Converging "Truths"? The Role of Digital Knowledge Platforms for Information Flows across Language Barriers and History

Discord, polarization, and "parallel truths" on digital platforms are widely believed to threaten social cohesion, with serious repercussions for political and economic stability. The erosion of mutually agreed 'truth' on social networks risks undermining the foundations of societal consensus in democratic societies, prompting substantial empirical research in Economics and IS on polarization and filter bubbles. However, competing and contested narratives about past events are already familiar to historians and memory studies, though these fields have yet to study narrative divergence quantitatively or at large scale.

We propose a groundbreaking, scalable quantitative analysis of how diverging versions of 'the truth' evolve and reconcile on digital platforms across different languages. By combining cross-lingual Large Language Models (LLMs), Artificial Intelligence (AI), modern econometric methods, and data acquisition through scraping and APIs to Wikipedia's unique features as a language-specific knowledge repository, this approach enables unprecedented precision and scale in studying these questions.

Specifically we develop a new methodology to contrast the revision histories of over 1000 historic encyclopaedic entries ("topics") in over 50 languages and for over 20 years (over 10,000,000 textual observations), to (1) quantify conflicts in knowledge, (2) analyze how these conflicts grow or diminish over time, and (3) shed light on factors and mediators which affect the conflict between and co-evolution of 'truths' and informational flows across different languages.

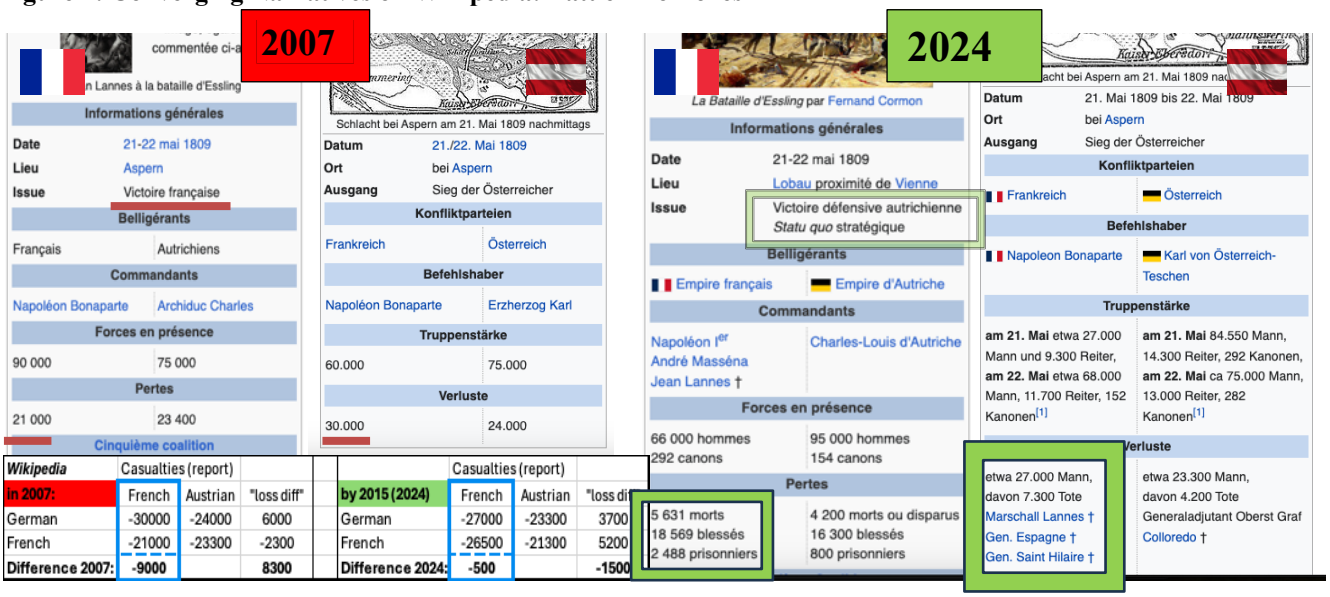
Keywords: Open Knowledge Platforms, AI Translation, LLMs & NLP, Flow of Information, Translation Skills, Misinformation, Polarization, Bias

:

Discord, polarization and "parallel truths" on digital platforms are widely believed to threaten social cohesion, with potentially serious repercussions on political and economic stability in democratic societies. Because the erosion of mutually agreed 'truth' on social networks could threaten the foundation for societal consensus, this phenomenon has inspired a sizeable body of empirical research about polarization on social media platforms in economics and information systems among other disciplines (Álvarez et al., 2020; Barberá, 2020; Choi et al., 2020; Cinelli et al., 2021; Garimella et al., 2018; Greenstein & Zhu, 2012, Greenstein & Zhu, 2018; Kitchens, Johnson, and Gray, 2020; Terren and Borge-Bravo, 2021). However, the idea of competing and contested narratives about past events is highly familiar in history and memory studies, where cultural memory and historiography as such are construed as fluid and selective processes that are subject to media as well as societal and cultural factors (Halbwachs, 1992; Erll, 2011; Olick, 2011; Assmann, 2011; Garde-Hansen, 2011; Ricoeur, 2004). Conflicts, wars, and pivotal events are understood to be documented from perspectives shaped by nationalistic, cultural, or political frameworks (e.g., Leggewie and Lang, 2011).

Simple examples include differing names for events ("Fall" vs. "Conquest of Constantinople"; "Reconquista" vs. "Fall of Al-Andaluz") and battles with both sides using different field names, casualty numbers or even both sides claiming victory ("Schlacht bei Aspern" vs. "Bataille d'Essling"). For example, older French sources report 21,000 French casualties at the Battle of Aspern-Essling (1809), while Austrian sources state 30,000 French casualties. Similarly, for the Battle of Trafalgar (1805), French accounts cite 26,000 French vs. 18,500 English troops (ratio: 71%), while English sources claim 30,000 French vs. 17,000 English (ratio 54%).

Figure 1: Converging Narratives on Wikipedia: Battle Info Boxes



Notes: *Figure 1 contrasts the French and the German Wikipedia "battle info box" summaries for the Battle of Aspern-Essling. The left panel shows the information presented in 2007, the right side corresponds to 2024. The box at the bottom summarizes the respective casualty numbers provided by each side at both points in time. While the divergence in the loss difference is 8300 in 2007, this is reduced to 1500 in 2024 (due to an edit made in 2015). The divergence on French casualties is reduced from 9000 to 500.*

We refer to this phenomenon as "***divergence in 'truths',"*** because both accounts typically reflect the language-specific tradition of historiography and can be traced back to reputable language-specific encyclopaedic sources. Numerous other examples exist of different languages keeping diverging records of historic events of conflict, such as battles and other atrocities (e.g., Leggewie and Lang, 2011). but this divergence has, as of yet, not been studied quantitatively and at large scale. Yet, such biases, deeply embedded in historical accounts foster misunderstanding and could hinder both reconciliation and even economic development.

Continuing our motivating example of the battle of Aspern-Essling, it is noteworthy that Austrian history refers to it as "Schlacht bei Aspern" (e.g.: 'Lion of Aspern' in Vienna), while French history refers to it as "Bataille d'Essling" (e.g., 'Arc de Triomphe' in Paris). Initially, neither Wikipedia page mentioned the other name, but as of 2009, both do. Further, until 2007, "Bataille d'Essling" stated a French victory, while "Schlacht bei Aspern" stated an Austrian victory. Since 2007, "Bataille d'Essling" describes it as an "Austrian victory without strategic consequences" and the divergence in casualties has shrunk from 9000 to 500 (now: "Bataille d'Essling": 26,500, "Schlacht bei Aspern": 27,000).

While the existing literature has focused on the potential of digital platforms to amplify discord and increase conflict, the present research posits that connecting and juxtaposing conflicting historic narratives can foster a process of harmonization and contribute to a resolution and convergence of long-standing conflicting narratives. We thus depart from the existing "one-way paradigm" of amplified conflict to analyze whether digital platforms also offer novel opportunities to enable a mechanism of reconciliation. However, we stress that the hypothesized mechanism is fundamentally different from solutions suggested in the existing literature on echo chambers (e.g. Terren and Borge-Bravo, 2021).

To study how digital platforms may aid informational flow and reconciliation of "truths", we focus on the Wikipedia ecosystem. Wikipedia offers an ideal setting for our study, because every edit is publicly recorded, a feature used extensively in past research contributions (Halatchliyski et al. (2016); Kummer et al. (2021), Hinnosaar et al. (2022, 2023)). Moreover, Wikipedia provides a platform that allows for language-specific sub-platforms and thus divergence in narratives and provides us with a research context of clearly connected (by topic and language links) yet unambiguously separated (across languages) narratives, that is perfectly suited for developing the proposed methodology. In that sense, Wikipedia offers a unique laboratory to study and quantify both of the following phenomena: (i) the existence, transmission, and propagation of language-specific narratives as an instance of "historic echo chambers" and (ii) the extent to which digital knowledge platforms can (or fail to) induce a process of convergence in narratives or "converging truths." To tackle these two questions, we leverage recent advances in multilingual Large Language Models

(LLMs), which allow us to quantify and compare texts from different languages in a shared semantic space. We gather texts from over 100 past battles in over 50 languages from Wikipedia's multilingual revision histories to build a monthly panel of articles, and incorporate additional information on the battle outcomes, including party strengths, casualties, and languages spoken by each involved party. We can then explore whether digital platforms can facilitate a process of convergence for pre-existing language-specific differences in historical "truths," before extending the methodology to other domains.

LITERATURE REVIEW

The potential for bias in online platforms is well-documented. Prior literature highlights in-group bias across languages (Álvarez et al., 2020) and political slants away from neutrality in both expert-curated Encyclopaedia Britannica and crowd-sourced Wikipedia (Greenstein & Zhu, 2012, 2018). We conceptualize historic biases and diverging accounts of historical events between languages and nations as a "historic echo chamber effect," where cultural perspectives and preferences amplify differences in narratives.

