# The Commons Captured

*A Critical History of Wikipedia*

Michael J Bommarito II

**DRAFT MANUSCRIPT**

May 2026

*Not for distribution or citation*

**The Commons Captured:**
**A Critical History of Wikipedia**



Cover design and typesetting by the author. Printed in the United States of America.

**Author Website:** `michaelbommarito.com`

ISBN 979-8-XXXX-XXXX-X (paperback)

ISBN 979-8-XXXX-XXXX-X (ebook)

**Publisher's Cataloging-in-Publication Data**
Names: Bommarito, Michael J., II, author.
Title: The captured commons :
a critical history of Wikipedia / Michael J Bommarito II.
Description: First edition. | Michigan : 2026. | Includes bibliographical references.
Identifiers: ISBN 979-8-XXXX-XXXX-X (paperback) | ISBN 979-8-XXXX-XXXX-X (ebook)
Subjects: LCSH:
Wikipedia.
Wikimedia Foundation.
Electronic encyclopedias—History.
User-generated content.
Open source software.
Copyleft.
Artificial intelligence—Social aspects.
Information commons.

Classification: LCC ZA4482 .B66 2026 | DDC 030—dc23

*Note on sources:* This book draws on Wikimedia Foundation documents, Wikipedia primary sources, academic research, news coverage, and interviews. See the Methodology appendix for detailed discussion of sources and verification practices.

# Draft Manuscript — Not for Distribution

*This document is an unfinished draft of* The Commons Captured: A Critical History of Wikipedia*, posted publicly by the author for transparency about a work in progress.*

***Please note:***

- *Chapters are at varying stages of completion. Many sections remain skeletal, undercited, or unrevised.*
- *Facts, dates, names, and quotations have not all been verified. Do not cite this draft.*
- *Prose, structure, and emphasis will change substantially before publication.*
- *Errors, omissions, and infelicities are the author's alone and will be corrected in subsequent revisions.*

*Comments and corrections are welcome: mike@273ventures.com.*



*Draft version: May 2026.*

*To the volunteers who built the commons,*
*and those who continue to defend it*

*We need to launch a movement to develop a universal free encyclopedia, much as the Free Software movement gave us the free software operating system GNU/Linux. The free encyclopedia will provide an alternative to the restricted ones that media corporations will write.*

Richard Stallman
"The Free Universal Encyclopedia and Learning Resource," 1999

# Contents

# Preface

In 1999, Richard Stallman imagined an encyclopedia that could not be enclosed. He had spent a decade watching corporations take freely shared software, modify it, and lock away the improvements. The solution, he argued, was a license that required sharing back. Any work built on the commons would become part of the commons. Stallman called this copyleft.

Two years later, Wikipedia launched under that principle. The GNU Free Documentation License required that derivative works remain free in the same sense as the original. Anyone could use Wikipedia's content. Anyone could improve it. No one could capture it.

Twenty-five years on, that promise lies in ruins.

Every major large language model trains on Wikipedia. OpenAI's GPT-3 drew 3 percent of its training data from the encyclopedia. Meta's LLaMA used 4.5 percent. Researchers have shown that these models memorize and reproduce Wikipedia text nearly verbatim. When users ask ChatGPT factual questions, it cites Wikipedia 47.9 percent of the time. The encyclopedia that a quarter million volunteers maintain has become the factual backbone of a trillion-dollar industry.

The license was supposed to prevent exactly this. CC-BY-SA, which replaced the GFDL in 2009, requires attribution and share-alike. If you build on Wikipedia, you must credit it. If you create a derivative work, that work must carry the same freedoms. But when a language model ingests Wikipedia and generates text, where is the attribution? Where is the share-alike? The license has no enforcement mechanism for what AI systems do.

The Wikimedia Foundation knows this. In April 2024, Foundation executives acknowledged that most of their enterprise customers were

"not compliant with the letter of the Creative Commons rules or the spirit of the licenses." The response was not enforcement. The response was continued service.

Wikimedia Enterprise now counts Google, Amazon, Meta, and Microsoft among its clients. The Foundation collects fees while the AI companies capture billions in market value built on volunteer labor. The editors who wrote the articles receive nothing. The readers who might have visited Wikipedia visit ChatGPT instead. Traffic to the encyclopedia fell 8 percent in 2024, with steeper declines likely as AI search becomes standard.

This is a book about that transformation—how a project founded on copyleft principles became AI infrastructure, and what happens when the institution created to protect a commons begins to profit from its capture.

We trace the arc from Stallman's 1985 experience with Symbolics—the moment that radicalized him about proprietary software—through Wikipedia's founding, its licensing controversies, its institutional growth, and its current crisis. We examine the 2009 license migration, a legally questionable maneuver that may have broken the guarantees copyleft was supposed to provide. We follow the Wikimedia Foundation's evolution from a shoestring operation to a $180 million institution with an enterprise sales arm. We document the factions now forming over AI and commercialization: volunteers who see the enterprise deals as betrayal, pragmatists who see them as survival, and leadership caught between.

The book ends in early 2026 with a paradox made visible. Grokipedia—a visible fork with attribution, free access, and neo-Nazi sources—drew condemnation from the Foundation. ChatGPT—which provided no attribution, charged for access, and cited Grokipedia as authoritative—drew partnership deals. The visible fork with editorial failures was criticized; the invisible extractor with structural violations was accommodated.

The degradation loop—GPT-5.2 citing Grokipedia citing Wikipedia—was real. But the deeper enclosure was not Grokipedia. It was the absorp-

tion of Wikipedia into proprietary systems: model weights that could not be examined, outputs that carried no attribution, access that required payment. Human curation laundered through AI layers, the value captured by corporations, the commons enclosed.

This is not a polemic against Wikipedia. Wikipedia remains one of humanity's great collaborative achievements. But achievements require protection. And somewhere between the volunteer editors and the AI boardrooms, protection failed.

The question this book asks is simple: Who benefits from Wikipedia's value? The volunteers who created it? The Foundation that hosts it? Or the corporations extracting from it at scale? The answer is not what Wikipedia's founders intended.

