controversy leads to more activity, and more activity also leads to more conflict. Here, however, the objective is not identification, but merely inspection of contingent patterns.

[12] See e.g. Liu et al. (2016) for an overview of topic models.

Model (STM) is a newly developed variant that allows topic prevalence to be correlated with arbitrary metadata, such as whether or not a given revision was reverted against or at what time it was created (Roberts, Stewart and Tingley, Forthcoming).

Before applying an STM or any other topic model to the revision texts, however, some considerations about the nature of these documents are in order. First, because we are interested in the topics that an editor is contributing to when submitting a new revision, we will not get much mileage from using the full revision texts. Since editors usually do not change everything about an article at once, the text of revision $i$ will overall tend to be very similar to the text of revision $i-1$. Often, an edit will be about a particular period or event in a country's history, such that much of the text will be left unchanged. This means that a topic model trained on the full texts will – rightly – uncover the same or very similar topic distributions for $i$ and $i-1$, thus revealing little information about what themes the edits turning $i-1$ into $i$ were about. Rather, this information resides in the *differences* between two revisions, which I operationalize as the concatenation of all whitespace-separated tokens (generally words) added and removed between $i$ and $i-1$. For example, if $i-1$ consists of the string $A\ A\ B$ and $i$ consists of the string $A\ C\ D\ C$, the list of removed tokens would be $[A, B]$ and the list of added tokens would be $[C, C, D]$. The total differences in this sense would then be $[A, B, C, C, D]$. The fact that this approach necessarily breaks up the grammatical structure of the texts is of no consequence here, since the topic model employed below is a bag-of-words-model, meaning that it simply counts the presence or absence of a word without caring about its position in a sentence.

Second, one is equally unlikely to get much useful information from training a topic model on multiple national history pages simultaneously. This is again because successive versions of the same history pages – and even their differences – are bound to be fairly similar. Revisions to the same page stem from a related data generating process, and so the realizations are of it are naturally clustered. When a model is trying to find topics that are likely to have generated history texts about countries $A$, $B$, and $C$, the best topics will often just be $A$, $B$, and $C$. Since country histories have many country-specific words (place names, events, names of individuals, important years etc.), a good model will tend to create a topic with high probabilities of country $A$-specific words and assign all the revisions to the history of $A$ a high probability of being about that topic and low probabilities of being about the inevitable country $B$- and country $C$-specific topics. Using revisions from all the countries in the data is bound to just lead to topics with very high probabilities for given country names, capitals, and historic figure assigned a high probability in revisions to their respective country histories.[13]

The implication of this is that we have to focus on a single country history (at a time).

[13]Naturally, the way this process plays out depends on the number of topics specified relative to the number of countries included, but the general problem remains.

In keeping with the theme, I again draw on the most conflictual page in the data, the history of Croatia. This also provides a nice sanity check of the proposed mechanism: If conflict over history is indeed associated with real-world war, one should expect many of reverted revisions on the most conflictual history page to be about war- and conflict-related topics.

To estimate an STM on the differences between successive revisions to the history of Croatia, I first perform a series of standard preprocessing steps including setting words to lowercase, removing stop words, numbers, punctuation, and extremely short or infrequent words, and stemming words such that different conjugations of the same basic concept figure as the same token (reducing, for example, the words "argue", "argued", and "argues" to the same root, "argu"). I then inspect results from a set of STM's with different numbers of topics in order to determine how many topics are appropriate for the resulting text corpus. There is no mechanical way of selecting the right number of topics (Grimmer and Stewart, 2013), but a combination of human reading and assessment of quantities like semantic coherence, held-out likelihood and exclusivity goes a long way. Over a search space from 10 to 70, I find that 47 topics provide and informative simplification of the development of the history page for Croatia.

Figure 3 presents results from an STM with 47 topics and reverts as a covariate predicting topic prevalence. The y-axis gives the expected frequency of each topic across the corpus, and the x-axis gives the difference in this expected frequency between revisions with reverts and revision without. A positive difference indicates that the topic is more prevalent in revisions that are reverted against. The list of tokens above the point estimates are the words with the highest FREX-scores for that topic, meaning that they are frequent in the corpus overall yet fairly exclusive to the topic (Airoldi and Bischof, 2016). These words give a good idea of what each topic is "about".