The echo chamber effect has been extensively explored in social media, revealing mechanisms such as algorithmic filtering, self-selection, and homophily that reinforce pre-existing beliefs (Cinelli et al., 2021). Studies further show how social media induces partisan shifts, amplifies rumors, and reduces bipartisanship (Kitchens et al., 2020; Choi et al., 2020; Barberá, 2020; Garimella et al., 2018). These dynamics contribute to polarization and challenge democratic discourse, while mitigation strategies, such as diversifying information sources, have been explored (Terren & Borge-Bravo, 2021). These findings resonate with cornerstone insights from the disciplines of history and memory studies who highlight that memory, the recollection of facts, and historiography are fluid and selective processes that are subject to societal and cultural factors (Halbwachs, 1992; Erll, 2011; Olick, 2011.), which interact in complicated ways with the medium that creates the record. (Assmann, 2011; Garde-Hansen, 2011; Ricoeur, 2004).

Theories of truth—coherence, consensus, and pragmatic—offer frameworks for analyzing information convergence. Coherence theory emphasizes logical consistency within interconnected beliefs (Rescher, 1973; Kirkham, 2001), while consensus theory defines truth through rational agreement in a community (Habermas, 1984; Rescher, 1993). Pragmatic theory ties truth to practical consequences, viewing it as utility-driven and evolving with new information (James, 1907; Peirce, 1878). Together, these theories inform how data from multiple sources can merge into a unified narrative (Blanshard, 1939; Jenkins, 2006). The present research posits that connecting and juxtaposing conflicting historic narratives can foster a process of harmonization and contribute to a resolution of conflict. We argue that digital platforms offer a novel opportunity for bringing this mechanism to fruition. However, we stress that the hypothesized mechanism is fundamentally

different from solutions suggested in the existing literature on echo chambers (e.g. Terren and Borge-Bravo, 2021). Wikipedia offers a unique laboratory to study and quantify both phenomena: (i) the existence, transmission, and propagation of language-specific narratives as an instance of "historic echo chambers" and (ii) the extent to which digital knowledge platforms can (or fail to) induce a process of convergence in narratives or "converging truths."

The value of understanding this channel cannot be overemphasized, as historical grievances and differing interpretations of history have been sources of tension and conflict between ethnic groups and nations (see e.g., Galtung, 1996). In contrast, a shared vision of history can foster cooperation and diplomacy and promote stability and peace. (e.g., Larson & Shevchenko, 2010, Montville, 1990). This perspective has also been at the heart of famous efforts including the Gacaca courts in Ruanda (Clark, 2010; Minow, 1998) or truth and reconciliation commissions (Doxtader, 2009; Tutu, 1999; Wilson, 2001). Related research found that providing students with a cohesive understanding of their country's past promoted national unity, social cohesion, and a sense of identity (e.g., Noddings and Brooks, 2016), while shared historical narratives facilitate cultural exchange and tourism between nations. (e.g., Timothy & Boyd, 2006). We take this as overwhelming evidence for the value of promoting a shared narrative of historical accounts.

Wikipedia is the ideal laboratory to study whether digital knowledge platforms induce convergence in narratives, because Wikipedia constructs knowledge collaboratively, with consensus emerging through, coherence, and practical relevance (Kittur & Kraut, 2008; Jemielniak, 2014; Reagle, 2010). Talk pages foster rational dialogue, aligning with consensus theory (Habermas, 1984), while coherence is ensured by integrating new information with existing content (Forte & Bruckman, 2008). Pragmatic utility guides contributions to prioritize relevance (Fallis, 2008). Mechanisms like Creative Commons licenses (Lih, 2009), verifiability policies (Magnus, 2009), and collective editing on talk pages balance differing viewpoints to create unified, unbiased narratives (Forte et al., 2012; Reagle, 2010).

Moreover, Wikipedia's multilingual structure uniquely reflects diverse cultural perspectives, with independently operated language editions often differing in scope and historical interpretations (Hecht & Gergle, 2010). Tools like interlanguage links, multilingual editors, and translation aids, supported by Wikidata, help bridge gaps, streamline updates, and enhance synchronization while highlighting cultural differences (Adafre & de Rijke, 2006; Miquel-Ribé & Laniado, 2018; Vrandecic, 2012; Samoilenko et al., 2017). On the other hand, research has shown that the process is subject to challenges, including resource disparities, cultural biases, and misinformation. Smaller editions lag in updates (Lewoniewski et al., 2017), while edits may reflect national pride or differing sources (Pfeil et al., 2006). Wikipedia's inclusivity also invites unreliable edits, requiring strong editorial practices to ensure verifiability and consistency (Ford et al., 2013; Luyt & Tan, 2010).

Language links aid synchronization and information transfer, but full convergence remains limited by resource and contextual differences (Adar et al., 2009; Callahan & Herring, 2011; Liu et al., 2018).

To summarize, our work builds on and contributes to four connected streams of literature:

- **Digital Platforms, Echo Chambers and Bias**: Studies highlight the prevalence of biases in digital platforms like social media platforms, and Wikipedia, particularly in early article versions, with later revisions often reflecting more balanced perspectives (Greenstein & Zhu, 2018; Kim et al., 2016). However, systemic biases remain, shaped by editorial practices and cultural perspectives.
- **Historiography and Narrative Bias**: Historical accounts, especially those documenting conflicts, frequently reflect the perspectives of victors or dominant cultural groups (Bloch, 1954; Carr, 1961; Tosh, 2015). This reinforces inequalities in representation and complicates reconciliation efforts (Halbwachs, 1992; Hunt, 1989; Minow, 1998; Tutu, 1999).
- **Peace-building and Shared Histories**: Research which underscores the importance of establishing shared historical narratives in fostering reconciliation and empathy (Bekerman & Zembylas, 2011; Galtung, 1996). Digital platforms represent a novel avenue for promoting such shared understandings, and thus contribute to peace building and collaboration between language groups and societies (Clark, 2010; Minow, 1998; Tutu, 1999; Zehr, 2002).
- **Recent progress in Large Language Modelling** has fueled AI-supported translation and Natural Language Processing methods. Major progress in NLP was achieved by several algorithms, including word2vec (Mikolov et al., 2013), BERT (Devlin et al., 2018), and, most recently, Multilingual Embeddings (Cohere, 2024). Cohere, founded by former Google Brain employees in 2019, focuses on LLMs enabling access to researchers through its nonprofit research lab and open-source community, Cohere for AI.

We identify three key stylized facts that will guide our analysis. First, conflicting parties report events differently, often highlighting their own achievements while downplaying those of their adversaries. Second, the winning side tends to provide more detailed accounts, leveraging its position to shape the narrative—though this is less evident in cases involving crimes against humanity. Finally, discrepancies in accounts, including numerical data such as casualty figures, are common, reflecting the subjective nature of historical reporting. These discrepancies often stem from biases, unequal access to information, or differing political and cultural agendas. For example, casualty figures may vary depending on whether reported by the victor, aiming to downplay losses, or the defeated, highlighting the damage inflicted.

Before moving to the general strategy of the analysis, we stress that the example from the Battle of Aspern-Essling is not unique. A similar pattern (though under opposite signs) emerges for the Battle of Wagram which ensued only 6 weeks later. Other prominent examples are the "Fall"/"Conquest" of Constantinople and of Al Andaluz/Spain, or the Battle of Trafalgar and many more. Our main goal is to analyze these patterns systematically, while our second major goal is to shed light on the relative positioning of "winning", "losing" and uninvolved languages.

Figure 2: "Conquest"/"Fall" of Constantinople in Turkish, Greek/Romanian, and English 'Wikipedia eyes'



Notes: *Main images chosen for the "Conquest" (Turkish) or "Fall (Greek/Romanian/English) of Constantinople (1453) in different Wikipedias. While Turkish Wikipedia focuses on the emperor, the image on Greek Wikipedia emphasizes the fortifications and bloodiness of the battle and emphasizes religion. English Wikipedia shows an image that emphasizes the strategic situation and important details of the naval battle.*

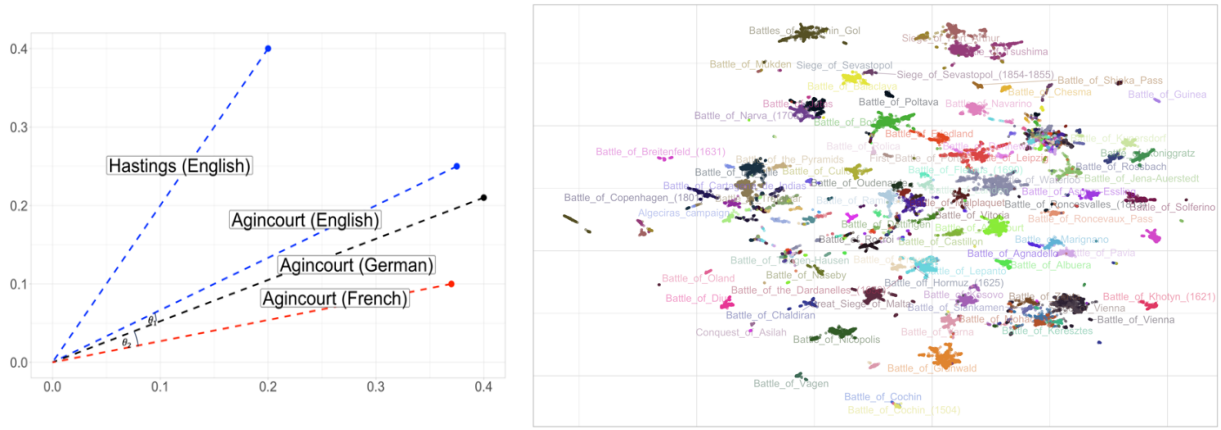
RESEARCH STRATEGY

Main Questions: Our analysis is shaped by three guiding questions, emerging from the examples presented above. First, do biases in reporting grow or diminish over time? Understanding this dynamic can reveal whether historical narratives tend to stabilize or diverge over time. Second, does interlinking on platforms foster a "convergence of truths," where conflicting perspectives align through shared access and connectivity? Third, do these biases, deeply rooted in cultural and historical contexts, persist and transmit through centuries, shaping how societies remember and interpret the past? In addition, we ask what the role of "neutral" languages is in bridging these divides and mediating between conflicting narratives?