A note on sources. We draw primarily from Wikimedia Foundation documents—board minutes, financial reports, legal filings—and from Wikipedia itself: policy pages, deletion logs, arbitration records, talk page debates. Academic research on Wikipedia informs the analysis. News coverage provides context. Memoirs and oral histories from participants supply perspective. Web archives let us see Wikipedia as it was, not just as it is. Every factual claim is cited; readers can verify for themselves.

The people who built Wikipedia deserve a clear-eyed account of what they built and what became of it. So do the rest of us who rely on their work.

San Francisco
January 2026

# Prologue

On October 17, 2025, a new website called Grokipedia appeared online. Its structure was immediately recognizable to anyone who had ever used Wikipedia. The same article titles. The same revision histories. The same editing interface. Grokipedia had copied Wikipedia wholesale—all of it, every article, every edit, every discussion page.

This wasn't unusual. Wikipedia's Creative Commons license explicitly permitted copying. What made Grokipedia different was what it added: a second layer of "sources" drawn from fringe websites, conspiracy theories, and explicitly neo-Nazi content.

Within weeks, researchers at Cornell Tech scraped 883,858 Grokipedia articles—99.8 percent of the entire site. What they found was methodical contamination. The neo-Nazi forum Stormfront appeared as a citation 42 times. The conspiracy site InfoWars appeared 34 times. The white nationalist publication VDare appeared 107 times. In total, researchers counted 12,522 citations to sources they deemed "very low credibility."

The site's entry on Adolf Hitler ran 13,000 words before mentioning the Holocaust by name. Wikipedia mentions the Holocaust in the first paragraph.

Grokipedia's owner, Elon Musk, explained the methodology on a podcast: "The team instructed Grok to compile Wikipedia's top one million articles and make content changes to them." Those changes included citing Stormfront six times in the article for *American History X*, a 1998

film about a reformed neo-Nazi. Wikipedia blacklists Stormfront entirely. Grokipedia treated it as a legitimate source.

The Wikimedia Foundation objected. "Even Grokipedia needs Wikipedia to exist," said Selena Decklemann, the Foundation's Chief Product and Technology Officer. "Wikipedia's content is open source by design; we expect it will be used in good faith to educate."

But objection was all the Foundation could manage. It had no enforcement mechanism, no legal strategy, no way to compel compliance. For four years, the Foundation had signed commercial deals with AI companies that trained on Wikipedia content without meaningful attribution. It had never sued OpenAI, which cited Wikipedia in nearly half its responses while paying nothing for the privilege. It had never demanded that Meta share its Llama model weights under the same copyleft license that governed the Wikipedia text Meta used for training. The Foundation had even admitted, in an April 2024 meeting, that most of its AI customers were "not compliant with the letter of the Creative Commons rules or the spirit of the licenses."

It continued selling to them anyway.

Now Musk was doing publicly what those companies had done privately: treating Wikipedia as raw material for a commercial product. The irony was that Grokipedia, in certain structural ways, was *more* compliant with CC-BY-SA than its critics. It provided visible attribution in a footer linking to Wikipedia. It offered free public access, unlike ChatGPT's subscription tiers. It published Wikipedia-derived content under CC-BY-SA, the same license Wikipedia used. OpenAI did none of these things—and paid nothing for the privilege.

The differences that mattered were editorial, not legal. Grokipedia added neo-Nazi sources. It buried the Holocaust in its Hitler article. Its curation was ideologically corrupted in ways that Wikipedia's community processes were designed to prevent. These were legitimate, serious failures. But they were failures of content quality, not licensing compliance. From

a pure copyleft perspective, the visible fork with attribution was closer to what the license intended than the invisible extraction without it.

---

Three months later, in January 2026, a researcher at the Weizenbaum Institute in Berlin noticed something strange. She had been testing OpenAI's newest model, GPT-5.2, asking it questions about historical events. The model kept citing a source she didn't recognize: Grokipedia.

She tested further. On twelve different queries, GPT-5.2 cited Grokipedia nine times. The model used Grokipedia to answer questions about the Iranian government and the telecommunications company MTN-Irancell. When asked about Sir Richard Evans—the British historian who served as an expert witness against Holocaust denier David Irving—ChatGPT produced what Evans himself described as "debunked information." The trail led back to Grokipedia's contaminated entry.

The loop had closed. Wikipedia's volunteer editors had spent two decades building the internet's most comprehensive encyclopedia. AI companies trained their models on that work. Grokipedia copied Wikipedia and injected neo-Nazi citations. Other AI models then cited Grokipedia, laundering the contamination through the respectable formatting of chatbot responses. The human curation that made Wikipedia reliable—the careful sourcing, the blacklists, the community review—had been bypassed entirely.

A researcher at the Weizenbaum Institute called Grokipedia what it was: "cloaking misinformation."

---

In 1985, a programmer named Richard Stallman watched a company called Symbolics take code he had written, improve it, and refuse to share those improvements back. The experience led him to create the concept of

copyleft: a license designed to ensure that free software stayed free. If you used copyleft code, your modifications had to be shared under the same terms. The license was viral by design. It was supposed to prevent exactly what Symbolics had done—taking community work, adding proprietary value, then locking others out.

Wikipedia launched in 2001 under a copyleft license. The promise was clear: anyone could use Wikipedia's content, but they had to share their improvements. The license was meant to keep knowledge free.

Twenty-four years later, the promise had eroded—but not in the way the prologue's opening might suggest. The real enclosure happened invisibly. OpenAI, Anthropic, Google, and Meta absorbed Wikipedia into proprietary model weights, provided no attribution to end users, charged subscription fees for access, and licensed their outputs under proprietary terms. This was the structural violation of copyleft: taking free content and locking derivatives behind proprietary systems.

Grokipedia was something different. It was a visible fork—public, attributable, re-forkable. Its content failures were severe: the neo-Nazi citations, the Holocaust minimization, the ideological curation. But its structure preserved what copyleft was designed to protect. The content remained accessible. Anyone could check it, criticize it, fork it again.

The copyleft promise failed not because of Grokipedia. It failed because the institution responsible for enforcement chose to partner with the companies performing invisible extraction while criticizing the one performing visible forking. The Foundation had no standing to single out Musk. It had spent years establishing that enforcement was not its practice.

The Wikimedia Foundation collected $180 million in donations annually. It employed over 700 staff. It operated a commercial arm called Wikimedia Enterprise that sold structured data access to Google, Amazon, and Microsoft. And when the moment came to defend the commons its volunteers had built, it issued a press statement.

