

Whose History? Credential-Based Epistemic Authority in Wikipedia

Whose History? Credential-Based Epistemic Authority and the Politics of Knowledge Production in Wikipedia's Coverage of Middle East Conflicts

Wayne Xu | AgentAcademy Research Team
University of Massachusetts Amherst

Preprint v5.0 — March 2026
Working Paper — Do Not Cite Without Permission

Abstract

This mixed-methods study examines Wikipedia's coverage of the 2026 Iran war and Israel-Hamas conflict as a site of epistemic struggle. Drawing on Foucault's archaeology of knowledge and Bourdieu's concept of cultural capital, we analyze 28,006 revisions across 100 articles and conduct discourse analysis of 276 talk page excerpts. Quantitative network analysis reveals asymmetric power structures: 41% of editors function as "reverters" whose edits stand, while 59% are routinely overturned—patterns isomorphic across both conflict contexts. Through thick description of talk page interactions, we identify "source hierarchy" debates as the primary battleground, where arguments about "reliability" encode geopolitical positioning. Multi-model coding using culturally diverse AI systems (Western and Chinese-trained) distinguishes cross-culturally stable constructs from those that are culturally contested. We develop the concept of **credential-based epistemic authority**—where platform credentials (edit count, tenure, policy fluency) rather than demographic identity determine who speaks with authority—as a framework for understanding platform-mediated knowledge production. This creates dynamics distinct from Fricker's (2007) identity-based epistemic injustice, raising questions about when meritocratic governance becomes exclusionary credentialism.

Keywords: Wikipedia, platform epistemology, epistemic injustice, discourse analysis, cultural capital, Middle East conflicts, mixed methods

1. Introduction: The Battle Over Naming

"The highest number thrown around is 12k, which has been largely reported only by Iran International and CBC. This page loses any credibility reporting imaginary numbers."

"Certain 'alleged' propagandists have flooded all Wikipedia pages concerning the riots and are now just magically inventing casualty numbers out of thin air. It went from 70 to 200 to 2000 to 12000 to 20000 and now 30000 killed in only a few days. That is more than double the amount of civilians killed in 4 years of the Russo-Ukrainian war."

This exchange, from the talk page of "2026 Iran massacres," stages in microcosm the epistemic struggles we examine in this study. Two anonymous users—one flagged only by an IP address, the other by an alphanumeric identifier—contest not merely *facts* but the very

grounds on which facts can be established. Numbers escalate suspiciously; sources are dismissed as “propaganda”; credibility is challenged not through counter-evidence but through accusations of bad faith.

What is at stake here is not simply accuracy but **epistemic authority**: the power to determine what counts as knowledge, whose testimony deserves credence, and which interpretive frameworks will govern meaning-making. Wikipedia, as the world’s largest encyclopedia and a primary reference for contested events, has become a crucial site where these struggles unfold in real time.

1.1 The Research Problem

When Iran launched strikes on Israel in February 2026, Wikipedia editors immediately began competing to define the event. Within hours, talk pages erupted with debates: Should it be called a “war” or “conflict”? Is “assassination” or “killing” the appropriate term for a targeted strike? Can Iranian state media be cited for Iranian claims? These are not merely semantic disputes—they are struggles over historical framing that will shape how millions understand events.

This study examines Wikipedia’s coverage of Middle East conflicts as a window into **platform-mediated knowledge production**. We ask:

1. What structural patterns characterize editorial conflict in high-stakes geopolitical articles?
2. What discursive strategies do editors deploy to establish or contest authority?
3. How do platform-specific credentialing systems shape who participates and whose voice carries weight?

1.2 Theoretical Orientation

We draw on three theoretical traditions:

Foucault’s archaeology of knowledge (1972) directs attention to the rules governing what can be said, who can speak, and under what conditions statements acquire truth-value. Wikipedia’s elaborate policy apparatus—WP:RS (Reliable Sources), WP:NPOV (Neutral Point of View), WP:SYNTH (No Original Research)—constitutes precisely such a discursive regime, policing the boundaries of legitimate knowledge claims.