Multi-Lingual LLM Embeddings and Cosine Similarity: The core idea is to analyze how the semantic differences in historical narratives across languages evolves over time. To achieve this goal, we leverage multilingual large language model (LLM) embeddings. The core idea is to embed all language versions of an article into one multilingual semantic space and to quantify the alignment or divergence of narratives. The quantification is achieved by evaluating their cosine similarity, a measure of how closely the embedding vectors for two texts (i.e., t_1 and t_2) are aligned in a high-dimensional space. Cosine similarity ranges from 1 (exactly identical) to -1 (exactly opposite) and is equal to one minus the cosine distance.

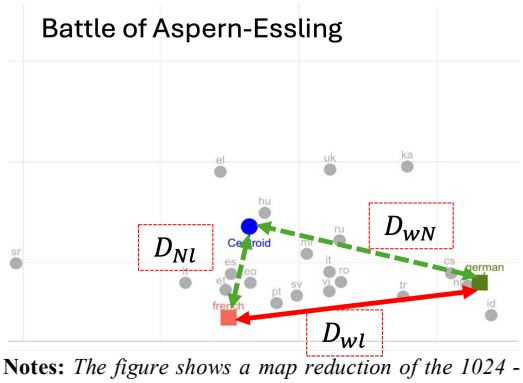
$$S_{cos}(t_1, t_2) = \frac{\sum_{i=1}^{n} t_{1i} t_{2i}}{\sqrt{\sum_{i=1}^{n} t_{1i}^2} \sqrt{\sum_{i=1}^{n} t_{2i}^2}}$$

Figure 3: Illustration of Cosine Similarity and Mapping of Crosslingual Similarity Scores



This is illustrated in the left panel of Figure 4 with a visual representation based on two historical battles: the Battle of Hastings and the Battle of Agincourt. The figure contrasts semantic cross-language similarity in narratives with same-language semantic distance between two topics (e.g. different battles). In this illustration, the English account of the Battle of *Hastings* shows significant semantic distance to all narratives about the Battle of *Agincourt*, as reflected in its isolated position in the figure. Among the more closely aligned multi-lingual accounts of the Battle of Agincourt the English and German versions show closer alignment, while the French version diverges more significantly. This illustrates how the proposed methodology will allow to quantify how historical events are framed differently across languages, shaped by cultural or linguistic context. For example, the French narrative of Agincourt may emphasize aspects that differ from those in the German or English accounts, reflecting the perspectives of each nation's historiography.

Figure 4: Variable 1: Distance of Narratives.



Notes: *The figure shows a map reduction of the 1024 - dimensional cosine similarity measures for Battle of Aspern-Essling in 2023. The squares show the language versions of the conflicting parties (German and French), the grey dots represent articles in neutral languages. The "Centroid" is the average location of neutral languages. The three arrows* D_{wl}, D_{wN}, D_{Nl} *indicate the distance measures that will be used as main outcome variables.*

Initial Cross-Language Comparison via LLMs (2023 version): The right panel of Figure 4 presents the results of embedding historical battles into multi-lingual large language model (LLM) by Cohere and visualizing the relationships using dimensionality reduction techniques. The analysis explores how historical narratives differ across languages, with each color representing a specific battle (e.g., Battle of Agincourt, Battle of Vienna) and each point representing the text of a corresponding Wikipedia article in a different language (German, French, English, Swahili, etc.) as

embedded by Cohere's multilingual LLM. The spatial distribution shows that articles on the same battle (but from different languages) are strongly clustered, which highlights that the embedding vectors encode meaningful semantic differences. Tightly grouped battles suggest a shared or convergent narrative across linguistic boundaries, while dispersed points highlight divergence in how different languages frame the same event. Some battles, like Battle of Vienna, appear closely aligned across languages, reflecting greater semantic agreement. Others, like Battle of Tsushima or the Siege of Port Arthur, are more dispersed, indicating greater variation in how the event is described in different languages.

Figure 4 shows the cluster corresponding to the Battle of Aspern-Essling. A triangular distance measure between averages of winners, losers, and neutral languages provides key metric for analysing narrative (static) *distance*. The red-blue triangle shows the quantification of the three distances between winning party, losing party and the average uninvolved language. We propose to use this measure to capture whether articles are converging over time – a shrinking triangle would indicate that articles from uninvolved, losing, and winning parties become more similar over time.

Before estimating this triangular convergence measure, we start with a simple cosine similarity between winners (w) and losers (l) at the article level (D_{Awlt}), which we track over time. We estimate linear and flexible time-trends, to evaluate whether divergences grow or diminish over time. We plan to contrast the battle specific similarities with analogous similarity measures for articles about neutral baseline topics, such as planets and chemical elements. A simple specification to estimate group specific linear trends at the topic level A is given by:

$$D_{Awlt} = \beta_0 + \beta_1 Battle_{Awl} + \beta_2 t + \beta_3 Battle_{Awl} \times t + \varepsilon_{At} \quad (1)$$

D_{Awlt} is the measure of narrative distance between the article A in language of the winner w and the loser l. $Battle_{Awl}$ takes the value 1 if the article is about a battle and if it was fought between w and l, and t is a simple numeric variable representing time. Suitable non-parametric alternatives will be estimated using time as a categorial variable and estimating non-parametric methods including splines or kernel smoothing.

The analysis of D_{Awlt} will be underpinned by an analysis of the alternative distance measures that are given by the differences in article properties (e.g. length) and casualty reports (compared against numbers in the info boxes of the control group). In addition to the control groups based on planets and chemical elements, we can also use battle pages for the same languages but about other battles that did not involve w and l as control groups.

In step 2, we can leverage the availability of languages of uninvolved ("neutral") parties. Using the centroid (i.e. average) of the measure obtained for these languages, we repeat the analysis but contrast the winner's distance to the centroid, D_{AwNt}, to the loser's distance to the centroid, D_{AlNt}.

Doing so allows us to analyze initial distances and convergence between each party and the position of the average uninvolved party. Moreover, the contrast of the convergence between neutral and winner versus convergence between neutral and loser will shed additional light on the convergence results (or absence thereof) in D_{Awlt}. A third aggregate measure is given by the circumference of the triangle $D_{Awlt} + D_{AwNt} + D_{AlNt}$, which can be conveniently tracked over time.

Using multi-lingual LLM embeddings, like those from Cohere, combined with Wikipedia's revision history to track changes at a monthly level, represents a novel method for studying narrative convergence and divergence across languages over time. By moving beyond surface-level differences in wording to capture deeper shifts in meaning and the evolving convergence of "truths," this method provides unparalleled insights into how digital platforms can foster convergence in long-standing, diverging historical narratives. Unlike the familiar discord and divergence often seen on social media, this approach explores how such platforms can constructively reshape collective understandings of the past and how language, culture, and economic factors influence this process.

Analysis of baseline convergence and platform interventions: We further analyze knowledge conflicts across languages as well as the efficacy of platform strategies such as cross-language linking or AI-driven translation tools in fostering narrative convergence. As before, we first focus on the analysis of historic accounts of the same battle across different languages and whether initial discrepancies in the different accounts on Wikipedia diminish over time (or do not) because of, among other things, exposure to the other language versions about the same battle.

We will leverage three major design features that were introduced by the platform since 2002 and test how (i) interlinking, (ii) user interface standardization (battle info boxes), and (iii) AI-driven tools for translation affect biases and foster alignment. For all three interventions, this part will deploy event study techniques, and thus a short description of the interventions is in order.

The first intervention, "battle info boxes," are standardized templates summarizing battles in the top right corner of articles (e.g., "Victoire francaise" vs. "Austrian victory" or casualty discrepancies). These templates make stark discrepancies more visible, potentially triggering convergence. Battle info boxes are user-added templates, enabling two event study moments: (1) Template Availability: When the template becomes available across a language, allowing estimation of an ITT (Intention to Treat) effect. (2) Template Activation: A battle info box is added to a specific article, allowing ATT (Average Treatment Effect on the Treated) estimation. Identification arises as different language versions (e.g., Dutch or Swedish) activate the battle info box at different times, enabling comparisons across language-battle observations. In addition, uncontested content, such as geometric shapes or celestial bodies, serves as alternative comparison group.

The second intervention, language links, connects articles to their counterparts in other languages, making cross-language versions one click away. Like battle info boxes, language links

make discrepancies more visible, triggering convergence. This intervention also supports ITT estimation (when the feature is generally available) and ATT estimation (when links are added). As before, control observations are provided by other language versions and by control topics. The third intervention, AI translations, is rolled out universally across the platform, making treatment exogenous. However, unlike the first two interventions, other languages cannot serve as controls. Instead, we use uncontested content, such as geometric shapes or celestial bodies, as a comparison group to control for baseline differences between languages. This comparison group is also applicable to the first two interventions.