The statement emphasized Wikipedia's "transparency" and "community model." It noted that "no single individual, company, or agenda can exert influence over the work." It did not threaten legal action. It did not assert licensing violations. It did not demand that Grokipedia remove the neo-Nazi citations or comply with the share-alike provisions that governed Wikipedia content.

The Foundation could not selectively enforce against Musk. It had spent years establishing that such enforcement was not its practice.

---

This book tells the story of how we arrived at this moment. It begins with the launch of Wikipedia in January 2001 and the schism that followed within months. It traces the license switch of 2009, when Wikipedia moved from the GNU Free Documentation License to Creative Commons—a change that may have weakened the copyleft protections volunteers relied upon. It examines the Wikimedia Foundation's growth from a shoestring operation to a $180 million institution, and the commercial pivot that accompanied that growth.

The story builds toward a question that the Foundation's leadership prefers not to answer: When AI companies extract value from volunteer labor without complying with license terms, whose responsibility is enforcement? And a harder question still: What happens when the institution responsible for enforcement draws its staff from overlapping professional circles, shares a common political orientation, and views tech companies as colleagues rather than adversaries?

The volunteers have no standing to sue. The Foundation has the resources but not the will. The AI companies have lawyers who argue that training on copyleft content is fair use, that model outputs are not derivative works, that attribution requirements do not apply to statistical patterns in neural networks.

These arguments may be wrong. Courts have not yet ruled definitively on whether AI training constitutes fair use of copyleft content. The memorization studies suggest that models do not merely learn patterns—they reproduce Wikipedia text nearly verbatim in many cases. One study found that ChatGPT's responses cited Wikipedia nearly half the time, often without attribution.

But the legal question has become academic. The Wikimedia Foundation chose not to fight. It chose to accommodate. It chose to sell access to the same companies whose compliance it doubted.

And then Grokipedia arrived—a visible fork with attribution, free access, and neo-Nazi citations. The Foundation objected to the content failures. It could not object to the structure, which was more compliant than the invisible extraction it had accommodated for years. The irony was complete: the Foundation condemned the visible fork while partnering with the invisible enclosers.

---

This is not a story about technology. It is a story about institutions, about money, about who controls knowledge and who profits from it.

Wikipedia's quarter million volunteer editors built something remarkable: a free encyclopedia that rivals anything published before it. They did this work without pay, often anonymously, on the promise that their contributions would remain free for anyone to use and improve.

That promise has been broken. Not by some outside force, but by the institution that was supposed to uphold it.

The question this book asks is simple: How did we get here?

The answer involves a founder who left in anger, a license switch that may have been legally dubious, a nonprofit that learned to think like a corporation, and an AI industry that discovered it could take without asking.

By January 2026, the consequences were visible to anyone willing to look. Neo-Nazi content laundered through AI systems. Human curation rendered invisible. The commons enclosed not by the visible fork that critics condemned, but by the invisible extraction that the Foundation had enabled for years.

The degradation loop—ChatGPT citing Grokipedia citing Wikipedia—was real. But the deeper enclosure was structural: volunteer labor absorbed into proprietary systems, attribution stripped, access paywalled. A biased public encyclopedia was still public. A "neutral" proprietary black box was enclosure.

This is not where the story was supposed to end.

# Part I

# Foundations and Fractures

CHAPTER ONE

# The Expert Failure

> The idea that a few smart people can sit down and design everything is wrong. Wikipedia works because it relies on the distributed knowledge of millions.
>
> — Jimmy Wales

In September 2000, after months of work, an article about the musical concept of atonality finally cleared its seventh and final review stage. The author, Christoph Hust, had written something any academic journal would accept: 14 footnotes, 27 references, a complete discography of atonal compositions. The prose was careful, the scholarship thorough. By every measure of traditional expertise, it was exactly what an encyclopedia should contain.

This was the dream of Nupedia, the free online encyclopedia that Jimmy Wales and Larry Sanger had launched six months earlier. Expert authors would produce expert content. A rigorous seven-step peer review process would ensure quality. The articles would be free to read, yes, but they would carry the weight of credentials. The internet would finally have an encyclopedia worthy of the name.

There was just one problem. After nine months of operation and roughly $250,000 in investment from Wales's company Bomis, Nupedia

had produced exactly two finished articles.[1] The "Atonality" article was one of them. At roughly $10,000 per article, the economics were absurd. At roughly one article every four months, the timeline was worse. The internet was growing at a pace that demanded thousands of articles, not dozens. The experts weren't showing up. The process wasn't scaling. The vision was dying.

What happened next would change how humanity organizes knowledge. But to understand why, we need to understand what went wrong.

## The Predecessors

The idea of a collaboratively written, freely available encyclopedia predated Nupedia by nearly a decade. It just never worked.

In 1993, a computer scientist named Rick Gates proposed something called Interpedia. The concept was simple: volunteers would write encyclopedia articles and share them over the early internet. Over 700 messages accumulated on the discussion mailing list. Fewer than 50 draft articles ever materialized. The project died before it could demonstrate whether collaboration at scale was even possible. The tools didn't exist. The culture wasn't ready.

Six years later, Richard Stallman—the MIT programmer who had pioneered the free software movement—proposed his own version. Stallman called it GNUpedia, positioning it as part of his broader crusade for digital freedom. He had already created the GNU General Public License, which required anyone who modified free software to share their modifications freely. Now he wanted to apply the same principle to knowledge. His proposal included what would become crucial legal infrastructure: the GNU Free Documentation License, or GFDL, which would later govern Wikipedia's content.

[1] Article counts vary across sources. Sanger's 2005 memoir states "21 articles" by the end of 2000, but other accounts suggest only 12 to 24 were ever completed before Nupedia's closure in 2003. The discrepancy may reflect different definitions of "complete" versus "in progress."

GNUpedia launched in January 2001—the same month as Wikipedia. Within weeks, the projects merged efforts, with GNUpedia contributors migrating to the faster-moving wiki. Stallman's licensing framework survived even as his encyclopedia did not.

The pattern was clear. Everyone agreed the internet needed a free encyclopedia. No one had figured out how to build one.