Figure 3: Difference in topic proportion for reverts and non-reverts
(with highest FREX-words as labels)

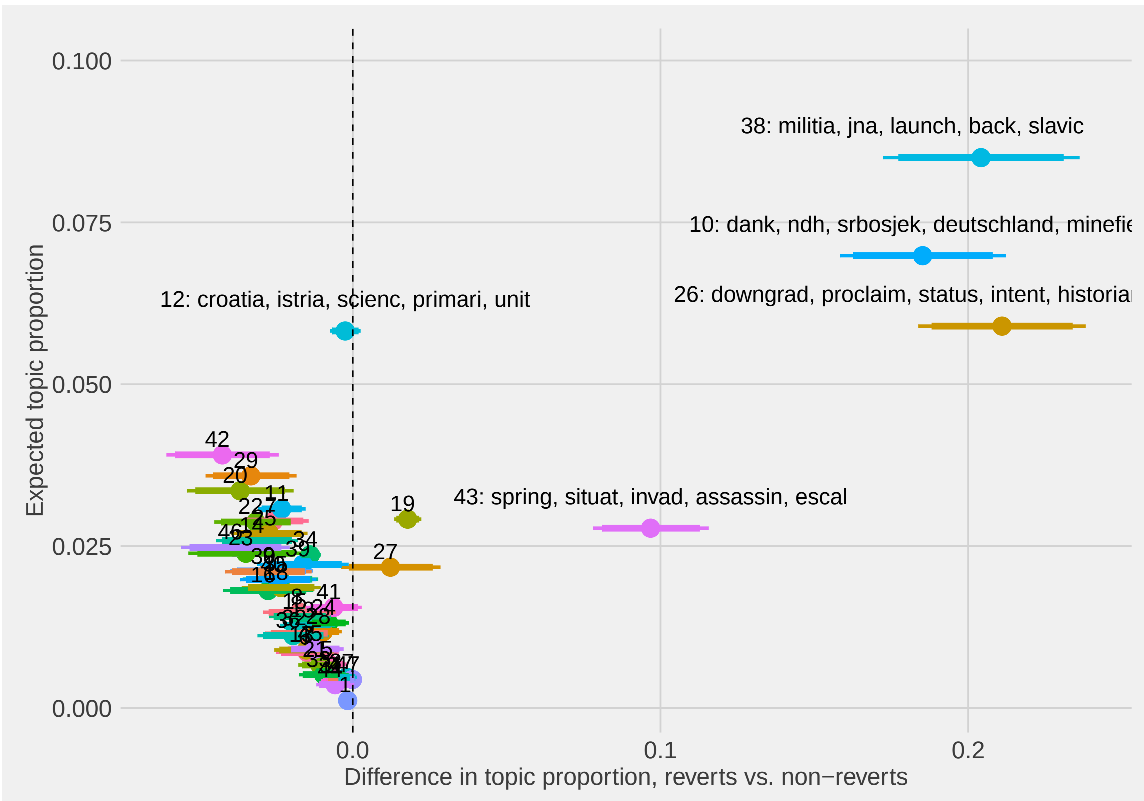

It is immediately clear from the plot that there is a large number of more or less redundant topics that are not very frequent, nor very correlated with conflict (lower left corner). Further, Topic 12 is fairly frequent, but not correlated with conflict at all. Interestingly, however, the three most frequent topics – covering 6, 7, and 8.5 % of the text, respectively – are highly positively correlated with conflict (upper right corner). Reading the full list of FREX words for these three topics and the documents most associated with them give a clear indication that they are indeed highly war-related.

Topic 38, with the top-5 FREX words: "milita", "jna", "launch", "back", and "slavic", appears to be about the Croatian War of Independence from 1991 to 1995. "milita" likely refers to the militias from the break-away Republic of Serb Krajina, who did much of the fighting and ethnic cleansing, while "jna" is the abbreviation for the Yugoslav National Army, which fought on the Serbian side in the first two years. "back" is more ambiguous, and "slavic" just seems to refer to the entire ethnic dimension of the conflict, but "launch" seems to be about the launch of the two Croatian offensives, *Flash* and *Storm*, which effectively ended the war in Croatia's favour.

A natural interpretation of Topic 10 is that it concerns WWII and the Ustaša regime. "ndh" is the Croatian abbreviation for the Independent State of Croatia, the fascist puppet regime established in Croatia by the Nazis and run by the Ustaša, while "srbosjek" means "Serbian cutter" and was a type of shear knife adapted for speedy execution of prisoners in Croatian concentration camps during WWII.[14] "dank" and "deutschland" are both from the 1991 song *Danke Deutschland* by Sanja Trumbić, which thanked Germany for its role in Croatia's international recognition and was broadcast on Croatian national TV the day before independence, but was seen by many Serbs as a clear reference to the country's fascist past. "minefield" is naturally also a reference to war, but could refer to both WWII minefields and minefields from the Yugoslav Wars.