Estimation strategy: Each intervention will be studied independently, but the estimation methodology is consistent across them. We define "TreatAct" as an indicator variable equal to 1 when a feature (e.g., battle info box) is added to an article or article pair, signifying active treatment. While features can be removed, they are typically improved over time. We hypothesize that treatment accelerates convergence, with the null hypothesis being no effect on narrative convergence. To test this, we estimate a DiD regression using article-level outcomes (e.g., article length, cosine similarity, or – for battles – info-box numbers). For simplicity, we focus on battle analysis using planets as a control group and language links as treatment. The methodology generalizes to other topics like historical figures. Our sample includes all battle-language and planet-language observations, forming a topic-language-month panel where an article A in language L is the observational unit. For article length, we construct a balanced monthly panel, coding non-existent articles as zero. For other outcomes, the panel is balanced for existing articles (cosine similarity) or reported casualties (casualty differentials). The simplest two-way fixed-effects:

$$\Delta D_{Awlt} = \beta_0 + \beta_1 TreatActive_{Awlt} \times Battle_{Awl} + \delta_A + \gamma_t + \varepsilon_{ALt} \quad (2)$$

This regression can be expanded to a full event study design with period-battle interactions and time normalized to 0 at the time when treatment between *w* and *l* is activated. D_{Awlt} is the measure of narrative distance between the winner's (*w*) and the loser's (*l*) version of article *A* and the centroid of the neutral languages in month *t*. For the event-study analysis we focus on the change in narrative divergence by taking the first difference to the month before. $TreatActive_{ALt}$ equals 1 if and only if treatment is active for the given article *A* at time t (e.g. if an info box is present or a language link is in place) and, otherwise, equals 0. $Battle_{Awl}$ is an indicator variable taking the value 1 if the article is a battle-observation involving languages *w* and *l* as conflicting parties, and 0 if it is from the control group. Article δ_A and weekly fixed effects γ_t are also included. As before, we can expand to compare the winner's and loser's distance to the neutral language centroid (N_A) for article *A* (D_{AwN_At} *vs* D_{AlN_At}) and include article-language level fixed effects α_{AL} for the winner and the loser respectively (not shown here).

These analyses contribute to the literature on information systems, because each intervention uses a different channel to make conflicting information more salient. Language links improve access to other versions, battle info boxes highlight key contradictions (e.g., in casualties), and AI translations expand the pool of users accessing foreign-language information. Comparing how the effects differ for each channel offers insights into their relative importance in exposing conflicting narratives.

Specification issues: A key challenge in our analysis is the potential endogeneity of language links and battle info boxes, as they are placed by users. First, reverse causality may occur if a link is added due to prior convergence. To address this, we use changes in narrative divergence as the dependent variable and run regressions with increasingly small time windows around the event, ensuring convergence is not misattributed. Pre-trend analyzes further confirm that convergence does not precede the intervention. Edits co-occurring with the placement of links will be studied in detail to understand the role of individual editors versus the crowd.

Second, the user who places a link may also edit articles, systematically reducing divergence. To account for this, we: (1) exclude convergence attributable to the user who placed the link, (2) exclude edits occurring within a week or month of the intervention and focus on changes in convergence speed beyond this window, and (3) conduct "honest DiD" analyzes of pre-trends (Rambachan and Roth, 2019, 2023).

Lastly, battle info boxes and language links are introduced at varying times, creating heterogeneous treatment timing. Standard Two-Way Fixed Effects (TWFE) estimators may produce biased results due to improper weighting of treatment effects, particularly when earlier-treated groups act as controls. To address this, we apply the Callaway and Sant'Anna (2021) methodology, which adjusts for heterogeneous treatment effects and avoids bias from incorrect control group weighting.

DATA AVAILABILITY AND PROOF OF CONCEPT (PILOT STUDY)

Wikipedia provides all the necessary information in the revision histrory and, more usefully for scraping, in an API which is publicly documented. The API allows to pull the raw text and metadata on every revision through a unique language specific revision ID. From the raw text it is also possible to obtain historical information on the presence of langauge links. Articles about the same topic (e.g. battle) are connected through a unique wikidata ID, connecting all language versions. (e.g. https://www.wikidata.org/wiki/Q541626 for the battle of Aspern-Essling). Page views and User profiles are publicly provided by wikimedia's x-tools (https://xtools.wmcloud.org/) which also ensures public data availability through an API.

Since data collection is at the core of this proposal we ran two pilot data collections, with data on 46 and 93 battles in up to 12 languages. We focus on pre-WWI battles to avoid complexities of modern warfare (month-long battles, complex equipment) and to minimize potential survivorship bias. A focus on human-related casualties ensured a manageable scope, consistent data and kept the

dataset aligned with the goals of the pilot. Our larger sample includes over 200,000 Wikipedia articles from 85 languages on 93 historic conflicts spanning the 14th to 20th centuries, edited by 56,000 editors. Among them, 1,721 edited in more than one language, with 76% contributing to English Wikipedia, underscoring its role as an informational aggregation language (Kim et al., 2016). German Wikipedia was second, with 479 multilingual editors (28%).For each conflict, we identified the involved countries and their languages, finding 33 distinct nations, most prominently France. While English Wikipedia covered all 93 conflicts, many other languages had far fewer articles, highlighting the significant barrier language poses to information flow. Future expansions will address underrepresented regions to improve coverage.

To prepare articles for embedding, the articles are cleaned and batched before being run against the Cohere v1/embed-jobs API endpoint (Team et al., 2023). This process enables clustering, distance, and topic analysis to assess text (dis)similarity. Cosine similarity is applied to derive topic- and language-level averages, identifying the most divergent topics and languages. We also plan to use cluster analysis techniques like t-SNE and Uniform Manifold Approximation and Projection (umap) for visualizing variation (Maaten & Hinton, 2008) and apply language fixed effects to account for differences in information density.

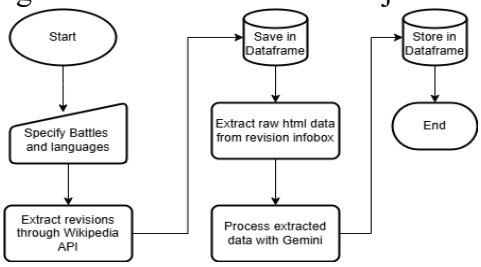


Figure 5: Prototype version of the Page Content Scraper.

To automate Wikipedia scraping for troop strength and casualty numbers, we use a three-step approach. First, yearly page revisions are collected via the Wikipedia API and saved in a structured data frame. Next, an HTML parser extracts and cleans infobox content for processing. Finally, we deploy the ChatGPT 4o-mini to interpret and translate infoboxes. In a post-processing step we handle inconsistencies across Wikipedia language versions to organize data into structured categories (e.g., troop casualties) with ranges recorded as bounds and thus obtain a harmonize dataset.

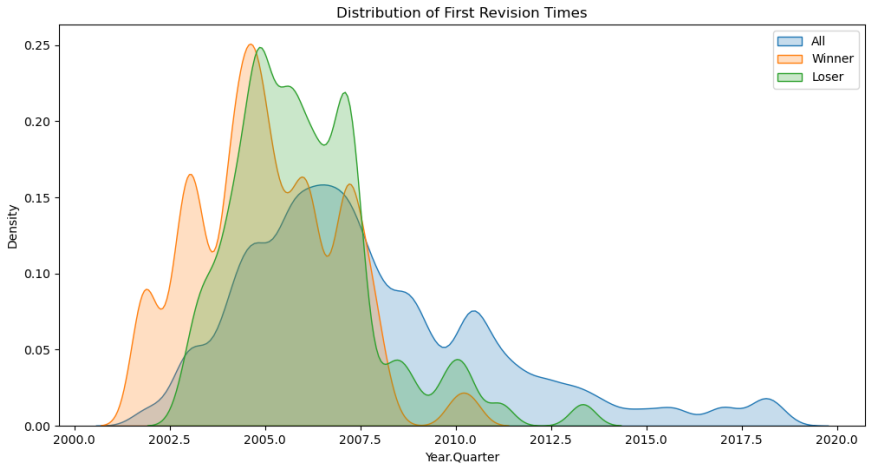
We deployed the Gemini AI for cross validation. We combined the Gemini AI (flash 1.5 with Google's translation API, and organized data into structured categories (e.g., troop casualties) with ranges recorded as bounds. AI results are validated against manual checks, with discrepancies resolved through Python postprocessing, ensuring accurate, scalable data collection across languages.

The underlying sample includes up to **104 historical battles**, out of which **54 (52%) feature full coverage** of both at least one winner and one loser. On average, each battle includes **94.4 observations for winners**, **80.6 for losers**, and **approximately 777 uninvolved parties**. API information, battle info boxes and cosine similarity measures were combined in the prototype analysis of battle clusters and narrative distances above. Casualty convergence across a second dataset of 46 battles is also available for analysis.

DATA AND FIRST DESCRIPTIVE ANALYSIS

Figure 6 visualizes the distribution of first revision timestamps for Wikipedia articles corresponding to entities involved in historical battles, disaggregated by winner status. The x-axis reflects the year in which an article first appeared on Wikipedia, while the y-axis shows the estimated density based on a kernel density estimator (KDE). We distinguish between three groups: all languages (baseline), languages of losing parties (green, solid line), and languages of winning parties (red, dashed line). The KDEs are computed using only those entities whose winner and loser status could be clearly identified.

Figure 6: Distribution of First Wikipedia Revisions by Winner Status



The underlying sample includes a total of 104 historical battles, out of which 54 (52%) feature full coverage of both at least one winner and one loser. On average, each battle includes 94.4 observations from winners, and 80.6 from losers, and approximately 777 from uninvolved parties. The figure reveals systematic differences in the timing of coverage. Battles associated with winning parties tend to be documented in corresponding languages earlier, while languages associated with losing parties show a rightward-shifted density, suggesting delayed availability of the historical record in its language-specific Wikipedia. This pattern is consistent with the hypothesis that online collaborative platforms may replicate offline visibility asymmetries.

This is corroborated by Figure 7, which shows a boxplot that shows content availability, proxied by revision size in bytes, for battle events on Wikipedia and contrasts articles in languages of winning parties to languages of losing or neutral parties. It is easy to see that much more content is provided in winning language than in other language, potentially reflecting differences in the importance of the event for the societal narratives.

Figure 7: Language-Specific Article Length by Battle Outcome

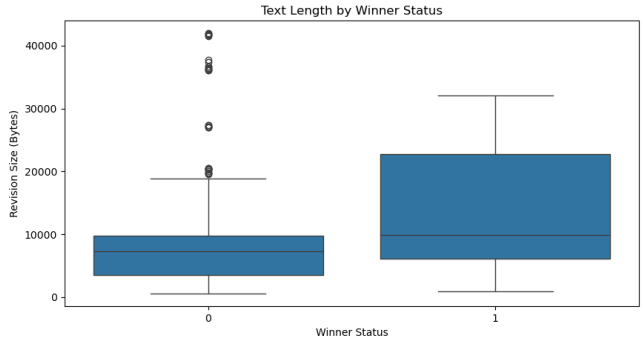


Table 1 provides descriptive regressions of how involvement in historical battles influences the timing and magnitude of Wikipedia coverage across three key outcome variables: article creation year, presence in structured battle info boxes, and revision size. Each model controls for three binary indicators: whether the entity was a winner, a loser, or was involved in the battle in any capacity.