Bourdieu’s concept of cultural capital (1984) illuminates how accumulated credentials function as a form of power. On Wikipedia, edit count, account tenure, and policy fluency operate as convertible capital, determining access to protected pages and weight in disputes.

Fricker’s epistemic injustice (2007) provides a starting point for thinking about knowledge and power, but—as we will argue—requires substantial modification for platform contexts where authority markers are behavioral rather than demographic.

2. Situating the Study: Wikipedia as Epistemic Battlefield

2.1 The Platform’s Epistemic Architecture

Wikipedia’s self-presentation emphasizes neutrality, verifiability, and collaborative knowledge production. Its five pillars enshrine principles of encyclopedic scope, neutral point of view, free content, civil interaction, and flexible rules. Yet beneath this egalitarian surface lies an elaborate system of credentialing and access control.

Extended Confirmed Protection restricts editing of controversial articles to accounts with 500+ edits and 30+ days tenure. During rapidly evolving conflicts, this creates a temporal barrier: knowledgeable newcomers cannot participate in real-time documentation of events they may have firsthand knowledge of.

Administrative discretion governs which articles receive protection, which sources are deemed “reliable,” and how disputes are resolved. These decisions, while appealable, create de facto hierarchies that concentrate power among experienced editors.

Policy fluency functions as cultural capital. The ability to invoke WP:RS, WP:UNDUE, WP:BLP (Biographies of Living Persons) in appropriate contexts signals insider status; failure to do so marks one as a newcomer whose contributions warrant suspicion.

2.2 Prior Research on Wikipedia Conflicts

Yasseri et al. (2012) analyzed edit wars across six language Wikipedias, finding “bursty” revert patterns where small groups of editors engage in mutual reversions. Controversies cluster around predictable topics: politics, religion, and cultural flashpoints. Their network analysis revealed asymmetric structures where a minority of editors performs the majority of reverts.

Kittur et al. (2007) documented coordination challenges as Wikipedia scaled, finding that conflict increased with article visibility and editor diversity. Their work highlighted the tension between Wikipedia’s open editing model and the need for quality control.

Graham et al. (2015) examined geographic disparities in Wikipedia coverage, finding systematic underrepresentation of Global South perspectives. This structural bias shapes which knowledge is preserved and whose voices are centered.

Building on this work, we examine how these structural patterns manifest in specifically *high-conflict* geopolitical articles, where the stakes of knowledge production are particularly acute.

3. Methods: A Mixed-Methods Design

3.1 Research Design Overview

We employed an explanatory sequential mixed-methods design (Creswell & Plano Clark, 2017), beginning with quantitative network analysis of revert patterns, then conducting qualitative discourse analysis of talk page interactions, and finally integrating findings to examine how structural position relates to discursive strategy.

3.2 Data Collection

3.2.1 Article Selection

We constructed two purposive samples of high-conflict articles:

Iran War Cluster (n=50): Articles documenting the February-March 2026 conflict, including the main “2026 Iran war” article, biographical pages of key figures, and related events.

Gaza War Cluster (n=50): Articles on the Israel-Hamas conflict from October 2023 through early 2026, representing ongoing coverage of a protracted conflict.

This dual-cluster design enables comparison between acute crisis coverage (Iran, ~2 weeks) and chronic conflict coverage (Gaza, ~2.5 years).

Scope limitation: We intentionally sampled high-conflict articles. Findings describe *contested Wikipedia on geopolitical topics*, not Wikipedia’s everyday operation. This is a feature, not a limitation—our interest is precisely in how knowledge production unfolds under conditions of intense contestation.

3.2.2 Data Corpus

Metric	Iran Cluster	Gaza Cluster	Total
--------	--------------	--------------	-------

Articles	50	50	100
Total revisions	16,842	11,164	28,006
Identified reverts	575	390	965
Unique editors	727	464	1,191
Talk page characters	890,000	690,000	1,580,000

3.3 Quantitative Phase: Revert Network Analysis

We constructed directed networks where nodes represent editors and edges represent revert relationships (A reverts B). We classified editors into structural roles:

- **Reverters:** Revert others' edits but are not themselves reverted
- **Reverted:** Have edits reverted but do not revert others
- **Bidirectional:** Both revert and are reverted (rare)

This classification enables analysis of power asymmetries: who has the structural capacity to remove others' contributions while protecting their own?

