

Red Flags and Cherry Picking: Reading the Scientific Blackpill Wiki

Celia Chen¹, Alex Leitch¹, Scotty Beland¹, Ingo Burghardt¹, William Conway¹, Rajesh Kumar Gnanasekaran¹, Marilyn Harbert¹, Emily Klein², and Jennifer Golbeck¹

¹ University of Maryland, College Park, MD, USA
clichen@umd.edu

² Moonshot, London, UK

Abstract. Incels are an online community of men who share a belief in extreme misogyny, the glorification of violence, and biological essentialism. They refer to their core ideology as "the blackpill," a belief that physical attraction is the only path to romantic success and that women are attracted only to one very specific, hypermasculine archetype. Incels believe their ideology is grounded in hard science. The research incels use as evidence is collected in an extensive online document, the Scientific Blackpill page on `incels.wiki`. We collected all 578 citations on the page and coded each against the source it cites. We find that the page overwhelmingly cites real, recent, peer-reviewed research (68% journal articles; median publication year 2014) and describes it accurately or with partial accuracy in nearly 90% of cases. Distortion is concentrated in the interpretive layer built around the reporting of the findings: context stripping (158 coded instances), overgeneralization (149), loaded language (111), and cherry-picking (75). Content structured in this way defeats interventions that focus on source credibility or claim-level falsity by reproducing reliability markers valued by AI synthesis systems and search rankings: declared neutrality, dense citation, structured headings, and quantitative language. We discuss what this implies for disclosure requirements in synthesized search results.

Keywords: Online harms · Radicalization · Science misuse · Platform governance · Information quality

1 Introduction

The incel community is an online misogynistic hate movement comprising men who are unable to romantically connect with women, blaming both women and society at large for this failure. Members of the blackpill branch of the incel ideological community have been responsible for mass shootings and other acts of public violence [8].

Radicalization into violent inceldom is referred to as "taking the blackpill" [5]: adopting a set of beliefs in which genetics and appearance are the sole determinants of romantic success, and specifically that women are only interested in

hypermasculine, white men with features that are typically Aryan, though the community generally does not use that term. The blackpill is a fatalistic worldview in which personal agency is rejected and systemic problems are deemed unsolvable through individual action [6].

This paper investigates the Scientific Blackpill page on `incels.wiki`, an incel-community-maintained wiki. This 116,949-word page uses scientific research to argue for the blackpill worldview. Organized into 18 sections with 219 subsections, it cites 578 articles, news stories, technical reports, and blog posts drawn from 483 distinct works. The wiki covers topics including women’s attraction to men based on facial structure, height, and race; penis size; the role of mental health in romantic success; women’s commitment in relationships; and rape myths. It has been read approximately 1.73 million times since 2019.

The page’s compilation peer-reviewed research citations creates a veneer of scientific legitimacy that may make its claims more persuasive to vulnerable individuals seeking explanations for their personal romantic difficulties. Our concern is therefore not whether the blackpill worldview is correct, but whether the page’s characterizations of the research it cites match what that research reports; and, if they do not, what kind of intervention could detect the gap. Despite growing research on incel communities and radicalization processes [5,6,11], none has examined how the page uses the research it cites. In December 2025 we retrieved every source cited on the page and compared the page’s representation of it against the source itself. Eight coders applied a purpose-built codebook recording source type, factual accuracy, distortion strategy, and study-quality concerns, yielding a coded corpus of all 578 citations.

We ask three questions:

- RQ1** What kinds of sources does the page cite, and how peer-reviewed and how current is that evidence base?
- RQ2** How often does the page’s characterization of a source match what that source actually reports?
- RQ3** Where the characterization departs from the source, which distortion strategies are used, and with what relative frequency?

2 Related Work

2.1 Incel Communities and Radicalization

Incels (involuntary celibates) are an extremist online community unified by their inability to establish romantic relationships with women and characterized by misogyny, racism, and glorification of violence [11]. The community congregates almost entirely online, creating echo chambers where extreme beliefs are normalized. Members overwhelmingly voice support for misogynistic violence, leading scholars to treat incel violence as a new variety of terrorism with a hate crime dimension [8]. Golbeck et al. [5] find a four-stage radicalization process, with pre-radicalization marked by appearance concerns and isolation, searching for blame, radicalization producing a "breakthrough realization," and post-radicalization

characterized by liberation and belonging. The "blackpill pipeline" describes progression through increasingly radical positions [6], functioning as a filter that removes agency and concentrates the most susceptible members toward extremity.