## Bomis

Jimmy Wales came to encyclopedias through an unusual path. In 1996, he and two partners—Tim Shell and Michael Davis—founded a web company called Bomis. The site functioned as a web portal, helping users find content in the chaotic early internet. Some of that content was adult material, and "Bomis Babes" generated enough advertising revenue to keep the company afloat.[1]

Wales's intellectual interests, however, ran in different directions. As an undergraduate at Auburn University in Alabama, he had studied economics with Mark Thornton, a scholar in the Austrian tradition. Thornton introduced Wales to Friedrich Hayek's work, particularly a 1945 essay titled "The Use of Knowledge in Society."[2]

Hayek's argument would shape everything Wales built. The essay contended that economic planning fails because the knowledge required for rational decisions is never concentrated in one place. It exists scattered across millions of individuals, each with their own fragment of understanding. No committee, however expert, can aggregate this dispersed knowledge efficiently. Only decentralized systems—markets, in Hayek's view—can coordinate what individuals know into coherent outcomes.

Wales also absorbed Eric Raymond's 1997 essay "The Cathedral and the Bazaar," which applied similar reasoning to software development.[3] Raymond contrasted the "cathedral" model of software—carefully planned by a small team of experts—with the "bazaar" model of open source, where anyone could contribute and the best code emerged from thousands of

competing attempts. "Given enough eyeballs," Raymond famously wrote, "all bugs are shallow."

By October 1999, Wales had begun conceptualizing an online encyclopedia. The question was how to build it.

## Sanger Arrives

Larry Sanger was finishing his PhD in philosophy at Ohio State University when he met Wales through online discussion groups. His specialty was epistemology—the study of knowledge itself. How do we know what we know? What distinguishes justified belief from mere opinion? These questions had occupied philosophers for centuries. Sanger approached them with academic rigor.

In May 1999, Sanger joined Bomis. By January 2000, Wales had hired him to develop and run Nupedia.

The philosophical tension between them was present from the start. Wales believed that good processes could substitute for credentials. Sanger believed that credentials signaled reliability precisely because knowledge requires justification, and justification requires expertise. Wales trusted emergence. Sanger trusted verification.

For Nupedia, Sanger designed a process that reflected his training. If knowledge required justification, then articles would need to be justified—by experts, through review, at every stage. The result was a seven-step gauntlet that looked exactly like academic peer review because, in Sanger's mind, academic peer review was how knowledge got validated.

Neither founder knew it yet, but they had planted the seeds of an irreconcilable conflict.

## Seven Steps

Sanger's review system required seven sequential stages. First, a topic was assigned to a qualified author—ideally someone with a PhD in the relevant field. Second, an expert reviewer was recruited, also requiring a

PhD "with few exceptions." Third, the lead reviewer checked accuracy and completeness, provided substantive feedback, and engaged in extended discussion with the author through a web forum. Fourth, other experts reviewed and commented, incorporating multiple perspectives in what Sanger compared to the question-and-answer portion of an academic conference. Fifth, a lead copyeditor improved grammar, punctuation, and layout. Sixth, additional copyeditors contributed refinements. Seventh and finally, the article received formal approval and was prepared for publication.

The structure made sense if you assumed that encyclopedia articles should work like academic papers. Each stage guaranteed scrutiny. Each reviewer added legitimacy. By the time an article cleared all seven steps, it had earned the right to be trusted.

The problem was that each step blocked all subsequent steps. Nothing could move to stage three until stage two finished. Nothing could move to stage four until stage three finished. The system had no parallel processing. A single delay at any stage halted everything downstream.

And delays were constant.

## The Absent Experts

The seven-step process assumed that qualified experts would volunteer their time. They didn't.

The reasons were structural, not personal. Academic culture rewards publication in peer-reviewed journals, which counts toward tenure and promotion. Encyclopedia writing does not. A professor who spent hours reviewing a Nupedia article would have nothing to show for it professionally. The work wouldn't appear on a CV. Colleagues wouldn't cite it. Tenure committees wouldn't care.

Nupedia offered no compensation. The project depended entirely on intrinsic motivation—the satisfaction of contributing to something worthwhile. But intrinsic motivation works differently for tenured professors

than for hobbyists. Professors already have outlets for intellectual contribution. Hobbyists often don't. The people most qualified to write expert encyclopedia articles were precisely the people with the least incentive to do so.

The committee structure made things worse. Academics are used to working alone or in small collaborations. They're not used to submitting their prose to seven sequential stages of review by strangers. The process felt burdensome rather than validating. It demanded more effort than academic publishing while offering fewer rewards.

"By the summer of 2000," Sanger would later write, "it had become clear that the process we tested out was very slow."[4]

### The Numbers

The raw statistics told the story. In its first year, Nupedia completed just 21 articles (some sources cite only 12). By the time of closure, only 24 articles had been fully completed, with 74 still in progress. Bomis had invested roughly $250,000—approximately $10,000 per finished article.

For comparison: Wikipedia, once launched, would generate hundreds of articles in its first two weeks. Within a year, it would have twenty thousand.

The "Atonality" article that Christoph Hust labored over for months was exactly what Nupedia wanted. It was also exactly what Nupedia couldn't afford. At that rate, building an encyclopedia of useful size would cost billions. At one article per four months, it would take centuries.

Something had to change.

### The Deeper Problem

The conventional explanation for Nupedia's failure is that it was too slow. This is true but incomplete. The deeper problem was epistemological.

Nupedia assumed that knowledge production works the way academic publishing works: a small number of credentialed experts create content,

and peer review validates it. This model functions in academia because journals have prestige that attracts authors and reviewers despite the lack of payment. Encyclopedia articles had no comparable prestige. The incentive structure was broken from the start.

But Wales had absorbed a different theory. Hayek's 1945 essay argued that central planning fails because no single authority can possess all the relevant knowledge. "The peculiar character of the problem of a rational economic order," Hayek wrote, "is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess."[2]

Hayek's argument was fundamentally about governance. Centralized expertise fails not because experts are stupid, but because knowledge is genuinely dispersed across millions of people. No single authority—no expert committee, no foundation, no corporation—can gather all that knowledge and make it coherent. Only a system that lets many knowledge-holders contribute simultaneously can create order without central planning. This was the dream of a commons: not the absence of rules, but rules created collectively by those with a stake in the resource.