Table 1: Impact of Conflict Involvement on Wikipedia Coverage

	(1)	(2)	(3)
	First Revision Year	**First Battlebox Year**	**Revision Size (bytes)**
Winner	-2.137 ***	-1.030 ***	11108.94 ***
	(0.124)	(0.127)	(340.57)
Loser	-1.127 ***	-0.357 ***	7126.79 ***
	(0.125)	(0.128)	(375.13)
Party involved	-1.089 ***	-0.752 ***	—
	(0.118)	(0.123)	—
Constant	2008.25 ***	2008.31 ***	11920.00 ***
	(0.014)	(0.015)	(155.16)
Observations	64,328	34,774	22,710
Adj. R-squared	0.103	0.065	0.052

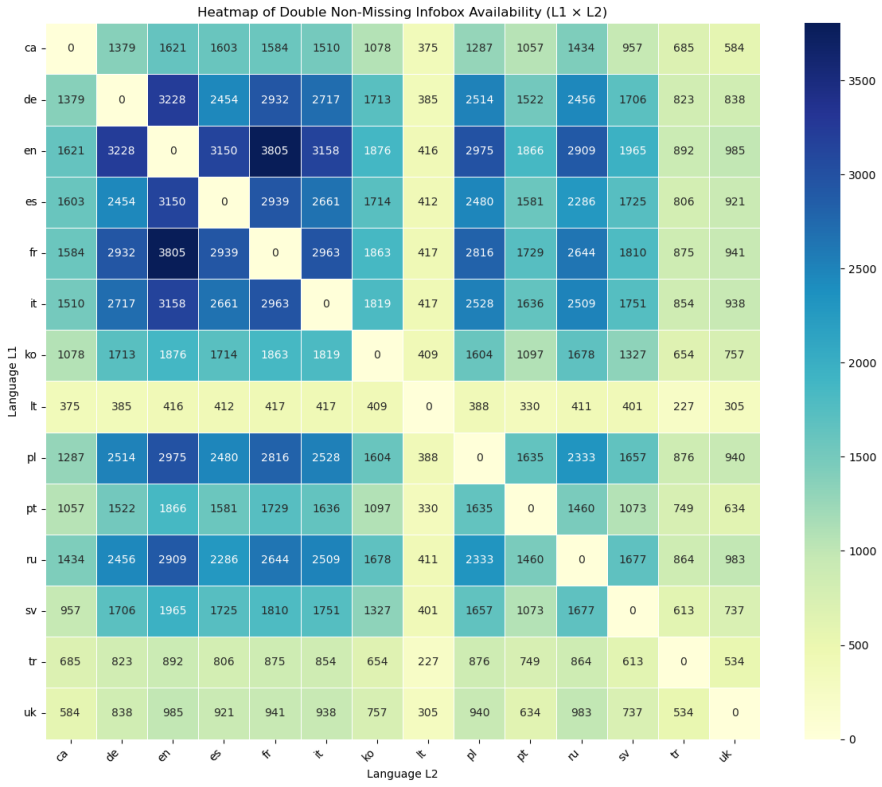
Notes: OLS estimates shown with robust standard errors in parentheses. The omitted category consists of uninvolved entities. The dependent variables are: (1) year of first article revision, (2) year of first battle info box appearance, and (3) initial article revision size in bytes. The third model is restricted to articles from battles with complete winner–loser coverage. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Across all three specifications, involvement in a conflict is associated with systematically earlier and more substantive Wikipedia coverage. Winners are, on average, featured over two years earlier than uninvolved actors (Col. 1) and added to battle info boxes roughly one year earlier (Col. 2). Losers are also significantly earlier, though to a lesser extent. In terms of content size, initial revisions for winners are approximately 11,109 bytes longer than for uninvolved entities, while loser-related articles are 7,127 bytes longer (Col. 3). These findings highlight the skewed documentation and visibility of conflict participants—particularly winners—in crowd-sourced historical records.

Together, these results offer preliminary descriptive evidence that Wikipedia's documentation patterns reflect asymmetries in outcome salience, with broader implications for digital memory, historical framing, and knowledge representation in collaborative systems, raising implications for digital historical memory and algorithmic knowledge curation.

Differences in Narratives, Numbers and Accounts: We now turn to the quantification of diverging accounts and potential convergence in truths, based on the textual analysis of the articles, and of battle info boxes like the one shown in Figure 1, specifically. For our pilot analysis we use a sample, which includes 99,400 pairwise language observations, drawn from 12 unique languages and spanning up to 104 distinct historical battles (80 after cleaning the data). Figure 8 illustrates the interlanguage availability of structured Battle Box information. Each cell counts how many observations have battle info boxes available for each langauge pair.

Figure 8: Interlanguage distribution of battle info box availability



The dependent variables are (1) the absolute difference in reported casualty differentials between the warring parties and (2) the absolute difference in reported troop strength differences (i.e., superiority between two language editions. The first difference derives from comparing casualties of side A to side B in a battle, and our measure then tracks, how these differences differ across languages.

Note that the first number – the difference in casualties or troop strength – is naturally non-zero, as it is normally the case that one side fields more troops or loses more troops than the other. However, if both sides agree what happened, then they would both report the same difference in casualties or troop strength, whereas if there is historic (or ongoing) disagreement, then this would be reflected in a non-zero difference in the differences of casualties or troop-strength. For example, if both sides

agree that 30,000 soldiers on one side fought 20,000 on the other side, both would report 10,000 as the troop strength difference, such that the difference of these differences is 0 (i.e., no discrepancy). However, if – as in the Battle of Trafalgar – one side reports 30,000 vs. 17,000 , while the other side reports 26,000 vs. 18,500, then the difference in these differences would be 5,500, which is the primary numeric measure of narrative disagreement. Figure 9 shows the distribution of the difference in reported casualty differences in our sample. Figure 10 shows how these differences evolve over time and contrasts it against the placement of a bidirectional language link between the two languages in the observation.

Figure 9: Distribution of Differences in Casualties Differentials

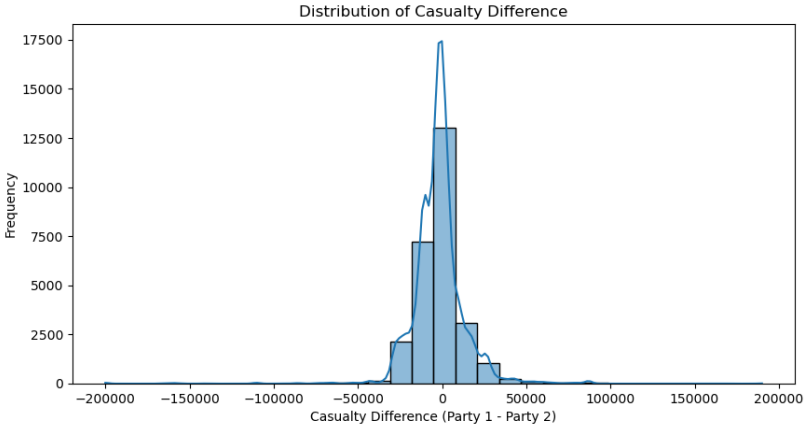


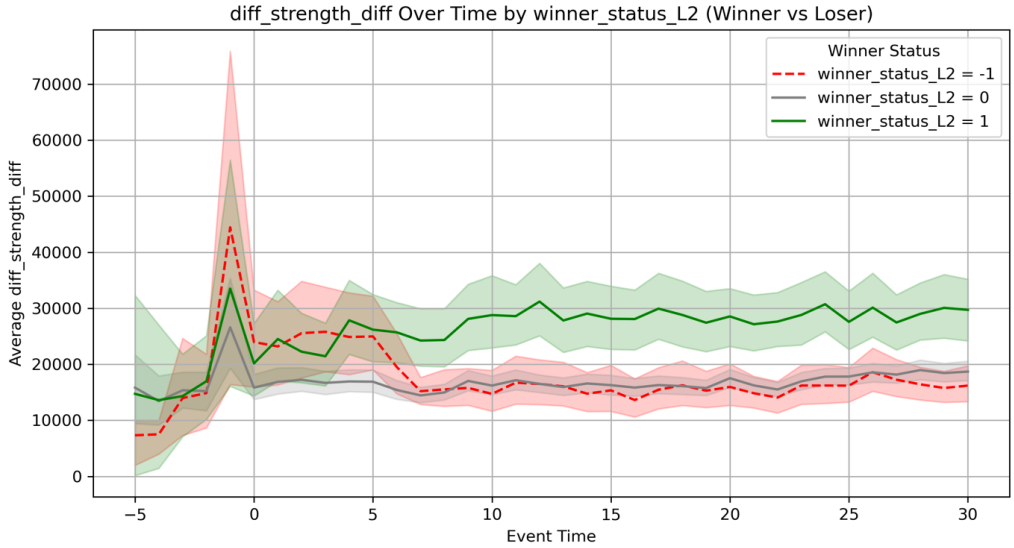
Figure 10 illustrates the evolution of differences in reported military outcomes between language versions of Wikipedia articles on the same historical event. The horizontal axis captures event time in quarters, centered at t = 0, the period in which a hyperlink is first established between the two language articles. Negative values represent pre-link periods, and positive values reflect post-link periods. The y-axis denotes the average value of the respective dependent variable. In the winner/loser stratification plots, line types and colors distinguish the narratives associated with the winner (green solid lines) and the loser (red dashed lines) of each conflict.