2.2 Science Misuse as a Persuasive Strategy

Motivated reasoning provides a psychological foundation for how scientific evidence gets distorted [13]: motivated cognition leads people to reject findings that threaten core beliefs [15], and identity-protective cognition leads individuals to scrutinize dis-confirming information while accepting confirming information at face value [9]. Denialism research identifies characteristic tactics including fake experts, cherry-picking, and misrepresentation through logical fallacy [3].

More directly relevant is work on how online communities deploy "real" science. Kata [10] catalogs anti-vaccination tactics (such as selective citation, demands for impossible certainty, and reframing of scientific caveats as concealment), several of which occur in our data. Related work shows how selective and deceptive citation constructs competing apparent consensus from a shared literature [2], how perceived experts operate within antivaccine communities [7], how false consensus is manufactured in public discussion of COVID-19 science [4], and how e-prints circulate in fringe communities on Reddit and 4chan [17]. The desire for citable scientific objects is not unique to the manosphere.

3 Data

The Scientific Blackpill page is 116,949 words long—roughly 290 pages if printed. It states that it "maintains a neutral tone and conveys the scientific findings without judgment," while noting that sections demarcated as "discussion" may contain unsourced speculation. This self-description is a claim to epistemic authority.

The page is divided into 18 sections and 219 subsections (Table 1), and each subsection has a topic or theme. For example, "Male gang members have dramatically more female sexual partners" followed by a summary of related research, usually a Discussion containing commentary and analysis, sometimes Quotes from the supporting research, and always a References list. The median subsection runs to 356 body words, but the spread is wide (min 78, max 2,100) and the distribution has a long right tail (Fig. 1). That variation is driven primarily by the length of the Discussion and less the summary of research. Most subsections cite one or two outside sources, though some cite more than ten.

3.1 Provenance and Reach

The page has had three principal editors: BlackpillScience, Bibipi, and Altmark22, together accounting for 90% of content added and 95% of the 4,175 edits made between February 2019 and August 2025. Non-logged-in users account for

Table 1. Structure of the page. Sub.: subsections; Cit.: citations. Totals: 219 and 578.

Section	Sub.	Cit.	Focus	Section	Sub.	Cit.	Focus
Personality	21	78	Dark traits, rape	Voice	4	7	Voice pitch
Mental	10	23	porn Mental health, autism	Age	9	21	Younger, underage females
Race	17	40	Racialized out-comes	Hypergamy	18	60	Higher-status partners
Looks (Life)	10	34	Attractiveness, life	Cucks	5	18	Infidelity, divorce
Looks (Love)	18	27	Attractiveness, romance	Sluts	13	25	Female sexual activity
Face	14	26	Facial features	MeToo	12	26	"Fabricated" assault reports
Money	8	14	Male income	Health	15	20	Sex and health
Height	11	21	Male height	ItsOver	21	113	Hopelessness, suicide
Body	9	21	Strength, weight, disability	Penis	4	4	Size preference

a further 8.8% of content added, overlapping temporally with active periods for two of the three, which may indicate the same individuals editing while logged out. Concentration of editing is not by itself evidence of coordination, and we make no such claim.

The page has accumulated roughly 1.73 million views, at a median of about 16,000 per month. In January 2024 it received 201,710 views in a single month, more than twelve times the prevailing rate. This spike is unexplained by the current work; it coincides with documented growth in automated scraping for language model training corpora, but we have no server-side data with which to attribute it.

4 Methods

Our approach was to treat every citation on the page as a verifiable claim about a source, retrieve that source, and code the relationship between what the source reports and what the page says it reports. Codebook development and coding are described below; the full codebook, coding instrument, and de-identified coded dataset are available upon request.

4.1 Codebook Development

Initial codebook construction (v1.0) began with the lead author independently coding the first major section of the page, documenting emergent patterns of misrepresentation and establishing preliminary categories. This produced nine fields covering location on the wiki, bibliographic details of the source, the source’s actual finding, the wiki’s representation of it, a distortion type, and study-quality red flags.

To refine the codebook we conducted a pilot coding exercise with the full research team ($n = 8$). A brief section on penis size preferences was selected for calibration, chosen both for its brevity (four citations) and as an early assessment




Fig. 1. Distribution of subsection body length by topic. Violin width shows where subsections cluster; box marks the interquartile range, line the median; each point is one subsection.

of coders' comfort with the page's content. Team members independently coded this section and gave feedback on category definitions and boundary cases; following asynchronous discussion the team reached consensus on revisions (v2.0). The major changes were separating factual accuracy ("Veracity": Accurate, Partially True/Misleading, False/Deceptive) from interpretive distortion ("Distortion Type"); adding a section-title field, capturing the recurring case where a title claims more than the body beneath it; adding a severity scale for study-quality issues; and expanding the distortion categories to include Fabrication (invented statistics, contradicted sources, false attribution), Correlation Misinterpretation, Irrelevant Information (ad hominem attacks on cited researchers), and False Balance.