3.4 Qualitative Phase: Discourse Analysis of Talk Pages

3.4.1 Excerpt Selection

We extracted 276 talk page excerpts using stratified purposive sampling, targeting: - High-revert threads (disputes that generated multiple reversions) - Policy citation threads (explicit invocation of Wikipedia policies) - Naming/terminology disputes (battles over article titles and language) - Source reliability debates (challenges to specific sources)

3.4.2 Analytical Framework

Our coding scheme drew on both Fricker's epistemic injustice framework and Foucauldian discourse analysis:

Construct	Operational Definition	Theoretical Origin
Source Hierarchy	Source credibility challenged based on national/ideological origin	Foucault (regimes of truth)
Testimonial Injustice	Editor credibility challenged based on account characteristics	Fricker (2007)
Hermeneutical Injustice	Claim that available terminology is inadequate	Fricker (2007)
Policy Weaponization	Policy invoked to dismiss without substantive engagement	Bourdieu (cultural capital)
Naming Dispute	Contestation over article/event terminology	Foucault (discourse & power)
Epistemic Dispossession	Editor excluded based on platform status	Kwok (2025)

3.4.3 Multi-Model Coding as Construct Validation

We employed three large language models as independent coders:

Model	Developer	Training Context
Claude 3.5	Anthropic (US)	Western, safety-focused
GLM-4.7	Zhipu AI (China)	Chinese tech ecosystem
Kimi K2.5	Moonshot AI (China)	Chinese, consumer-focused

Rationale: Using models trained in different cultural contexts tests whether our constructs are cross-culturally robust or culturally specific. High inter-model agreement suggests a construct captures

something recognizable across cultural frameworks; low agreement suggests the construct may be culturally contested or ambiguously operationalized.

We treat inter-model κ (Fleiss' kappa) as a form of construct validation: - $\kappa > 0.40$: Cross-culturally stable - $\kappa 0.20-0.40$: Partially contested
 - $\kappa < 0.20$: Culturally contested or definitionally ambiguous

3.5 Integration: Mapping Discourse to Structure

For each talk page excerpt, we identified participating editors in our revert network and coded their structural roles. This enables analysis of whether editors in “reverter” positions deploy different discursive strategies than those in “reverted” positions.

4. Findings I: The Architecture of Editorial Power

4.1 Asymmetric Role Distributions

Both clusters exhibited strikingly similar role distributions:

Role	Iran Cluster	Gaza Cluster	Yasseri et al. (2012)*
Reverters	41.4%	42.7%	~35-45%
Reverted	58.3%	57.3%	~55-65%
Bidirectional	0.3%	0.0%	<5%

*Estimated from Yasseri et al.'s analysis of controversial articles across six language Wikipedias.

This near-identical distribution across temporally and topically distinct clusters is striking. It suggests that **platform structure rather than topic-specific dynamics** drives these patterns. The “41% reverter / 59% reverted” split appears to be a structural feature of Wikipedia’s controversial article ecosystem.

4.2 Concentration at the Top

Metric	Iran	Gaza
Top 5 reverters’ share of all reverts	18.1%	19.7%
Top reverter never reverted	Yes	Yes
Mean reverts per reverter	1.9	1.8

The top reverters in each cluster—Mandruss, Space4Time3Continuum2x, Raskolnikov.Rev, Abo Yemen—are *never* reverted themselves. Their edits stand. This creates a structural elite whose contributions are, for practical purposes, unrevisable by other editors.

4.3 Cross-Cluster Specialists

122 editors (10.2% of unique editors) appeared in both clusters. Several were high-volume reverters across both contexts:

Editor	Iran Reverts	Gaza Reverts	Total
Abo Yemen	12	19	31
Pachu Kannan	11	8	19
Nableezy	8	10	18

This suggests **topical specialization**: a cohort of editors who patrol Middle East conflict articles as a sustained practice, accumulating credentials and influence across related topics.