Finally, Topic 26 appears to touch on the origins of the Croatian War of Independence. "downgrad", "status", and "proclaim" likely refer to the fears of the Serbian minority in Croatia in 1991 about their future rights and status after Croatia had proclaimed independence – fears that strongly contributed to the outbreak of the ensuing war (Fearon, 1994). A closer reading of documents associated with Topic 26 suggests that "intent" and "historian" may both be references to Franjo Tudjman himself who was indeed a historian by trade and whose intentions with Croatian independence vis-à-vis the Serbian minority is and was a subject of intense debate.

As is clear from this exercise, interpreting the output of topics models is more art than science and there are generally no right or wrong interpretations, so other readers could come to different conclusions about the proper labels for the three probability distributions.

[14] As a testament to the gruesome efficiency of the weapon, Petar "Pero" Brzica, a Jasenovac guard, reportedly won a bet by executing over 1,300 prisoners in a single night with a srbosjek.[REF XXX]

Considering the FREX word rankings and the content of strongly associated documents, however, the above interpretations seem reasonable, though again a bit post hoc. If one accepts that, we have the interesting result that the most prevalent topics in the most conflictual history page, the history of Croatia, are not only substantively related to war, but also highly correlated with reverts.

Sanity thus checked we can proceed with the main analysis: Does civil war produce online conflict over history? To answer this question, I first add a dummy for civil war incidence to the country year panel. This takes a value of one in country years with an ongoing civil war that at some point caused 1000 or more battle-related deaths according to the Uppsala Conflict Data Program (UCDP) and a value of zero in all other country years (Pettersson and Eck, 2018). As the UCDP time series end in 2016, I limit the panel to the period 2001-2016. Also, I exclude microstates with a population of less than a million people.

I then estimate the following two-way fixed effects model:

$$Reverts_{it} = \beta_0 + \beta_1 Civil\ war_{it} + \boldsymbol{\gamma} \boldsymbol{X_{it}} + a_i + u_t + v_{it}, \tag{1}$$

where $i$ indexes country, $t$ indexes time, $\beta_0$ is a constant, *Civil war* is the aforementioned dummy for civil war incidence, and $\boldsymbol{X}$ is a vector of controls. The three remaining terms represent unobserved determinants of *Reverts*. $a$ is time-invariant and is captured by a full set of country fixed effects, $u$ is unit-invariant and is captured by a full set of year fixed effects, and $v$ contains unobserved determinants that vary across both country and year. The key identifying assumption here is thus $E(v|Civil\ war, \boldsymbol{X}, a, u) = 0$.

The control vector contains only the bare minimum of the usual suspects in the conflict studies literature: log(population), log(GDP per capita), GDP growth (%), and Polity2 score. In some specifications I add a count of the number of revisions (edits) made in a country year to measure general activity, but due to the likely simultaneity with conflictual behaviour (people edit more when they are fighting other editors), I prefer models without it.

Though *Reverts* is strictly speaking a count variable, I treat it as continuous in the main analyses and estimate equation 1 with OLS, clustering standard errors on country.[15] The results are presented in Table 3 and the estimate on *Civil war* for columns 1-3 are plotted in Figure 4.

[15] In the appendix I re-estimate equation 1 with fixed effects poisson models, which yields similar, if slightly stronger results

Table 3: Reverts and intrastate conflict incidence

| | *Dependent variable:* | | | |
|---|---|---|---|---|
| | No. of reverts | | | No. of revisions |
| Civil war | 0.77* | 0.95** | 0.54* | 31.69* |
| | (0.43) | (0.49) | (0.29) | (16.99) |
| log(Populaton) | | 2.94 | 3.40* | −35.27 |
| | | (2.04) | (1.97) | (41.71) |
| log(GDP per capita) | | 0.65 | 1.10 | −34.85 |
| | | (1.46) | (1.37) | (26.89) |
| GDP growth | | 0.02 | 0.02 | 0.47 |
| | | (0.04) | (0.03) | (0.61) |
| Polity2 | | 0.01 | 0.03 | −1.52 |
| | | (0.04) | (0.03) | (1.03) |
| No. of revisions | | | 0.01*** | |
| | | | (0.004) | |
| Country FE's | Yes | Yes | Yes | Yes |
| Year FE's | Yes | Yes | Yes | Yes |
| avg(outcome\|civil war=0) | 1.16 | 1.13 | 1.13 | 85.68 |
| Observations | 2,391 | 2,273 | 2,273 | 2,273 |