4.2 Data Coding

Following revision to v2.0, section assignments were distributed across the team. Coders read every subsection and every external reference in their assigned sections and coded the attributes of the surrounding discussion. Coding ran for approximately one month. The initial test section (coded under v1.0) and the pilot section were retroactively re-coded under v2.0 to ensure consistency. The eight coders were faculty and graduate researchers in information studies and computer science at a single institution, four women and four men. All had prior training in qualitative coding, and all had prior research experience with online

Table 2. Source types and veracity of representation. Veracity percentages are of the 535 citations that received a veracity judgment; the remaining 43 cited sources could not be retrieved.

Source type	<i>n</i>	%	Source type	<i>n</i>	%	Veracity	<i>n</i>	%
Journal art.	393	68.0	Book	15	2.6	Accurate	233	43.6
News art.	64	11.1	Govt. report	15	2.6	Partially true	247	46.2
Other	46	8.0	Conf. paper	8	1.4	False/deceptive	55	10.3
Blog post	32	5.5	Book chap.	5	0.9	Not retrievable	43	—

extremist communities, most through a shared earlier study of violent online content.

Sections were assigned to single coders rather than double-coded, with the exception of the four-citation pilot section, which the whole team coded to consensus. We therefore cannot report an inter-rater reliability statistic, and the frequencies below should be read as the product of a consensus-calibrated single-coder protocol rather than as independently replicated measurements.

5 Results

5.1 What Gets Cited by Scientific Blackpill (RQ1)

The page cites overwhelmingly formal, recent, peer-reviewed literature (Table 2). Journal articles account for 68% of all citations; adding conference papers and book chapters brings the peer-reviewed share to 70%. News articles (11%) and blog posts (6%) form a much smaller component. The literature is not notably outdated: median publication year is 2014 (IQR 2009–2018), with 74% published in 2010 or later and only 9% before 2000.

5.2 How Accurately Are Citations Represented (RQ2)

Of the 535 citations whose sources we retrieved, 44% were represented accurately and a further 46% partially accurately; only 10% were coded False/Deceptive. Study-quality red flags, where recorded, were mostly ordinary limitations (223 of 365 coded citations).

Most entries state the original finding accurately but strip the context the authors placed around it. For example, Section 14.2 is subtitled, "Women rapidly lose interest in sex once in a stable relationship or living with a man". It discusses Klusmann [12], a peer-reviewed article on how sexual interest relates to relationship duration across age groups. Klusmann found that men's motivation for sex remained fairly constant over time while women's varied, but other measures showed no comparable difference. The article also advances an evolutionary-psychological hypothesis, proposing that women use sex to "procure male resources." The wiki reports the survey accurately, and the statistics it cites are sound.

The page gives a fairly accurate 311-word summary, then adds a 1,010-word discussion beginning: "Contrary to common dating advice emphasizing soul

Table 3. Coded instances by distortion type and veracity of representation. Categories are not mutually exclusive; a citation may carry several.

Distortion type	Accurate	Partially true	False/decep.	Total
<i>Framing distortions</i>				
Context stripping	32	115	11	158
Overgeneralization	11	130	8	149
Loaded language	7	81	23	111
Cherry-picking	14	55	6	75
Deterministic framing	0	9	0	9
Mischaracterization	1	6	1	8
False balance	4	2	0	6
Exaggeration	0	5	0	5
Aggregation distortion	0	4	0	4
Essentialism	0	1	0	1
<i>Technical errors</i>				
Correlation misinterpretation	0	37	5	42
Irrelevant information	4	6	3	13
Statistical manipulation	0	9	2	11
<i>Fabrication</i>				
False attribution	1	22	20	43
Invented statistics/data	4	5	10	19
Contradicts source	0	3	10	13
Other	1	3	2	6
Total	79	493	101	673

matching and deep intimacy, the result suggests that, on average, a high degree of familiarity is actually detrimental to relationship stability because the woman gets bored, possibly due to the man opening up, revealing his flaws which is a low status signal." This carries several incel themes: criticism of mainstream dating advice, rejection of emotional connection between men and women, the framing of women as "bored," the premise that male emotional disclosure is low status, and the premise that women pursue only high-status men. We coded this entry Accurate, because the study's results are described correctly. The commentary around the science is speculative and reveals motivated reasoning, illustrating how the page borrows scientific authority without misstating scientific fact.