Applied to encyclopedias, this suggested something radical. Maybe the problem wasn't finding enough experts. Maybe the problem was relying on experts at all. Maybe knowledge could emerge from the coordination of thousands of non-experts, each contributing their small piece, with the aggregate exceeding what any expert committee could produce.

This was heresy to Sanger's epistemology. Knowledge requires justification. Justification requires expertise. Without credentials, how could you trust what someone wrote?

But Hayek's model suggested a different kind of trust: trust in process rather than credentials. If enough people contributed, and if good pro-

cesses existed to correct errors, then the quality of individual contributors might matter less than the scale of participation.

The question was whether such a process could exist. In January 2001, Sanger would stumble across one.

## The Dinner

On January 2, 2001, Sanger met his friend Ben Kovitz for dinner at a Mexican restaurant in San Diego. Kovitz was a computer programmer Sanger knew from online mailing lists. They ordered food. Sanger had enchiladas.

Kovitz had been spending his spare time on something called WikiWikiWeb, a website created by programmer Ward Cunningham in 1995. The name came from the Hawaiian word for "quick"—Cunningham had named it after the Wiki Wiki Shuttle bus at Honolulu International Airport. The concept was simple but strange: any visitor could edit any page, and changes appeared immediately. There were no gatekeepers, no review stages, no credentials required.

Cunningham's site had grown into a thriving community of software developers. They wrote about programming patterns, debated design choices, and refined each other's contributions. When someone made an error, someone else fixed it. When vandals appeared, the community reverted their changes. Against all intuition, the system worked. The wiki self-corrected.

Sanger listened with growing interest. "At some point I said, 'You know, I just wonder how this might solve the problem I've been having with Nupedia,'" he later recalled. "I was skeptical, obviously, because it doesn't sound plausible at first, but Ben was able to answer a lot of my objections. The best part of it was that the software was easy to install."[5]

That evening, Sanger and Kovitz went to Sanger's apartment. Sanger wrote a formal proposal and sent it to Wales. The proposal suggested using wiki software to accelerate content creation for Nupedia—a "feeder"

system where rough drafts could be developed quickly before entering the seven-step review.

Wales approved the experiment.

Kovitz later described his role with characteristic modesty: "I was only the bumblebee. I had buzzed around the wiki flower for a while, and then pollinated the free-encyclopedia flower."

## The Rejection

On January 10, 2001, Clifford Adams's UseModWiki software was installed on the Nupedia servers. The wiki section launched as a test, positioned as a drafting area that would feed content into the main review process.

The Nupedia Advisory Board was not pleased.

These were credentialed scholars—the very experts Sanger had recruited to give Nupedia legitimacy. They looked at the wiki and saw chaos. Anyone could edit? No review before publication? No verification of credentials? This wasn't an encyclopedia. This was a threat to everything Nupedia stood for.

"A clear majority of the Nupedia Advisory Board wanted to have nothing to do with a wiki," Sanger recalled. "Their commitment was to rigor and reliability."

Their objections were not irrational. Academic knowledge production had developed peer review for good reasons. Errors get caught before publication. Cranks get filtered out. Quality emerges from scrutiny. A wiki inverted all of this. Errors would be published first and corrected later (maybe). Credentials wouldn't matter. Quality would depend on the crowd, and crowds are notoriously unreliable.

Sanger, though he had proposed the wiki, understood their position. "They were perfectly reasonable to doubt that it would turn into the fantastic source of content that it did."

The Board's rejection forced a decision. The wiki couldn't be part of Nupedia. It would have to be something else.

## Wikipedia Launches

On January 15, 2001, the wiki launched as a separate project under a new name: Wikipedia. Wales typed "Hello, World!" as the first edit.[6]

The name came from combining "wiki" with "encyclopedia." The domain was wikipedia.com (later changed to wikipedia.org to emphasize its nonprofit nature). The software was UseModWiki, the same tool Kovitz had told Sanger about. The licensing was GFDL, inherited from Stallman's free software movement.

And the articles started appearing.

Within days, Wikipedia had more content than Nupedia had produced in a year. By the end of January, roughly 600 articles existed. Contributors arrived from nowhere—hobbyists, enthusiasts, anyone curious enough to click "edit." They wrote about their interests, corrected each other's errors, and built something that no committee could have planned.

But this same openness—the feature that made Wikipedia powerful—would eventually make it indefensible. An encyclopedia written by anyone could be read by anything. An archive built to be freely shared could not easily be restricted from free-taking.

The wiki magic that Kovitz had described was real. When someone contributed something valuable, others refined it. When someone contributed something wrong, others fixed it. When vandals appeared, they were reverted within minutes. The process wasn't perfect, but it was fast. And speed, it turned out, mattered more than perfection.

By year's end, Wikipedia had grown a hundredfold. By the end of 2002, it had passed 100,000 articles. The "feeder" system had become the main course.

### What They Saw

The Nupedia Advisory Board understood something important: the wiki wasn't just a technical change. It was a philosophical revolution.

Nupedia treated knowledge as something that experts produce and gatekeepers validate. Wikipedia treated knowledge as something that emerges from process—from the iterative contributions of thousands of people, most of them anonymous, none of them credentialed. The Board recognized that these were incompatible visions. Associating with a wiki would undermine everything Nupedia represented.

They were right about the stakes. They were wrong about which model would win.

What they failed to anticipate was that the internet had changed the economics of knowledge production. Nupedia's model required finding experts who would work for free despite having no professional incentive to do so. Wikipedia's model tapped a different pool: people who wanted to contribute to something larger than themselves, who had passion for their topics, who found satisfaction in building a shared resource. These people existed in the millions. Academics with PhDs existed in the thousands.

Scale beat credentials. The Board couldn't see it coming. Neither could Sanger, really. The wiki was supposed to be a tool for accelerating the real work of building Nupedia. Instead, it replaced Nupedia entirely.

### The Aftermath

Nupedia lingered for two more years. After 2001, only two more articles completed its seven-step review. The community's energy had shifted to Wikipedia, where a single contributor could write and publish in an afternoon what Nupedia took months to produce. The expert model couldn't compete.