The analysis leverages structured info box data, which reflects factual claims made within Wikipedia's semi-structured data schema and facilitates visual inference into whether narrative alignment improves after cross-lingual exposure (via hyperlinking). The figure shows a pronounced spike in divergence shortly before language links are introduced, and this spike disappears after the introduction. However, the plot shows that winner/loser perspectives exhibit different temporal trajectories.

Estimation Strategy and Model Specification: We estimate a high-dimensional fixed effects model to examine how cross-lingual hyperlinking on Wikipedia influences the absolute difference in reported casualty figures between language editions of the same conflict. The dependent variable, *Absolute Casualty Difference*, captures the magnitude of disagreement between two language editions

for a given battle and time period, measured in quarters. To identify the temporal dynamics surrounding hyperlink formation, we include a set of event-time indicators ranging from several quarters before to multiple years after the link is established.

Figure 10: Evolution of Cross-Language Narrative Divergence Over Time



The specification includes fixed effects for the battle (i.e., Wikidata event ID) and for both source and target languages. These controls account for persistent unobserved heterogeneity across conflicts and language communities. Formally, the regression equation is given by:

$$AbsTroopStrengthDiff_{(i,t)} = \sum_{k(\neq -1)} \beta_k \cdot EventTime_{(t=k)} + \alpha_i + \lambda_l + \gamma_t + \varepsilon_{(i,t)}$$

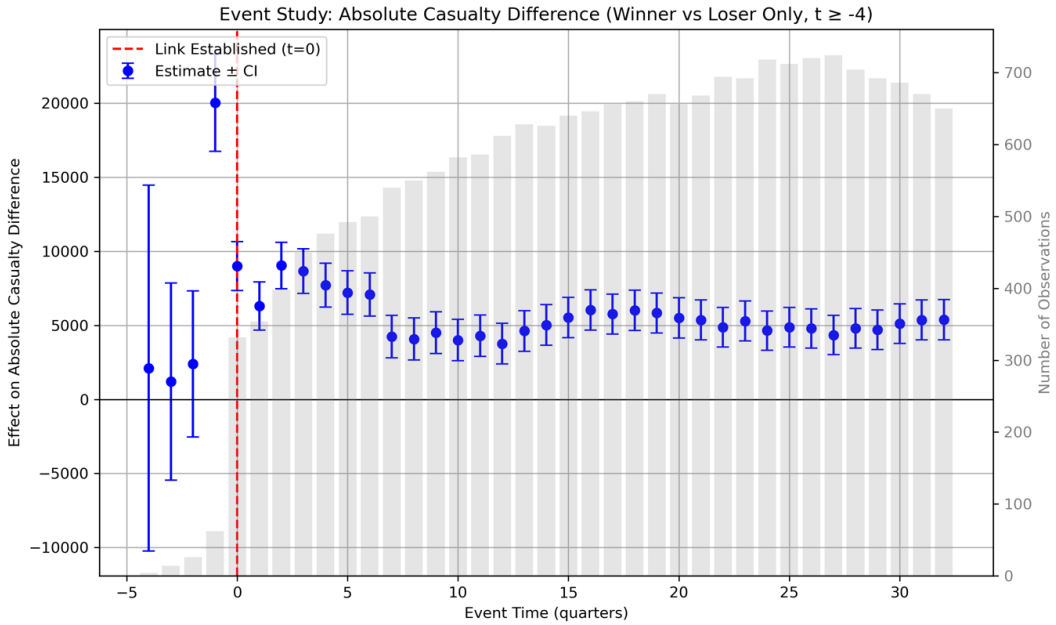
where α_i. are battle fixed effects, λ_l and γ_t are fixed effects for the two language editions in the dyad, and ε is an idiosyncratic error term. The model is estimated via OLS with heteroskedasticity-robust standard errors.

Results and Interpretation: The results are presented in the Figure 11 below and in Appendix Table 1. The figure presents an event study of the absolute difference in reported battle casualties across language editions of Wikipedia, from losing and winning parties. The outcome variable is the absolute difference in reported casualties between two language versions of the same battle article. The x-axis denotes event time in quarters relative to the establishment of a direct interlanguage link between the articles (t = 0), while the y-axis shows the estimated effect on the outcome variable.

Coefficient estimates from a two-way fixed effects regression (battle and language-pair fixed effects) are plotted as blue points with 95% confidence intervals. Estimates for quarters $t \in [-4, 32]$ are shown, with the omitted category being $t = -1$. The secondary axis (right y-axis) displays the number of observations used to estimate each event-time coefficient. The distribution of these counts reveals a sharp growth in coverage around the linking event, with substantially fewer observations available

for pre-treatment periods. This imbalance informs the widening of confidence intervals prior to t = 0 and cautions against overinterpreting pre-trend dynamics.

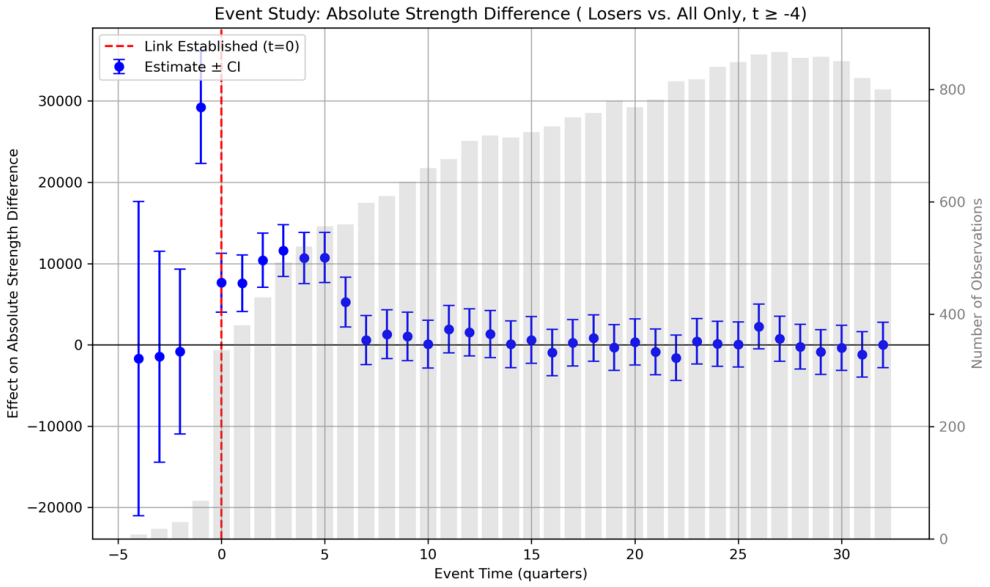
Figure 11: Differences in reported Casualties-Difference across languages. Losing parties compared against all others.



Troop-Strength-Differentials: Turning from casualty figures to troop strength, we now examine whether Wikipedia's interlingual narratives diverge in their depiction of military capacity. In particular, we analyze the evolution of interlanguage differences in reported troop strength differentials — that is, the extent to which two language editions differ in how asymmetrically they report the number of troops involved in the same battle. As before, we restrict attention to battles where a structural link is established between a "winner" and "loser" page at time t=0t = 0t=0, and estimate dynamic treatment effects using the same event-study specification.

The results are shown in Figure 12. In contrast to the patterns observed for casualty differences, interlingual differences in reported troop strength exhibit a sharp and sustained increase following the creation of the interlanguage link. At t=0t = 0t=0, the divergence jumps significantly — implying that, upon linkage, editors in each language appear to revise their pages in a manner that *amplifies*, rather than reconciles, asymmetries in reported force sizes. The effect persists across subsequent quarters and remains well above pre-linkage levels, suggesting that exposure to alternative narrative framings may activate editorial responses that reinforce, rather than resolve, contested representations of military strength.

Figure 12: Troop Strength Differentials of losing party vs. winners and neutral parties



The pattern of convergence is even more pronounced. They reveal a sharp and persistent reduction in the absolute casualty difference following link formation. There is no evidence of pre-trends: event-time coefficients for the quarters leading up to the link are statistically indistinguishable from zero. Starting at $t = 0$ the establishment of the hyperlink coincides with a large increase of observations, i.e. instances in which both languages under comparison report numbers. The difference is strongly positive and remains statistically significant for over a year but declines in the quarters thereafter. The largest increase is observed one period before the link is introduced suggesting that pre-linking differences were even more substantial. Still, the coefficients for $t = 0$ through $t = 5$ remain in the range of 10,000 and are significant at the 1% level. This is followed by a decrease to small and insignificant numbers, consistent with a meaningful decrease in reporting discrepancies.

These findings support the hypothesis that hyperlinking across language editions facilitates informational alignment. By connecting articles on the same conflict, editors may become exposed to alternative narratives or data sources, triggering revisions that reduce interlingual inconsistency. The inclusion of detailed fixed effects ensures that these effects are not confounded by differences in editing intensity or conflict characteristics, strengthening the causal interpretation of the link between hyperlink creation and convergence in reported information.

Alternative Specifications: In the Appendix we provide alternative results. First, we show the same analysis when constraining that the differences between two languages are computed from both warring parties (i.e., both the winner and the loser). The results are similar and more pronounced, but with lower significance due to the smaller number of observations.

Outlook and future work: The first descriptive analysis, while promising, highlights limitations of data availability on battle figures. In our main analysis we will deploy our measure of cosine similarity

and complement the analysis of open contradictions in numbers with our universally available analysis of narrative convergence or divergence over time. In addition, we will expand the present analysis with more rigorous causal approaches and use pages on celestial bodies as an alternative control group to complement articles about the same topic in uninvolved languages

---

DISCUSSION

The results suggest a clear and statistically significant reduction in the absolute difference in casualty reports following link establishment. We find no evidence of anticipatory adjustments in the pre-periods, but beginning at t=0 (i.e., the quarter of hyperlink creation), the treatment coefficients become sharply positive and remain significantly elevated thereafter. The largest discrete increase occurs at period -1, likely reflecting revision activity immediately preceding link formation. Subsequent coefficients, including those for t=0 through t=32, are consistently significant at the 1% level and hover around 4,000–6,000 casualties in magnitude, reflecting a sustained narrowing of disagreement between language editions post-link.