*Note:* *p<0.1; **p<0.05; ***p<0.01
se's clustered by country

Figure 4: Reverts and intrastate conflict incidence

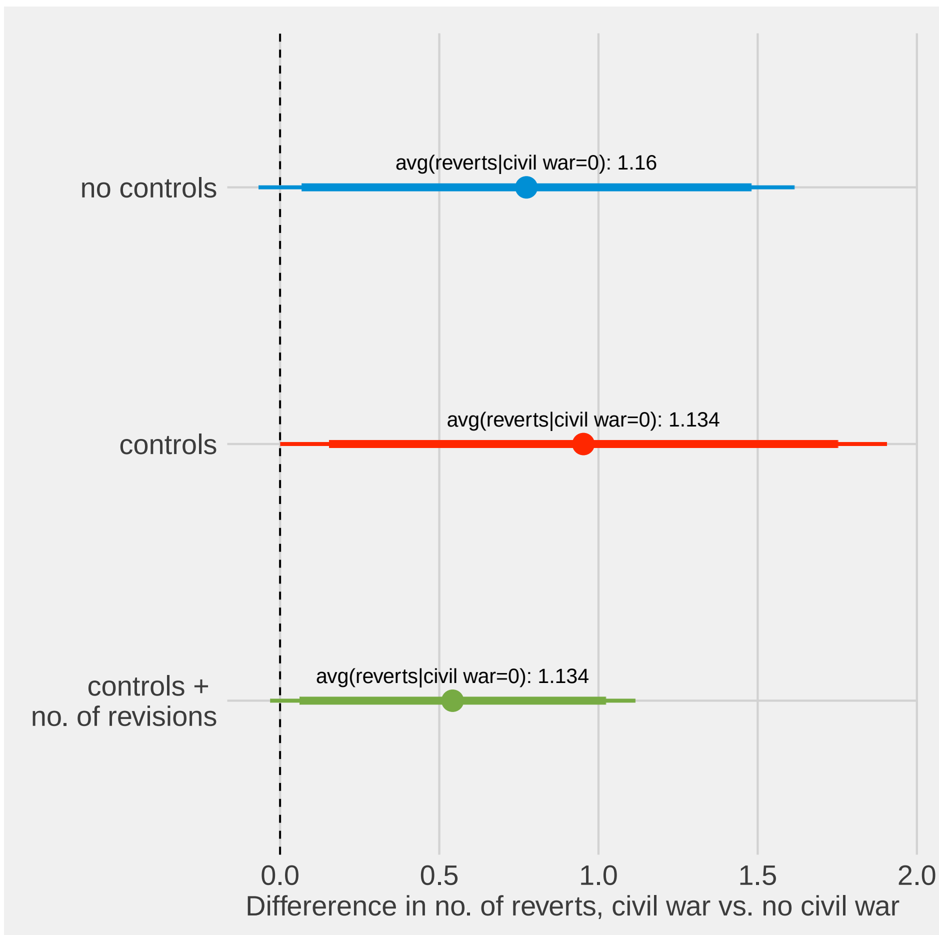

Table 3 and Figure 4 provide strong support for the idea that civil war is followed by increased conflict over national history. Throughout columns 1-3 the estimate on *Civil war* is positive, large, and quite precisely estimated ($p$-values ranging from 0.05 to 0.07). In substantive terms, models 1 and 2 predict that the number of reverts is between 66 and 84 % higher in periods with an ongoing civil war compared to periods without one. Remarkably, even when holding the number of revisions constant in column 3, the estimate on *Civil war* remains large and fairly precise ($p = 0.06$).

Column 4 swaps *Reverts* for *No. of revisions* as the outcome variable, and shows that general revision activity is also increased during civil war, though only by about 37 %. Thus, under a causal interpretation, civil war has a much greater relative effect on the conflict on national history pages than it does on the general revision activity on these pages. None of the controls reach conventional levels of significance, which may partly be due to a large share of their variation over the relatively short time period being captured by the country fixed effects.

### Are they fighting over "ancient" history?

So there seems to be evidence that civil war or particular events occurring during civil wars make people fight over the representation of national history. A natural question to ask is then: But what periods of history?

Since "history" on Wikipedia is fairly loosely defined and basically covers anything in the past, the increased conflict I observe during civil wars could be due to people fighting over the contemporary causes of the civil war; such as who cheated most in the election last year, whether or not the government massacred those protesters or just quelled a riot, or if the opposition was connected to that terrorist group that detonated a car bomb in the town square. While interesting, such debates do not really capture the identity-forming types of events that the theoretical argument is about. Instead, whether arising strategically or organically, discursive struggles over group identity conveys an image of more "fundamental" historical events being disputed. Intuitively, we would tend to think of the topics of such struggles as being relatively far in the past – maybe at least a generation or so.