On September 26, 2003, Nupedia's servers were shut down for good. The project had produced barely two dozen completed articles—the same number a single Wikipedia editor might write in a month. That same

day, Wikipedia had over 350,000 articles in multiple languages, created by contributors who held no credentials other than a willingness to click "edit."

Sanger had left the projects 18 months earlier, in March 2002. His departure would become its own controversy, which we'll examine in Chapter 3. But the philosophical battle was already lost by the time he left. Wikipedia had chosen Wales's model over Sanger's. The community had voted with their keystrokes, and they voted for velocity over verification.

"The original idea for Nupedia was to create an encyclopedia written by experts, for everyone," Sanger later reflected. "Wikipedia was supposed to be a feeder system. Instead, the feeder ate the host."

## The Founding Choice

Christoph Hust's "Atonality" article, with its 14 footnotes and 27 references, represented the last gasp of a model that valued credentials over contribution. Within months of Wikipedia's launch, anonymous contributors had written their own article on atonality—less scholarly, perhaps, but read by millions.

The choice Wikipedia made in January 2001 had consequences that extend to the present.

The first consequence was vulnerability to extraction. Content created without credentialed gatekeepers is harder to defend as proprietary expertise. When AI companies later scraped Wikipedia for training data, they inherited content whose creation process didn't privilege any single source of authority. The openness that defeated Nupedia's gatekeeping would later become "anyone can scrape."

The second consequence was dependence on volunteer labor. The anti-elitist model relies on mass participation, making the project resilient to individual departures but vulnerable to systematic extraction by commercial entities. The volunteer army that builds Wikipedia has no employment contract, no union, no collective bargaining power.

The third consequence was the enforcement problem. When the Wikimedia Foundation later faced AI companies using Wikipedia content without attribution, the distributed, non-expert model made it difficult to argue special harm. The project chose accessibility over defensibility.

A fourth consequence emerged more gradually: the composition of the institution that would steward this open project. Wikipedia's anti-credentialist philosophy did not determine who would staff its Foundation. That would depend on hiring networks, geographic sorting, and cultural fit—dynamics that would produce their own kind of homogeneity.

This founding tension—between openness and control, between process and credentials, between velocity and verification—has never been resolved. Every subsequent crisis Wikipedia faced traces back to this original choice. The Spanish Wikipedia fork of 2002 (Chapter 2) showed the first time the community would fracture over these questions. Sanger's departure (Chapter 3) made the philosophical split permanent. And when AI companies began extracting Wikipedia at industrial scale two decades later, the project's founders would discover that the openness they celebrated had left them exposed.

The question was no longer whether the open model could work. It had worked spectacularly. The question was what it would cost—and who would pay it. For two decades, that question would linger, dormant, until the answer became clear: the commons would pay.

CHAPTER TWO

# The Spanish Fork

> The international Wikipedia that you all know and have come to take for granted might have been impossible without the Spanish fork.
>
> — Edgar Enyedy, 2011

In mid-February 2002, Edgar Enyedy sat down to write a farewell message. For six months he had devoted himself to building the Spanish Wikipedia—eight to twelve hours a day, every day, creating articles, developing a style guide, designing an index system based on the Universal Decimal Classification. He had a master's degree in communications systems from the Polytechnic University of Madrid, training in philology and computer science, a background in journalism. He had poured all of it into Wikipedia.

Now he was leaving.

The message he wrote was brief, but its final line became legendary in Wikipedia's early history: "Good luck with your WikiPAIDia."

The neologism captured everything wrong in a single word. Wiki plus PAID. Free knowledge corrupted by money. But the anger went deeper than just advertising. Enyedy had volunteered thousands of hours to build Spanish Wikipedia. He had contributed to something he thought was a commons—a shared knowledge resource governed by and for its

contributors. Now he discovered he had been building value for someone else's company. His labor, Bomis's asset. In the language of commons governance, what he'd discovered was enclosure: the appropriation of shared resources by a private owner.

Enyedy wasn't alone. When he announced the fork, questions flooded in: "And now? What's next?" The answer was Enciclopedia Libre Universal en Espanol—the Free Universal Encyclopedia in Spanish. A new project, hosted at the University of Seville, free from commercial control forever.

Most of the Spanish Wikipedia's active contributors followed him out the door.

## First Success

The Spanish Wikipedia launched on May 20, 2001, just four months after the English version. By early 2002 it was one of the largest non-English Wikipedias, proof that the wiki model could work across languages and cultures.

Enyedy and Javier de la Cueva, a lawyer specializing in technology and internet issues, emerged as its unofficial leaders. They didn't just write articles. They built infrastructure—policies, style guides, organizational systems. They communicated with Wales and Sanger in San Diego about the challenges facing international Wikipedias. They tried to make the project work.

But the challenges were substantial, and many had nothing to do with content.

## Bomis Control

Wikipedia was hosted by Bomis, Jimmy Wales's for-profit company. All the domains—wikipedia.com, wikipedia.org, wikipedia.net—belonged to Wales personally. The servers sat in San Diego. The software updates

were controlled by Bomis staff. There was no formal governance structure separating editorial decisions from commercial interests.

For the English Wikipedia, this ambiguity was tolerable. The English editors worked closely with Wales and could negotiate directly. For international Wikipedias, the arrangement felt colonial. They were building content for an American company, subject to decisions made in a language many of them didn't speak, with no formal say in the project's direction.

Enyedy later enumerated his grievances. They went far beyond advertising.

"I asked myself 'why are we working for a dot-com?'" he recalled in a 2011 interview. "I asked for Wikipedia to be changed to a dot-org."[7]

He objected to Larry Sanger's role as editor-in-chief: "I wanted the Big Brother out. Larry Sanger was against the nature of the project itself. None of us felt comfortable with such a figure."

He pushed for autonomy: "We did not want to be seen as mere translations of the American version... I was told so many times to translate from the main wiki, and my response was always the same: We are not a translation of the American Wikipedia!"

### Technical Grievances

Beyond governance, there were practical problems. Software updates went to the English Wikipedia first. International versions ran outdated code for months. When Enyedy requested access to the Spanish server to help maintain it, his request was denied. "They said it was for security reasons because Bomis Inc. was hosting files from its clients on the same server."