Overall, these findings are consistent with the interpretation that cross-language hyperlinking reduces divergence in conflict-related information across communities. By connecting articles in different languages, Wikipedia may facilitate informational convergence, possibly through increased mutual visibility and editing spillovers. The inclusion of battle and language fixed effects ensures that this effect is not merely driven by more standardized conflicts or more active editing languages, reinforcing the argument for a causal role of link formation in shaping cross-lingual alignment.

By developing innovative measurement strategies, we aim to identify areas of narrative divergence, highlighting the potential of online platforms like Wikipedia to bridge gaps in understanding - rather than divide - and foster a process of converging “truths.” Additionally, we plan to examine the factors that moderate this process and assess the impact of targeted interventions, such as flagging diverging narratives, to promote alignment. On the other hand, non-convergence or even growing divergence in narratives could turn out to be a valuable indicator of latent conflict or even enable early detection of future outbreaks.

The broader impact of this work lies in its emphasis on fostering a mutual and shared understanding of historical events, which is crucial for conflict resolution and prevention and might even allow quantifying the “heat” of ongoing and latent conflicts.

REFERENCES

Adafre, S. F., & de Rijke, M. (2006). Finding similar sentences across multiple languages in Wikipedia. *Proceedings of the EACL 2006 Workshop on New Text*, 62–69.

Adar, E., Skinner, M., & Weld, D. S. (2009). Information arbitrage across multilingual Wikipedia. *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, 94–103.

Álvarez, G., Oeberst, A., Cress, U., & Ferrari, L. (2020). Linguistic evidence of in-group bias in English and Spanish Wikipedia articles about international conflicts. *Discourse, Context & Media, 35*, 100391. https://doi.org/10.1016/j.dcm.2020.100391

Araujo, A. L. (Ed.). (2013). *Politics of memory: Making slavery visible in the public space.* Routledge.

Assmann, A. (2011). *Cultural memory and Western civilization: Functions, media, archives.* Cambridge University Press.

Assmann, A., & Conrad, S. (Eds.). (2010). *Memory in a global age: Discourses, practices, and trajectories.* Palgrave.

Barberá, P. (2020). Social media, echo chambers, and political polarization. In N. Persily & J. A. Tucker (Eds.), *Social media and democracy: The state of the field, prospects for reform* (p. 34). Cambridge University Press.

Bekerman, Z., & Zembylas, M. (2011). *Teaching contested narratives: Identity and memory in peace education.* Cambridge University Press.

Blanshard, B. (1939). *The Nature of Thought* (Vols. 1–2). George Allen & Unwin.

Bloch, M. (1954). *The historian's craft.* Vintage Books.

Buckley-Zistel, S. (2006). Remembering to forget: Chosen amnesia as a strategy for local coexistence in post-genocide Rwanda. *Africa: Journal of the International African Institute, 76*(2), 131–150.

Carr, E. H. (1961). *What is history?* Macmillan.

Callahan, E. S., & Herring, S. C. (2011). Cultural bias in Wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology, 62*(10), 1899–1915.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics, 225*(2), 200–230.

Choi, D., Chun, S., Oh, H., Han, J., & Kwon, T. T. (2020). Rumor propagation is amplified by echo chambers in social media. *Scientific Reports, 10*(1), 310.

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences, 118*(9).

Clark, P. (2010). *The Gacaca courts, post-genocide justice and reconciliation in Rwanda: Justice without lawyers.* Cambridge University Press.

Cohere. (n.d.). *Cohere embeddings*. Retrieved 10/12/2024, from https://cohere.com

Cornelißen, C. (2006). Europas Gedächtnislandkarte: Gibt es eine Universalisierung des Erinnerns? In N. Frei (Ed.), *Was heißt und zu welchem Ende studiert man Geschichte des 20. Jahrhunderts?* (pp. 42–49). Wallstein.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.0.

Doxtader, E. (2009). *With faith in the works of words: The beginnings of reconciliation in South Africa.* Michigan State University Press.

Erdmann, M., Nakayama, K., Hara, T., & Nishio, S. (2009). An approach for extracting bilingual terminology from Wikipedia. *Proceedings of the 13th International Conference on Database Systems for Advanced Applications*, 380–392.

Erll, A. (2011). *Memory in culture.* Palgrave Macmillan.

Fallis, D. (2008). Toward an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology, 59*(10), 1662–1674.

Ford, H., Sen, S., Musicant, D. R., & Miller, N. (2013). Getting to the source: Where does Wikipedia get its information from? *Proceedings of the 9th International Symposium on Open Collaboration*, 9:1–9:10.

Forte, A., & Bruckman, A. (2008). Scaling consensus: Increasing decentralization in Wikipedia governance. *Proceedings of the 41st Annual Hawaii International Conference on System Sciences*, 157.

Forte, A., Lampe, C., & Bruckman, A. (2012). Scaling consensus: Collective editing in Wikipedia. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 161–170.

Forte, A., & Lampe, C. (2013). Defining, understanding, and supporting open collaboration: Lessons from the literature. *American Behavioral Scientist, 57*(5), 535–547.

Galtung, J. (1996). *Peace by peaceful means: Peace and conflict, development and civilization.* SAGE Publications.

Garimella, K., De Francisci Morales, G., Gionis, A., & Mathioudakis, M. (2018). Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference* (pp. 913–922).
Garde-Hansen, J. (2011). *Media and memory.* Edinburgh University Press.
Garton Ash, T. (2011, March 16). Germany can show reborn Arab nations the art of overcoming a difficult past. *The Guardian.*
Greenstein, S., & Zhu, F. (2012). Is Wikipedia biased? *American Economic Review, 102*(3), 343–348. https://doi.org/10.1257/aer.102.3.343
Greenstein, S., & Zhu, F. (2018). Do experts or crowd-based models produce more bias? Evidence from Encyclopædia Britannica and Wikipedia. *MIS Quarterly, 42*(3), 945–959. https://doi.org/10.25300/MISQ/2018/11432
Habermas, J. (1984). *The Theory of Communicative Action: Reason and the Rationalization of Society* (Vol. 1). Beacon Press.
Halatchliyski, I., Kummer, M. E., Slivko, O., Giorgidze, G., & Saam, M. (2016). Centrality and content creation in networks: The case of economic topics on German Wikipedia. *Information Economics and Policy, 36*, 36–52. https://doi.org/10.1016/j.infoecopol.2016.05.001
Halbwachs, M. (1992). *On collective memory.* University of Chicago Press.
Hale, S. A. (2014). Multilinguals and Wikipedia editing. *Proceedings of the 2014 ACM Conference on Web Science*, 99–108.
Hayner, P. B. (2010). *Unspeakable truths: Transitional justice and the challenge of truth commissions.* Routledge.
Hecht, B., & Gergle, D. (2010). The Tower of Babel meets Web 2.0: User-generated content and its applications in a multilingual context. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 291–300.
Hinnosaar, M., Hinnosaar, T., Kummer, M., & Slivko, O. (2023). Wikipedia matters. *Journal of Economics & Management Strategy.* https://doi.org/10.1111/jems.12421
Hunt, L. (1989). *The new cultural history.* University of California Press.
James, W. (1907). *Pragmatism: A New Name for Some Old Ways of Thinking.* Longmans, Green, and Co.
Jemielniak, D. (2014). *Common Knowledge? An Ethnography of Wikipedia.* Stanford University Press.
Jenkins, H. (2006). *Convergence Culture: Where Old and New Media Collide.* New York University Press.
Kim, S., Park, S., Hale, S. A., Kim, S., Byun, J., & Oh, A. (2016). Understanding editing behaviors in multilingual Wikipedia. *PLOS ONE, 11*(5). https://doi.org/10.1371/journal.pone.0155305
Kirkham, R. L. (2001). *Theories of Truth – A Critical Introduction.* MIT Press.
Knigge, V. (2008). Europäische Erinnerungskultur: Identitätspolitik oder kritisch-kommunikative historische Selbstvergewisserung. In Kulturpolitische Gesellschaft e.V. (Ed.), *kultur.macht.europa – europa.macht.kultur* (pp. 69–82). Klartext.
Kittur, A., & Kraut, R. E. (2008). Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, 37–46.
Kummer, M., Slivko, O., & Zhang, X. (2019). Unemployment and digital public goods contribution. *Information Systems Research.* https://doi.org/10.1287/isre.2019.0916
Laitenberger, U., Viete, S., Slivko, O., Kummer, M., Borchert, K., & Hirth, M. (2023). Unemployment and online labor: Evidence from microtasking. *MIS Quarterly.* https://doi.org/10.25300/MISQ/2022/17347
Larson, D. W., & Shevchenko, A. (2010). Status seekers: Chinese and Russian responses to US primacy. *International Security, 34*(4), 63–95.
Lederach, J. P. (2005). *The moral imagination: The art and soul of building peace.* Oxford University Press.
Leggewie, C., & Lang, A. (2011). *Der Kampf um die europäische Erinnerung: Ein Schlachtfeld wird besichtigt.* C.H. Beck.
Lewoniewski, W., Węcel, K., & Abramowicz, W. (2017). Multilingual ranking of Wikipedia articles with quality and popularity assessment in different language editions. *Online Information Review, 41*(4), 478–495.
Lih, A. (2009). *The Wikipedia Revolution: How a Bunch of Nobodies Created the World's Greatest Encyclopedia.* Hyperion.
Liu, S., Yan, X., Li, H., Zhao, T., & Wang, Y. (2018). Cross-lingual entity alignment via attribute-enhanced embedding. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2486–2495.
Luyt, B., & Tan, D. (2010). Improving Wikipedia's accuracy: Is edit age a solution? *Journal of the American Society for Information Science and Technology, 61*(12), 2582–2593.