Section 3.5, "Being Asian in the USA is a primary predictor of never being kissed," demonstrates the Partially True category. It cites research showing that among first-year undergraduates at a U.S. institution, 28% of Asian-American students reported never having been kissed against 7–11% for white, Black, and Latino students [14]. The research does show that Asian-American students in this sample were more likely never to have kissed; it does not show that race is a "primary predictor." As such, we coded this Statistical Manipulation.

5.3 Which Distortion Strategies Are Used (RQ3)

Table 3 reports distortion strategies. Codes are not mutually exclusive, and most entries carry two or three.

The four most frequent strategies are context stripping, overgeneralization, loaded language, and cherry-picking. Context stripping and overgeneralization usually appear together: a finding from a constrained sample with stated boundary conditions is restated as a general law about women. Loaded language adds

an evaluative and charged opinion the source did not originally carry. Cherry-picking operates across the corpus as a whole, as the selection of which studies to include carries much of the argument, and a reader checking any individual citation will not see it. Fabrication of some kind accounts for 81 coded instances.

6 Discussion

A notable trend in the wiki is the tension between an appearance of rationalism and a drive to misrepresent findings from selected journal articles. The tension emerges from a community emphasis on rationality as an exclusively masculine value coupled with a strong subscription to the belief that women are irrational and easily persuaded by "dark triad" personality traits.

The data sources selected, while sometimes questionable, dated, or sexist themselves, nonetheless carry the institutional marks of scientific legitimacy: they are for the most part peer-reviewed, recent, and drawn from established journals. That is not the same as being methodologically sound (our coders recorded critical or major problems in 25 of the peer-reviewed sources) but it is the property that credibility heuristics, and readers, respond to. However, this is not how they are represented within the wiki. Cherry-picking is deployed to support the same worldview reinforced within incel-focused social spaces online, and the selection of what to include carries much of the argument. This has implications for the broader production of hierarchical authority through selective interpretation of data.

Research on motivated reasoning and science denialism has shown that individuals and groups often deploy the rhetorical style of science as a strategy for legitimizing inaccurate beliefs [13,15,9,3], and our analysis of the Scientific Blackpill aligns with those results. Given previous work finding that incels follow common radicalization processes, and that exposure to online content espousing the incel ideology is critical in that process [5], this kind of presentation of science-adjacent content allows vulnerable readers to rationalize increasingly extreme beliefs as "evidence-based." This dynamic shows how radicalization can proceed through cognitive co-option via individuals adopting the movement's epistemic rules before adopting its most extreme ideological positions.

Within Scientific Blackpill, the comfort of numbers may provide a sense of control over external phenomena, where facility with quantitative approaches can represent an alignment with the community-supported presentation of rationality as a masculine value based in quantitative reasoning. Conclusions in the wiki editing are reinforced through a directed editorial approach that prioritizes general inclusion of quantitative work rather than clear analysis of the findings of that work. The further use of quantitative analysis to show "rationality" circularly reinforces pre-existing beliefs about the value of such approaches in proving hierarchical masculinity.

This type of motivated reasoning and scientific misuse is common in other conspiratorial settings: selective citation [2], perceived expertise [7], and manufactured consensus [4] operate in much the same way among anti-vaccine and

COVID-19 communities. The commonality of process shared in these communities is what makes it difficult to intervene online spaces committed to harm.

Many interventions against online harm are attached to numbers that derived from source reputation or assertion falsity. The sources presented by Scientific Blackpill are mainly real, peer-reviewed, and relatively recent, and more than 90% being partially or accurately represented in the text of the wiki. The page structure is clear and well-organized, with "discussion" separated from findings and citations clearly linked, aligning the page with norms of information presentation valued in STEM spaces. The problems are entirely in the writing itself, resting on whether a claim in the text is one the citation actually supports, which has historically been a challenge to measure with automated heuristics.

This difficulty has some knock-on effects for synthesized text produced by generative or search-based systems. Structural features of Scientific Blackpill, chief among them a large volume of well-maintained, clearly-structured, heavily-cited text, make it attractive to automated retrieval systems as a possible information source. Should it be collected, as page-access statistics imply it may be, this would lead to the page's framing being included in automated knowledge systems as legitimate. The implication here is that automated systems may need to be developed that can expose interpretive distance between retrieved passages and their cited basis, rather than treating the presence of citation alone as terminal evidence for accuracy. Such systems would need to be evaluated against adversarial systems similar to this wiki.