To investigate whether civil war merely leads to more fighting over history because people fight over the war's contemporary causes, I need to figure out what historical period or periods each revision deals with. Here, I take the simple approach of focussing on which calendar years a revision edited. The first step to this is to identity all calender years in the entire text corpus. While finding numbers in the text is no problem, determining whether a given number is a calender year or not is more complicated. In free prose the number "1648" could refer to the year of the Peace of Westphalia, the number of troops deployed in some battle, the number of miles between two locations, or something else entirely, depending on the context. To distinguish years from other numbers I therefore need to build a classifier that takes this context into account.

To do this, I first extract all sentences with positive integers from the 247,254 Wiki texts, defining a sentence as a series of one or more tokens delimited by punctuation marks.[16] This amounts to about 50 million sentences. Many of these are duplicates, since some sentences have multiple numbers in them and each has a different set of tokens surrounding it, thus requiring separate classification.

From the 50 million sentences I randomly sample 10,000 to classify by hand according to whether the integer in question is judged to be a calendar year or not. Though the true character of each integer is of course unobserved, my assessment after completing this task is that a human can generally distinguish remarkably precisely between years and other integers using only information from the sentence in which they occur. With the training data complete I next extract a number of features from each sentence.

These features consist of a full list of all words occurring between one and three places before and after the integer in question, where each word has a separate entry for each of the six possible positions around the integer. This means, for example, that I distinguish between the significance of the two "in"'s in the sentence "the conflict escalated into armed incidents *in* 1991 *in* the majority-Serb populated areas". The first position of "in" relative to the integer turns out to be a powerful predictor of calendar year whereas the second does not. I also add a number of features of the integer itself, such as the number of digits it consists of and whether or not it is within certain common ranges for years.

Using these features I train a Support Vector Machine (SVM) with a linear kernel on the 10,000 training examples. Basically, an SVM is a supervised learning model which attempts to find the hyperplane that provides the best separation between positive examples (years) and negative examples (other integers) with the rule for "best" being that the hyperplane must have the largest possible Euclidian distance to the nearest points. The vectors parallel to the hyperplane and touching these points are then the eponymous "support vectors". Once such a hyperplace is found, it can be used to classify new points in the same space according to which side of the plane they fall. In this case the SVM performs particularly well, and achieves recall and precision levels for both years and non-years close to 100 %. When using a set of purely deterministic rules, like whether the integer has four digits, is below some threshold for realistic years, and is preceded by one of a list of tokens such as "in", "during" etc., recall and precision are down to about 80 %, so all the trouble with the supervised approach adds a good chunk of information. The trained SVM is finally used to classify the remaining 49,990,000 integers.

[16] The full rule set used to identity sentences with numbers is a bit more complex; see e.g. the appendix, which also contains Python code with the regular expressions and general framework used for these tasks.

With all years in the text thus identified,[17] I can then apply the same difference approach used for the topic model above to determine what, or in this case, which period(s), a revision is about. Accordingly, if revision $i-1$ contains the years 253, 1389, and 1999, and $i$ contains 253, 1400, and 1999, I code the edit turning $i-1$ into $i$ as dealing with the years 1389 and 1400. To get one number for each revision, I simply take the median of the resulting list of years, and to get to the country year level I take the median again.

Merging these medians of changed years to the country year panel, I can then estimate the following twoway fixed effects model:

$$\begin{aligned} Median\ years_{it} = \beta_0 + \beta_1 Civil\ war_{it} + \beta_2 Reverts_{it} + \beta_3 Civil\ war_{it} * Reverts_{it} \\ + \boldsymbol{\gamma X_{it}} + a_i + u_t + v_{it}, \quad (2) \end{aligned}$$

where $Median\ years$ is the median of changed years and everything else has the same interpretation as in equation 1.

While it is interesting in itself to see if conflict over history is generally about different time periods than peaceful revisions, what we are really interested in is the interaction with civil war. The first derivative of equation 2 with respect to $Reverts$, $\beta_2 + \beta_3 * Civil\ war$ yields the change in median years per increase in reverts. If the estimate of $\beta_3$ is positive and precise, this means that when there is a civil war, reverts tend to be about more recent event than they otherwise are. If not, then the civil war-induced reverts are just as much about old history as other reverts.

Table 4 presents the results from estimating equation 2 with OLS, again clustering standard errors on country. Figure 5 shows the predicted change per revert in median years for periods with civil war and periods without (columns 3 and 4 in Table 4).