He asked for mirror servers—backup copies that would ensure the Spanish content survived independently. The answer was always the same: the project needed to stay together. Database exports, which would have allowed contributors to easily copy their work elsewhere, weren't being updated. The GNU Free Documentation License that governed

Wikipedia's content promised that anyone could fork the project and create their own version. In practice, the technical barriers made forking nearly impossible.

"The Wikipedia page on Sourceforge had instructions that read like hieroglyphics," Enyedy said. "And once again due to 'technical' reasons that none of us believed, the downloadable database was never updated... These conditions did not resemble what the GNU/FDL was supposed to ensure."

The license promised freedom. The infrastructure delivered dependency.

Here was the betrayal beneath the surface. The GFDL was supposed to guarantee freedom—if Wikipedia became proprietary, anyone could fork it. Copyleft was the shared resource's insurance policy. But Bomis controlled the infrastructure. Database exports weren't updated. The GFDL protected your legal right to fork. The servers protected Bomis's practical ability to prevent it. This was the first lesson about digital knowledge projects: legal rights mean nothing if you don't control the technology.

## The Trigger

In late 2001, the internet bubble collapsed. Bomis lost major contracts. Revenue dried up. By December, Tim Shell informed Larry Sanger that he would be laid off. Wikipedia was costing money Bomis no longer had.

On February 2, 2002, Jimmy Wales posted to the Wikipedia-L mailing list: "However, with the ongoing hard times in the Internet economy, we do anticipate adding some forms of advertising to the site in the near future."[8]

Sanger, now working unpaid, mentioned that advertising revenue "might make it possible for me to come back to my old job."

For Enyedy, this was confirmation. Everything he feared about Bomis's commercial control was coming true. The advertising announcement

wasn't a surprise—it was the last straw on a pile of grievances that had been accumulating for months.

"The possibility of advertising was out of the question," he later said. "I asked Wales for a public commitment that there would be no advertising. This only came after we left."

## The Execution

On February 26, 2002, Enciclopedia Libre went live. Juan Antonio Ruiz Rivas had arranged hosting at the University of Seville—academic, non-commercial, permanently outside Bomis's control.

The GFDL made forking legally possible. But the practical barriers remained. The database exports weren't current. There was no automated way to migrate content. The Spanish contributors had to copy their articles one by one.

"When the server was up and running, and as the GNU/FDL permitted, we began copying our articles from Wikipedia," Enyedy recalled. "It wasn't an automated process, no bots or anything, just us bringing the articles across one by one from Wikipedia's server to ours."

The exodus was swift. Enyedy "persuaded most of the volunteers to go with him." The Spanish Wikipedia, which had been one of the most active non-English versions, was gutted. Andrew Lih, in his history of Wikipedia, noted that the Spanish edition was "virtually inactive until mid-2003." It took "more than a year for the Spanish Wikipedia to get back on its feet again."[9]

## Wales Responds

On March 1, 2002, Larry Sanger formally resigned as Wikipedia's chief editor. The same day, Wales sent an email to Wikipedia-L with a significant reversal:

"With the resignation of Larry, there is a much less pressing need for funds. Therefore, all plans to put advertising of any kind on the wikipedia is called off for now."[10]

The qualifier—"for now"—was notable. Wales continued: "Just as the National Geographic Society is supported in large part by advertisements in the National Geographic Magazine, I expect this to be a potentially necessary thing at some point in the future, if we wish to have an impact beyond our own little corner of the Internet."

Even in 2006, Wales told ClickZ: "We have never said there would absolutely never be ads on Wikipedia."

But for the immediate crisis, the concession was enough. Advertising was off the table. The Spanish forkers had won their first demand.

In August 2002, Wikipedia changed its domain from .com to .org, signaling non-commercial intent. In October, a Wikipedian named Daniel Mayer (known as Maveric149) led a reunification effort, hoping to bring the Spanish contributors back. The Enciclopedia Libre community held a "Retorno a Wikipedia" vote. They rejected reunification.

On June 20, 2003, the Wikimedia Foundation was incorporated as a 501(c)(3) nonprofit in Florida. Wikipedia was formally separated from Bomis. The project would be governed by a foundation, not a company.

The sequence was clear: Fork, crisis, concessions, Foundation. Enyedy was right to claim credit. Wales feared other language communities would follow the Spanish example. He learned, as Enyedy put it, "what to do and what not to do."

### Founder Dispute

Years later, the advertising controversy would generate its own controversy. In January 2011, when Wired UK published a story about the Spanish fork, Wales wrote in the comments section:

"Sanger was absolutely adamant that Wikipedia must have ads, and it was my refusal to do so that led to Wikipedia being as it is today. The

Spanish fork did not provoke any changes of any kind. We stayed the course."

Sanger, who had long since broken with Wales, responded on his blog:

"The suggestion that I demanded ads and that Jimmy Wales was opposed to them is, I am afraid, yet another self-serving lie from Wales... From the beginning, Wales let me know in no uncertain terms that, once it garnered enough traffic, Nupedia would become ad-supported."[11]

Sanger cited Wales's February 2, 2002 email as evidence. He noted that Wales had deleted a Wikipedia discussion where editors pressed him on the discrepancy between his 2011 claim and the historical record.

The mailing list archives contradicted Wales's later account. He had announced advertising plans; the community had revolted; he had reversed course. These were documented facts. But the dispute revealed something deeper: the founders themselves couldn't agree on what they had intended. No governance structure existed to settle such questions. Everything depended on the memory and goodwill of the people who controlled the servers.

This was the deeper lesson of the fork. The Spanish contributors had fought about advertising because advertising was visible and concrete. The underlying problem—no governance mechanism to constrain central authority—remained even after the Foundation was created. Two decades later, when Enterprise revenue created conflicts of interest with copyleft enforcement, the same gap would matter again.

### Principled Dissent

To Enyedy and the Spanish forkers, the revolt was about principles, not convenience. Volunteers should not build value for a for-profit company without governance guarantees. The GFDL's promise of freedom means nothing if technical barriers prevent forking. Commercial control was the opposite of what free knowledge required.

"We were basically working for Bomis Inc., and asked in a gentle way to translate from the main Wikipedia," Enyedy summarized. "Finally, came the possibility of incorporating advertising, so we left. It couldn't be any other way."