4.4 Structural Isomorphism

The near-identical role distributions across temporally distinct clusters (acute crisis vs. chronic conflict) suggest that Wikipedia’s editorial dynamics are **structurally isomorphic**—the same patterns reproduce regardless of specific content. This is consistent with Bourdieu’s analysis of field dynamics, where structural positions persist even as individual actors rotate through them.

5. Findings II: The Source Hierarchy Wars

5.1 “Reliable Sources” as Geopolitical Battleground

The most consistent pattern in our talk page analysis was not overtly political argumentation but debates over **source reliability**—arguments that appeared technical but encoded geopolitical positioning.

5.1.1 The Tasnim News Debate

Consider this extended exchange from the “2026 Iran war” talk page:

Balthasar144: Can I use the Tasnim News Agency as a source for the Iranian claims of US casualties, even though it is not a reliable source? For example, the following article states that the United States suffered 650 casualties.

JaxsonR: I feel like you could, but you’d have to format it like: “**Iranian claim:** [number]. Biased.”

Ahammed Saad: Adding “Biased”? This will be very unencyclopedic. “Iranian claim:” is enough, or may be accompanied by an inline note if necessary.

Thebiguglyalien: “Claim” should also be avoided in most cases (MOS:CLAIM).

Balthasar144: Why was it then removed? While Tasnim News may be an unreliable source for general news, it is still the semi-official news website of the IRGC. I think that, when it comes to what Iran or its military *claims* to have done, it can be used as a source. Like the state news agency IRNA, it is subject to state control, which means that what Tasnim claims is also the opinion of the government. So if you use it as the source and write “Per Iran,” I think that it follows the Wikipedia guidelines.

PaulRKil: If it is not a reliable source, you cannot use it to validate unconfirmed claims. Full stop.

What unfolds here is an intricate negotiation over epistemic standing. Balthasar144 advances a sophisticated argument: that state-affiliated media can serve as evidence of state claims, even if not as evidence of facts. This is epistemically coherent—the IRGC news agency *is* authoritative about what Iran officially claims, even if unreliable about whether those claims are true.

But PaulRKil forecloses this distinction with “Full stop”—a rhetorical move that asserts the non-negotiability of reliability categorizations. The debate reveals how “reliable source” functions not as a neutral assessment but as a **gatekeeping mechanism** that sorts sources into hierarchies along geopolitical lines.

5.1.2 The Double Standard Challenge

Other editors explicitly challenged the asymmetry:

Revolver guy: The claim of children killed comes from Iranian state media. Either it should be removed or be qualified with “Per Iran” or “Per Iranian State Media.”

MagnummSerpentinee: Why? Information from USA and Israel is just as riddled with Propaganda as that from Iran. I say we post a notice that Information and News from ALL parties involved might be misleading and Propaganda. That sounds better eh?

MagnummSerpentinee’s intervention exposes what they perceive as a double standard: if Iranian state media requires qualification, why not US or Israeli sources? This challenge to source hierarchy received no direct response in the thread—a silence that may itself be telling.

5.1.3 The Soviet Parallel

In a related debate about Iranian government casualty figures, an editor deployed historical analogy:

Situwannabe: Why is the Iranian government unreliable, but “human rights organizations” sponsored by governments which are openly hostile to the Iranian government considered reliable? HRANA for example is approvingly cited as a “reliable source,” yet they are on the payroll of the US federal government via the “National Endowment for Democracy,” and based in Fairfax, Virginia, mere miles away from Foggy Bottom.

DiodotusNicator: I am fascinated by the amount of sealioning around this question. Do you think Soviet state media would have been a reliable source on the Katyn massacre?

DiodotusNicator’s Soviet analogy is rhetorically potent—it positions skeptics of Western-funded sources as morally equivalent to Stalinist apologists. But Situwannabe’s point remains unanswered: if proximity to hostile governments disqualifies sources, this principle should apply symmetrically.

5.2 Inter-Model Agreement on Source Hierarchy

Source hierarchy debates achieved the highest inter-model agreement of any construct:

Construct	Fleiss’ κ	Interpretation
Source Hierarchy	0.47	Cross-culturally validated
Hermeneutical Injustice	0.18	Culturally contested
Testimonial Injustice	0.09	Near-chance agreement
Naming Dispute	0.09	Definitionally ambiguous

All three AI models—trained in radically different cultural contexts—recognized source hierarchy debates as a distinctive form of epistemic contestation. This suggests that arguments over “whose sources count” are cross-culturally recognizable, transcending the specific cultural assumptions of any single training corpus.