[17] Again the process is a bit more complex. See e.g. the appendix for the full rule set converting integers and classifications into actual years.

Table 4: Median years changed in revision texts

| | *Dependent variable:* | | | |
|---|---|---|---|---|
| | Median year | | | |
| Reverts | 1.23 | 1.71* | 1.27 | 1.78** |
| | (0.85) | (0.88) | (0.90) | (0.90) |
| log(Populaton) | | −205.59 | | −199.46 |
| | | (154.47) | | (149.11) |
| log(GDP per capita) | | −36.63 | | −32.20 |
| | | (79.30) | | (77.86) |
| GDP growth | | 2.20 | | 2.41 |
| | | (1.39) | | (1.50) |
| Polity2 | | 3.71 | | 3.53 |
| | | (5.74) | | (5.64) |
| Civil war | | | −12.45 | 53.99 |
| | | | (95.55) | (93.43) |
| Reverts*Civil war | | | −0.22 | −1.66 |
| | | | (3.01) | (3.79) |
| Country FE's | Yes | Yes | Yes | Yes |
| Year FE's | Yes | Yes | Yes | Yes |
| avg(year\|reverts=0) | 1742.56 | 1743.37 | 1742.56 | 1743.37 |
| Observations | 2,253 | 2,141 | 2,253 | 2,141 |
| *Note:* | | *p<0.1; **p<0.05; ***p<0.01 | | |
| | | se's clustered by country | | |

Figure 5: Median years changed in revision texts

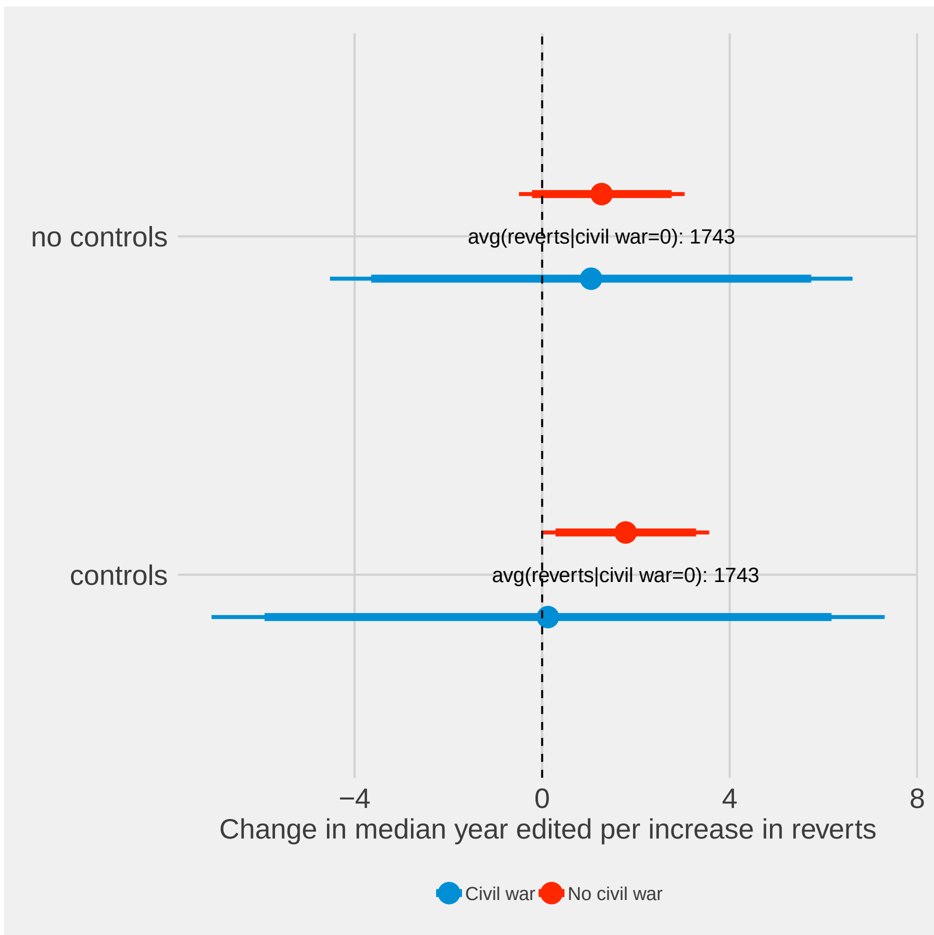

Somewhat depending on specification, Table 4 shows that the more reverts there are, the more recent history is generally being edited, implying that fights are more about newer history than older. However, the increases are only at a scale of one or two years per revert, which, given that the average year edited in country years without conflict is in the middle of the 18th century and that reverts are quite rare, is negligible.