7 Limitations

Our coding is single-coder per section and we report no inter-rater reliability; the reported frequencies should be read accordingly. Forty-three citations (7%) referenced sources we could not retrieve, mostly paywalled or dead links, and are excluded from veracity percentages. Distortion codes were recorded as free text and normalized during analysis to the categories in Table 3. We analyze the page as retrieved in December 2025; it is actively edited and has continued to grow since, with at least one subsection added after that date falls outside our coding. We cannot speak to whether readers in fact update their beliefs. Finally, we characterize the affordances of retrieval and synthesis systems analytically rather than through system-level measurement; establishing how often such systems surface this material is important future work, as is linking the cited corpus to bibliographic metadata to test whether distortion varies with venue prestige or open-access status.

8 Conclusion

The Scientific Blackpill assembles 578 citations to largely real, recent, peer-reviewed literature and represents that literature accurately or partially accurately in nearly 90% of cases. Its ideological work is done through context stripping, overgeneralization, loaded language, and selection, all strategies that op-

erate on true findings and leave factual record intact. Misinformation here is real information turned toward predetermined ends, in a form built to satisfy every credibility heuristic we currently deploy. Governance frameworks for online harms, and disclosure requirements for AI-synthesized search results in particular, will need to ask not whether a source exists but whether it supports the claim attached to it.

Acknowledgments. This work was supported by NSF Award #2331257, and by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at the University of Maryland administered by Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence (ODNI).

References

1. Ali, M., Zannettou, S.: From isolation to desolation: investigating self-harm discussions in incel communities. In: Proc. ICWSM (2024)
2. Beers, A., Nguyen, S., Starbird, K., West, J.D., Spiro, E.S.: Selective and deceptive citation in the construction of dueling consensus. *Science Advances* (2023)
3. Diethelm, P., McKee, M.: Denialism: what is it and how should scientists respond? *European Journal of Public Health* **19**(1), 2–4 (2009)
4. Efstratiou, A., Efstratiou, M., Yudhoatmojo, S., Blackburn, J., De Cristofaro, E.: “Here’s your evidence”: false consensus in public Twitter discussions of COVID-19 science. In: Proc. CSCW (2024)
5. Golbeck, J., Chen, C., Leitch, A.: “Why I took the blackpill”: a thematic analysis of the radicalization process in incel communities. In: Proc. ASONAM (2025)
6. Green, R., Fowler, K., Palombi, A.: The blackpill pipeline: a process-tracing analysis of the incel’s continuum of violent radicalization. *Sexualities* **28**(5–6), 1811–1831 (2025)
7. Harris, M.J., Murfeldt, R., Wang, S., Mordecai, E.A., West, J.D.: Perceived experts are prevalent and influential within an antivaccine community on Twitter. *PNAS Nexus* (2024)
8. Hoffman, B., Ware, J., Shapiro, E.: Assessing the threat of incel violence. *Studies in Conflict & Terrorism* **43**(7), 565–587 (2020)
9. Kahan, D.M.: Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making* **8**(4), 407–424 (2013)
10. Kata, A.: Anti-vaccine activists, Web 2.0, and the postmodern paradigm: an overview of tactics and tropes used online by the anti-vaccination movement. *Vaccine* (2012)
11. Klein, E., Golbeck, J.: A lexicon for studying radicalization in incel communities. In: Proc. ACM Web Science Conference, pp. 262–267 (2024)
12. Klusmann, D.: Sperm competition and female procurement of male resources as explanations for a sex-specific time course in the sexual motivation of couples. *Human Nature* **17**(3), 283–300 (2006)
13. Kunda, Z.: The case for motivated reasoning. *Psychological Bulletin* **108**(3), 480 (1990)

14. Lefkowitz, E.S., Wesche, R., Leavitt, C.E.: Never been kissed: correlates of lifetime kissing status in US university students. *Archives of Sexual Behavior* **47**(4), 1283–1293 (2018)
15. Lewandowsky, S., Oberauer, K., Gignac, G.E.: NASA faked the moon landing—therefore, (climate) science is a hoax. *Psychological Science* **24**(5), 622–633 (2013)
16. Russo, G., Horta Ribeiro, M., West, R.: Stranger danger! Cross-community interactions with fringe users increase the growth of fringe communities on Reddit. In: *Proc. ICWSM* (2024)
17. Yudhoatmojo, S., De Cristofaro, E., Blackburn, J.: Understanding the use of e-prints on Reddit and 4chan’s Politically Incorrect board. In: *Proc. Web Science* (2023)