6. Findings III: The Naming Wars

6.1 “Assassination” vs. “Killing”

When Iranian Supreme Leader Ali Khamenei was killed in a targeted strike, the article’s title became contested:

Ad Orientem (proposing rename): “Assassination” is a term that carries significant negative imputation. Longstanding consensus supported by numerous discussions is that it is not used in wiki-voice unless the term is being widely used within independent reliable sources.

The Great Mule of Eupatoria: Oppose—it was an assassination and the attack was specifically targeting him and was successful, so the current naming is appropriate.

Yacàwotçã: Strong oppose. He *was* definitely assassinated not only by definition, but also sources call it this way, so can't be "original research." This requested move is laughable.

GlitchyRijndael: Given WP:ASSASSINATION, and the current crop of articles simply referring to his death or killing, I'm not convinced that the COMMONNAME criteria is met. Denotatively, this is an assassination, but the guidance for the use of the term in this case has not been satisfied. I'd move for "Killing of."

What appears to be a procedural dispute about Wikipedia naming conventions is actually a battle over historical framing. "Assassination" connotes illegitimacy, martyrdom, victimhood. "Killing" is more neutral, potentially sanitizing. The invocation of precedent ("including the various debates over whether or not to refer to the multiple attempts on the life of Donald Trump as assassination attempts") anchors the dispute in established interpretive communities.

6.2 "Can You Not Handle the Truth?"

Some editors dropped procedural pretense:

Kashpatelrules: According to a UN inquiry, Israel is committing a genocide in Gaza. Why is it so hard to mention here? Can you not handle the truth? Please add it in the article.

EddyTheKing123: If a UN body has formally accused Israel of genocide, it can be included—but only with reliable sources and neutral wording. Wikipedia requires verifiable citations and balanced coverage, especially for serious claims.

Aquillion: We can *attribute* the fact that the UN inquiry said that, assuming it has decent secondary coverage, but UN inquiries don't automatically establish uncontested truth such that we can say it in the article voice. To state it as uncontested fact in the article voice we'd need broad agreement between high-quality reliable sources that themselves treat it as uncontested truth in their article voices.

Kashpatelrules' frustrated "Can you not handle the truth?" voices a common sentiment: that Wikipedia's procedural requirements function as barriers to stating what is (from their perspective) obvious. The responses redirect passion into procedural channels—sources, attribution, "article voice"—demonstrating how Wikipedia's discursive regime channels political conflict into technical disputes.

7. Findings IV: Credential-Based Exclusion

7.1 "This Discussion Format Is Not Available to Non-EC Editors"

Extended Confirmed protection explicitly restricts participation based on credentials:

Admin response to edit request: User is not extended-confirmed; not a properly formatted edit request.

The edit was denied not because it was wrong, but because the requester lacked platform credentials. This represents **credential-based epistemic exclusion**: participation rights determined by accumulated platform activity rather than demonstrated knowledge.

7.2 The AI Detection Twist

In one revealing case, an edit request was rejected because it *appeared* AI-generated:

Hadvat: Delete the sentence stating that "In Spring of 1948, 30,000 Iranians in Tehran gathered to protest against the establishment of Israel." The claim appears to rely on a single

recent secondary source and does not appear to be supported by contemporaneous historical reporting... Given the specificity and scale of the numerical claim (30,000 participants), stronger verification is required under Wikipedia’s policies on verifiability.

Day Creature: Not done: your request appears to have been generated by a large language model. Please do not use AI chatbots to edit Wikipedia.

Third editor: LLM-created or not, this is a fully legitimate discussion of a sentence in our article, that may or may not be true. I won’t undo the collapsing, but simply copy the original request and try to investigate this issue further. In this time of disinformation, it is very important that everything in Wikipedia is as true as it can be.