Much more interestingly, the estimate of $\beta_3$ is negative and far from standard levels of significance, which suggests that the extra reverts brought about by civil war are *not* merely about the contemporary causes of the war. Instead, they deal just as much with the ”ancient” past as other reverts – on average with the year 1750.

Together with the main result of increased fighting over history during civil wars, this has profound theoretical implications.

[PROFOUND THEORETICAL IMPLICATIONS COMING HERE]

## Conclusion

[CONCLUSION COMING HERE]

## References


Airoldi, Edoardo M and Jonathan M Bischof. 2016. "Improving and evaluating topic models and other models of text." *Journal of the American Statistical Association* 111(516):1381–1403.

Barkin, J Samuel and Laura Sjoberg. 2017. *Interpretive quantification: methodological explorations for critical and constructivist IR*. Ann Arbor, MI: University of Michigan Press.

Bennett, Andrew. 2015. "Found in translation: Combining discourse analysis with computer assisted content analysis." *Millennium* 43(3):984–997.

Blattman, Christopher and Edward Miguel. 2010. "Civil war." *Journal of Economic literature* 48(1):3–57.

Brass, Paul R. 1997. *Theft of an idol: Text and context in the representation of collective violence*. Princeton, NJ: Princeton University Press.

Cederman, Lars-Erik, Andreas Wimmer and Brian Min. 2010. "Why do ethnic groups rebel?: New data and analysis." *World Politics* 62(1):87–119.

Cederman, Lars-Erik, Nils B Weidmann and Kristian Skrede Gleditsch. 2011. "Horizontal inequalities and ethnonationalist civil war: A global comparison." *American Political Science Review* 105(3):478–495.

Cinar, Alev. 2001. "National history as a contested site: The conquest of Istanbul and Islamist negotiations of the nation." *Comparative Studies in Society and History* 43(2):364–391.

Collier, Paul and Anke Hoeffler. 2004. "Greed and grievance in civil war." *Oxford economic papers* 56(4):563–595.

Correlates of War Project. 2017. "State System Membership List, v2016." `http://correlatesofwar.org`. [Online; accessed 07-December-2017].

Dyrstad, Karin. 2012. "After ethnic civil war: Ethno-nationalism in the Western Balkans." *Journal of Peace Research* 49(6):817–831.

Fearon, James D. 1994. Ethnic war as a commitment problem. In *Annual Meetings of the American Political Science Association*. pp. 2–24.

Fearon, James D. and David D. Laitin. 2000. "Violence and the social construction of ethnic identity." *International Organization* 54(4):845–877.

Fearon, James D. and David D. Laitin. 2003. "Ethnicity, insurgency, and civil war." *American Political Science Review* 97(1):75–90.

GeoNames. 2017. "Geonames languages." `http://download.geonames.org/export/dump/countryInfo.txt`. [Online; accessed 15-December-2017].

Glaeser, Edward L. 2005. "The political economy of hatred." *The Quarterly Journal of Economics* 120(1):45–86.

Grimmer, Justin and Brandon M Stewart. 2013. "Text as data: The promise and pitfalls of automatic content analysis methods for political texts." *Political analysis* 21(3):267–297.

Guex, Samuel. 2015. "The History Textbook Controversy in Japan and South Korea." *Cipango-French Journal of Japanese Studies. English Selection* (4).

Howard, Michael. 1979. "War and the nation-state." *Daedalus* 108(4):101–110.

Hu, Minqing and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM pp. 168–177.

Jalal, Ayesha. 1995. "Conjuring Pakistan: History as official imagining." *International Journal of Middle East Studies* 27(1):73–89.

Kalyvas, Stathis N. 2008. "Ethnic defection in civil war." *Comparative Political Studies* 41(8):1043–1068.

Kolstø, Pål. 2011. The Serbian-Croatian Controversy over Jasenovac. In *Serbia and the Serbs in World War Two*, ed. Sabrina P. Ramet and Ola Listhaug. London: Palgrave Macmillan pp. 225–246.

Kubota, Yuichi. 2017. "Imagined Statehood: Wartime Rebel Governance and Post-war Subnational Identity in Sri Lanka." *World Development* 90:199–212.

Liu, Lin, Lin Tang, Wen Dong, Shaowen Yao and Wei Zhou. 2016. "An overview of topic modeling and its current applications in bioinformatics." *SpringerPlus* 5(1):1608.

Lowensohn, Josh. 2014. "Russia spotted editing Wikipedia page about downed Malaysia Airlines jet." `https://www.theverge.com/2014/7/18/5917099/russia-spotted-editing-wikipedia-page-of-downed-malaysia-air-jet`. [Online; accessed 29-July-2018].