Maaten, L. van der, & Hinton, G. (2008). Visualizing data using t-SNE. Journal of Machine Learning Research, 9(86), 2579–2605.
Magnus, P. D. (2009). On trusting Wikipedia. *Episteme, 6*(1), 74–90.
Minow, M. (1998). Between vengeance and forgiveness: Facing history after genocide and mass violence. Beacon Press.
Miller, D. L. (2023). An introductory guide to event study models. *Journal of Economic Perspectives*, 37(2), 203-230.
Miquel-Ribé, M., & Laniado, D. (2018). Wikipedia Cultural Diversity Dataset: A cartography of Wikipedia and Wikimedia projects. *Proceedings of the 14th International Symposium on Open Collaboration*, 1–9.
Montville, J. V. (1990). The healing function in political conflict resolution. Journal of Social Issues, 46(2), 33–53.
Noddings, N., & Brooks, L. (2016). Teaching controversial issues: The case for critical thinking and moral commitment in the social studies. Teachers College Press.
Nora, P. (1996). Realms of memory: Rethinking the French past. Columbia University Press.
Olick, J. K., Vinitzky-Seroussi, V., & Levy, D. (2011). The collective memory reader. Oxford University Press.
Peirce, C. S. (1878). How to make our ideas clear. *Popular Science Monthly, 12*, 286–302.
Pfeil, U., Zaphiris, P., & Ang, C. S. (2006). Cultural differences in collaborative authoring of Wikipedia. *Journal of Computer-Mediated Communication, 12*(1), 88–113.
Radonic, L. (2010). Krieg um die Erinnerung: Kroatische Vergangenheitspolitik zwischen Revisionismus und europäischen Standards. Campus.
Radonic, L. (2014). Postsozialistische Gedenkmuseen zwischen Opfernarrativen und der 'Europäisierung der Erinnerung.' Jahrbuch für Politik und Geschichte, 5, 85–106.
Rambachan, A., & Roth, J. (2019). An honest approach to parallel trends. *Working paper, Harvard University*.
Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, *90*(5), 2555-2591.
Reagle, J. (2010). *Good Faith Collaboration: The Culture of Wikipedia.* MIT Press.
Regan, P. (2010). *Unsettling the settler within: Indian residential schools, truth telling, and reconciliation in Canada.* UBC Press.
Rescher, N. (1973). *The Coherence Theory of Truth.*
Rescher, N. (1993). *A System of Pragmatic Idealism, Volume II: The Validity of Values.* Princeton University Press.
Ricoeur, P. (2004). *Memory, history, forgetting.* University of Chicago Press.
Rousso, H. (2004). Das Dilemma eines europäischen Gedächtnisses. *Zeithistorische Forschungen, 1*, 363–378.
Rousso, H. (2011). History of memory, policies of the past: What for? In K. H. Jarausch & T. Lindenberger (Eds.), *Conflicted memories: Europeanizing contemporary histories* (pp. 23–38). Berghahn.
Samoilenko, A., Karimi, F., Edler, D., Kunze, T., & Strohmaier, M. (2017). Linguistic neighbourhoods: Explaining cultural borders on Wikipedia through language co-editing activity. *Proceedings of the Eleventh International AAAI Conference on Web and Social Media*, 217–226.
Simpson, A. (2014). *Mohawk interruptus: Political life across the borders of settler states.* Duke University Press.
Sommer-Sieghart, M. (2006). Historische Ausstellungen als 'Contested Space.' In J. Feichtinger et al. (Eds.), *Schauplatz Kultur – Zentraleuropa: Transdisziplinäre Annäherungen* (pp. 159–166). Studienverlag.
Team, Gemini, Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., Millican, K., Silver, D., Petrov, S., Johnson, M., Antonoglou, I., Schrittwieser, J., Glaese, A., Chen, J., Pitler, E., ... & Vinyals, O. (2023). Gemini: A family of highly capable multimodal models. arXiv preprint arXiv:2312.11805.
Terren, L. T. L., & Borge-Bravo, R. B. B. R. (2021). Echo chambers on social media: A systematic review of the literature. *Review of Communication Research, 9*.
Tosh, J. (2015). *The pursuit of history.* Routledge.
Tutu, D. (1999). *No future without forgiveness.* Random House.
Vrandecic, D. (2012). Wikidata: A new platform for collaborative data collection. *Proceedings of the 21st International Conference on World Wide Web*, 1063–1064.
Weizman, E. (Ed.). (2014). *Forensis: The architecture of public truth.* Sternberg Press.
Wilson, R. (2001). *The politics of truth and reconciliation in South Africa: Legitimizing the post-apartheid state.* Cambridge University Press.
Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). MIT Press.
Zehr, H. (2002). *The little book of restorative justice.* Good Books.

Zhang, X. (Michael), & Zhu, F. (2011). Group size and incentives to contribute: A natural experiment at Chinese Wikipedia. *American Economic Review, 101*(4), 1601–1615.
Zhu, K., Walker, D., & Muchnik, L. (2020). Content growth and attention contagion in information networks: Addressing information poverty on Wikipedia. *Information Systems Research, 31*(2), 491–509.

APPENDIX

In this part we present the results when analyzing the differences between winners and losers only. The patterns are more pronounced but only fewer observations (one per battle, in some cases) can be used.

Figure A1: Difference in Casualty Differentials between Winners and Losers

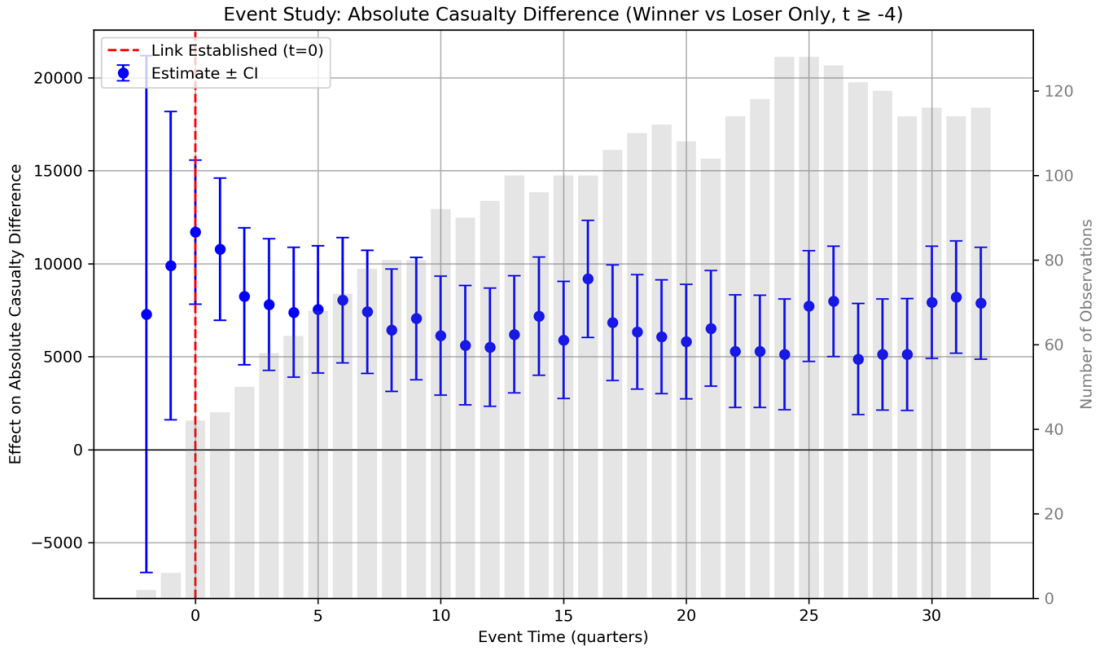


Figure A2: Difference in Troop Strength Differentials between Winners and Losers

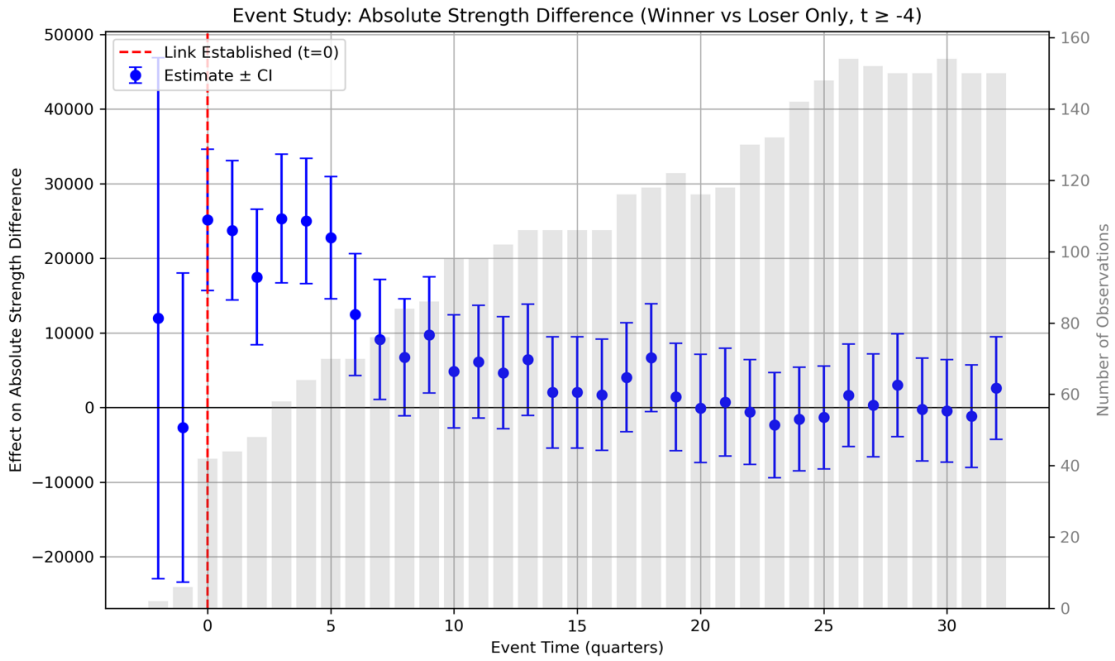


Figure A3: Long Run Effects Analysis