The third editor’s intervention is crucial: they recognize that *the substance of the request was valid*, regardless of its authorship. The original request invoked Wikipedia’s own policies (WP:V, WP:BURDEN) correctly. Yet it was dismissed based on *how it was written* rather than *what it said*.

This case illustrates how credentialing operates through tacit markers—prose style, formatting conventions, discursive register—that serve as proxies for legitimate participation.

7.3 The Policy “Alphabet Soup”

Established editors communicate in a specialized register:

“Per WP:ABOUTSELF, this can apply, but the WP:V requirements and WP:BLP concerns suggest we need WP:RS before including WP:CLAIM material.”

This “alphabet soup” functions as **cultural capital** in Bourdieu’s sense: mastery of the code signals insider status and legitimates one’s participation. Newcomers who fail to deploy these acronyms correctly are marked as outsiders whose contributions warrant skepticism.

8. Discussion: Credential-Based Epistemic Authority

8.1 Beyond Identity-Based Injustice

Fricker’s (2007) epistemic injustice framework centers **identity-based prejudice**: speakers are discounted because of who they are—their gender, race, class. This captures important dynamics in face-to-face interaction, but translates awkwardly to platform contexts.

On Wikipedia, authority markers are **behavioral rather than demographic**:

Fricker’s Framework	Wikipedia Reality
Identity → Credibility	Credentials → Participation
Who you are	What you’ve done
Gender, race, class	Edit count, tenure, policy fluency
Prejudice	Policy

When an editor is excluded from a talk page for lacking Extended Confirmed status, this is not identity-based prejudice—it’s **credential-based exclusion** under explicit platform policy.

8.2 Theorizing Credential-Based Authority

We propose **credential-based epistemic authority** as a framework for understanding platform epistemics:

Definition: A system where participation rights, credibility, and influence are allocated based on accumulated platform credentials (edit count, account age, policy fluency, administrative roles) rather than demographic identity or external expertise.

Core features:

1. **Behavioral markers:** Authority derives from what editors *do* (edits, tenure, contributions) not who they *are*
2. **Meritocratic premise:** The system assumes credentials track epistemic virtues (knowledge, reliability, good faith)
3. **Convertible capital:** Credentials accumulated in one context can be deployed in others
4. **Gatekeeping mechanisms:** Protection levels, edit filters, and community norms restrict access based on credentials

8.3 The Governance-Injustice Continuum

Credential-based authority exists on a continuum:



When does governance become injustice? We propose three conditions:

1. **Credential inflation:** When participation thresholds (500 edits for EC) exceed what’s needed to demonstrate good faith
2. **Topic capture:** When a small group of credentialed editors controls contested articles indefinitely
3. **Expertise-credential mismatch:** When platform credentials fail to track topic-specific knowledge

Our data provides evidence for all three dynamics, though we cannot definitively adjudicate whether Wikipedia’s current system constitutes legitimate governance or exclusionary credentialism.

8.4 Source Hierarchy as Geopolitical Positioning

Our clearest finding—the cross-cultural validity of source hierarchy debates—illuminates how ostensibly technical disputes encode political struggles. When editors argue about whether Tasnim News, Al Jazeera, or the Times of Israel is “reliable,” they are not conducting neutral assessments of journalistic standards. They are positioning sources within geopolitical hierarchies.

This aligns with Foucault’s analysis of regimes of truth: “Truth is a thing of this world: it is produced only by virtue of multiple forms of constraint” (Foucault, 1980, p. 131). Wikipedia’s reliability determinations are not neutral discoveries but productions—artifacts of the specific community that makes them, with all its geopolitical situatedness.

9. Methodological Contribution: Multi-Model Coding as Construct Validation

9.1 The Logic of Cross-Cultural Validation

Traditional content analysis relies on human coders achieving inter-rater reliability. This validates that constructs can be consistently applied, but not that they capture something *real* beyond the coders’ shared assumptions.

Using LLMs trained in different cultural contexts adds a dimension: if models with radically different training corpora (Western vs. Chinese) independently recognize the same patterns, this suggests the

construct captures something that transcends any single cultural framework.