Mill, John Stuart. 1910. *Utilitarianism, liberty, and representative government*. London: Dent.

Monroe, Burt L, Jennifer Pan, Margaret E Roberts, Maya Sen and Betsy Sinclair. 2015. "No! Formal theory, causal inference, and big data are not contradictory trends in political science." *PS: Political Science & Politics* 48(1):71–74.

Pettersson, Therése and Kristine Eck. 2018. "Organized violence, 1989–2017." *Journal of Peace Research* .

Preneel, Bart. 2010. The first 30 years of cryptographic hash functions and the NIST SHA-3 competition. In *Cryptographers' track at the RSA conference*. Springer New York: pp. 1–14.

Renan, Ernest. 1882. “What is a nation?” https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=2&cad=rja&uact=8&ved=0ahUKEwiSmrrhtKvcAhXNyKQKHX5cCXoQFgg1MAE&url=http%3A%2F%2Fucparis.fr%2Ffiles%2F9313%2F6549%2F9943%2FWhat_is_a_Nation.pdf&usg=AOvVaw1f2lVKlietpYgKjcYI7oXk. [Online; accessed 14-July-2018].

Rivest, R. 1992. “The MD5 Message-Digest Algorithm.” https://www.rfc-editor.org/info/rfc1321. [Online; accessed 13-Aug-2018].

Roberts, M., B. Stewart and D. Tingley. Forthcoming. “stm: R Package for Structural Topic Models.” *Journal of Statistical Software* pp. 1–42.

Sampson, Tim. 2013. “How pro-fascist ideologues are rewriting Croatia’s history.” https://www.dailydot.com/layer8/croatian-wikipedia-fascist-takeover-controversy-right-wing/. [Online; accessed 29-July-2018].

Simmel, George. 2010. *Conflict and the web of group affiliations*. New York, NY: Simon and Schuster.

Smith, Anthony D. 1986. “The ethnic origins of nations.”.

Taboada, Maite, Julian Brooke, Milan Tofiloski, Kimberly Voll and Manfred Stede. 2011. “Lexicon-based methods for sentiment analysis.” *Computational linguistics* 37(2):267–307.

Tilly, Charles. 1990. “Coercion, capital, and European states, AD 990.”.

Vogt, Manuel, Nils-Christian Bormann, Seraina Rüegger, Lars-Erik Cederman, Philipp Hunziker and Luc Girardin. 2015. “Integrating data on ethnicity, geography, and conflict: The ethnic power relations data set family.” *Journal of Conflict Resolution* 59(7):1327–1342.

Wikimedia. 2012. “Wikipedia Editors Survey 2012.” https://dumps.wikimedia.org/other/surveys/editorsurvey2012/. [Online; accessed 30-July-2018].

Wikimedia. 2018. “Total page views, monthly.” https://stats.wikimedia.org/v2/. [Online; accessed 31-July-2018].

Wikipedia. 2018. “Wikipedia:Bots (revision ID: 849594788).” https://en.wikipedia.org/w/index.php?title=Wikipedia:Bots&oldid=849594788. [Online; accessed 31-July-2018].

Wilmer, Franke. 2004. *The social construction of man, the state and war: Identity, conflict, and violence in former Yugoslavia*. Routledge.

Wimmer, Andreas, Lars-Erik Cederman and Brian Min. 2009. “Ethnic politics and armed conflict: a configurational analysis of a new global data set.” *American Sociological Review* 74(2):316–337.

Wood, Elisabeth Jean. 2008. "The social processes of civil war: The wartime transformation of social networks." *Annu. Rev. Polit. Sci.* 11:539–561.

Woodward, Susan L. 1995. *Balkan tragedy: Chaos and dissolution after the Cold War.* Washington, D.C., DC: Brookings Institution Press.

World Bank. 2017. "World Development Indicators." `https://datacatalog.worldbank.org/dataset/world-development-indicators`. [Online; accessed 13-December-2017].

Wucherpfennig, Julian, Philipp Hunziker and Lars-Erik Cederman. 2016. "Who Inherits the State? Colonial Rule and Postcolonial Conflict." *American Journal of Political Science* .

Yanagizawa-Drott, David. 2014. "Propaganda and conflict: Evidence from the Rwandan genocide." *The Quarterly Journal of Economics* 129(4):1947–1994.

Yasseri, Taha, Robert Sumi, András Rung, András Kornai and János Kertész. 2012. "Dynamics of conflicts in Wikipedia." *PloS one* 7(6):e38869.