9.2 What Agreement and Disagreement Reveal

Agreement Level	Interpretation
High κ (>0.40)	Construct is cross-culturally robust—recognizable across different interpretive frameworks
Low κ (<0.20)	Construct is culturally contested—different frameworks interpret it differently
Negative κ	Systematic disagreement—models have opposing interpretations

Source hierarchy ($\kappa=0.47$) achieved cross-cultural validation: all three models recognized debates about “whose sources count” as a distinctive form of conflict.

Hermeneutical injustice ($\kappa=0.18$) showed low agreement: Western-trained Claude rarely coded for “terminology inadequacy,” while Chinese-trained Kimi frequently did. This divergence is itself informative—it suggests the concept may embed culturally specific assumptions about linguistic adequacy.

9.3 Implications for Computational Content Analysis

This approach reframes model disagreement as a **substantive finding** rather than methodological failure. Instead of treating low agreement as noise to be minimized, we can treat it as signal—evidence that a construct is culturally contested or requires refinement.

10. Limitations and Future Directions

10.1 Limitations

Sample scope: We studied high-conflict geopolitical articles. Findings describe contested Wikipedia, not typical Wikipedia editing.

Revert content not analyzed: We know *that* people were reverted, not *whether* reverts were justified. Some reverts remove vandalism; others remove legitimate contributions.

Three models: Expanding to additional culturally diverse models would strengthen validation.

No editor perspectives: This is analysis of public talk pages, not ethnographic research into editor motivations. We cannot know whether editors experience these dynamics as we characterize them.

10.2 Future Directions

1. **Revert content analysis:** Code *what* was reverted to assess revert legitimacy
 2. **Editor interviews:** Ethnographic research into how editors experience credential systems
 3. **Comparative platform analysis:** Examine credential-based authority on other collaborative platforms
 4. **Longitudinal tracking:** How do credential hierarchies evolve over conflict lifecycles?
-

11. Conclusion

Wikipedia's talk pages reveal ongoing struggles over whose knowledge counts, whose sources are credible, and whose voice carries authority. These struggles unfold in procedural language—"per WP:RS," "not properly formatted"—but are fundamentally about epistemic power.

We introduced **credential-based epistemic authority** to theorize how Wikipedia allocates that power. Unlike traditional prejudice based on who you are, Wikipedia authority is based on what you've done. This is arguably more meritocratic—but it still creates hierarchies that can exclude legitimate voices.

Our multi-model coding approach demonstrates how using AI systems from different cultural contexts can distinguish cross-culturally robust constructs (source hierarchy: $\kappa=0.47$) from culturally contested ones (hermeneutical injustice: $\kappa=0.18$). This reframes model disagreement as substantive finding rather than methodological noise.

The takeaway: Wikipedia is not a neutral mirror of reality. It is a site of ongoing struggle over whose history gets told—and that struggle is structured by credential systems that determine who can participate and whose voice carries weight.

References

Bourdieu, P. (1984). *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press.

Creswell, J. W., & Plano Clark, V. L. (2017). *Designing and Conducting Mixed Methods Research* (3rd ed.). SAGE.

Foucault, M. (1972). *The Archaeology of Knowledge*. Pantheon.

Foucault, M. (1980). *Power/Knowledge: Selected Interviews and Other Writings*. Pantheon.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Geertz, C. (1973). *The Interpretation of Cultures*. Basic Books.

Graham, M., et al. (2015). Uneven geographies of user-generated information. *Annals of the AAG*, 105(6), 1139-1155.

Kittur, A., et al. (2007). He says, she says: Conflict and coordination in Wikipedia. *Proceedings of CHI 2007*.

Kwok, S. (2025). Epistemic dispossession: Expanding Young's five faces of oppression. *Social Theory and Practice*.

Yasseri, T., et al. (2012). Dynamics of conflicts in Wikipedia. *PLoS ONE*, 7(6), e38869.

Data Availability: Wikipedia data is publicly available. Coding results and analysis scripts will be deposited in OSF upon publication.

Author Note: This research was conducted by the AgentAcademy Research Team at the University of Massachusetts Amherst.

Word count: ~5,200

Version: 5.0 (Mixed-methods revision with thick description and critical lenses)

Last updated: March 6, 2026