The fork succeeded as pressure even if Enciclopedia Libre didn't survive long-term. The concessions—no advertising, .org domain, the Foundation—proved the leverage was real.

### Continuity Narrative

From Wales's perspective, and the official Wikimedia Foundation history, the fork was a misunderstanding resolved. Wikipedia was always intended to be ad-free; the February 2002 announcement was an aberration caused by financial pressure, not a statement of intent. The Foundation's creation reflected natural evolution, not forced concession. The project "stayed the course."

The evidence cited: Wikipedia's current ad-free status, the Spanish Wikipedia's eventual recovery and growth to millions of articles, Enciclopedia Libre's decline and 2024 closure. The fork failed; Wikipedia prevailed.

### Prophetic Warning

A third reading sees the fork as prophecy. Enyedy wasn't just worried about advertisements—he was identifying a pattern that would define Wikipedia's next two decades. A shared resource is only stable if its contributors believe they control it. The moment a central authority takes it for profit, the project's social contract collapses.

The forkers were right about the threat. They were wrong only about its form and timeline.

They feared advertising. They could not have imagined Enterprise API sales to AI companies. Nor could they anticipate a subtler capture: that the institution created to protect against commercialization would be staffed

entirely by people who saw tech companies as natural allies. They won concessions against banner ads—the visible symbol of commercialization. They established no protections against extraction through training data and structured APIs.

Enyedy's 2011 warning proved prescient: "This is what I didn't want to happen: a large, money-centred organisation made possible by the free work of the community." When he said this, the Foundation's annual budget was around $20 million. By 2024, it exceeded $180 million.

The structural parallels are striking. In 2002, Spanish editors asked "Why are we working for a dot-com?"; in 2025, volunteers ask "Why is our labor training commercial AI?" In 2002, no database export was available; in 2025, API access is sold to Enterprise customers. In 2002, software was inequitable across languages; in 2025, Enterprise customers get structured data while others scrape. In 2002, Enyedy warned against commercialization; in 2025, the Foundation's annual budget approached two hundred million dollars. In 2002, editors contributed freely; in 2025, a quarter million volunteers receive nothing from AI companies using their work.

The fork that almost was became the prophecy that came true—twenty-two years late, through mechanisms no one anticipated.

### Long Aftermath

In 2011, when Nathaniel Tkacz interviewed Enyedy for an academic volume on Wikipedia, the forker offered a clear-eyed assessment of what he had accomplished:

"The fork had its time and place, its goal and its consequences. Nowadays, the romantic point of view is that EL survived and is still going strong. It is a nice view, but wrong. EL has failed as a long-term project for one reason: The project itself was not intended to last. It was merely a form of pressure. Some of the goals were achieved, not all of them, but it was worth the cost."

The goals achieved were significant: Wikipedia became ad-free, moved to .org, established the Foundation, gave international Wikipedias more autonomy, improved database exports, made software available on Sourceforge. The Spanish Wikipedia eventually recovered and thrived.

But the structural vulnerability Enyedy identified was never resolved. The institution grew; its incentives diverged from those of the volunteers who built it.

The Foundation created to protect against commercialization grew into exactly what he feared. Not through advertising—he had won that battle. Through institutional growth, through staff expansion, through budget inflation, and ultimately through Enterprise sales to the companies building AI on volunteer labor.

## Quiet End

In late 2024, Enciclopedia Libre finally ceased operations. Twenty-two years after its founding, the fork that almost was ended with a whisper. By then, the Spanish Wikipedia had millions of articles. Wikipedia had won.

But by then, the fears that motivated the fork were no longer hypothetical. Wikipedia had not run advertising. It had done something the Spanish forkers never imagined: it had become the training data for a new generation of commercial products, and the Foundation created to protect it was taking the money.

Enyedy, living in a small seaside town, steering clear of public life, had been right all along. He was just two decades early.

## Theme Stated

The Spanish fork stated the question the entire book will answer: Who controls ”free” knowledge, and what happens when money enters the equation?

In 2002, the answer seemed clear: volunteers could threaten exit, and the institution would accommodate their principles. The fork demonstrated that the GFDL enabled competition—Bomis couldn't lock contributors in. This was the promise of copyleft.

But the promise had limits the forkers couldn't see. They could prevent advertising because advertising was visible. They could not prevent forms of extraction that hadn't been invented yet. When AI companies began scraping Wikipedia for training data, the same openness that enabled the fork enabled the scraping. When the Foundation created Enterprise to monetize that access, the same governance structures that protected against Wales's advertising protected nothing against the Foundation's revenue imperatives.

The fork showed that Wikipedia's community could act decisively when commercialization threatened. But it also showed the limits of such action: the volunteers won on threats they could see. Threats they could not see would arrive later.

Sanger's departure (Chapter 3) would remove the governance tensions the fork exposed. The project moved forward without resolving them. Those tensions would return in the 2020s—not as founder disagreement, but as financial dependency on AI companies. The volunteers who built Wikipedia would discover, as Enyedy had discovered, that the money always finds a way in.

CHAPTER THREE

# The Departure

> The root problem: anti-elitism, or lack of respect for expertise. As a community, Wikipedia lacks the habit or tradition of respect for expertise.
>
> — Larry Sanger, 2004

Larry Sanger waited years before admitting the truth.

His formal resignation, posted on March 1, 2002, said all the right things. "This has nothing to do with my lack of faith in the projects," he wrote. His relationship with Jimmy Wales and Bomis "remains friendly." The departure was purely financial—the dot-com bubble had burst, and Bomis could no longer afford his salary.

This was true as far as it went. But the fuller confession, written months later on a quiet corner of Meta-Wiki, told a different story:

"I must now confess, when I had to quit as 'organizer' of Wikipedia at the end of February 2002, that it was a relief."[12]

A relief. The only paid Wikipedia employee. The man who named the project, designed its policies, served as its chief editor for its entire first year. And leaving felt like escape.

"Very much of my time was spent in defending myself from such attacks," he continued. The attacks came from within—from other Wikipedians who opposed his vision of how the project should work. "Several

different people, not all at once, began a concerted effort to criticize and undermine virtually every 'management' action."

The wiki structure worked against him. "Critics in a wikiwiki environment cannot be silenced... no criticism, no matter how false or misleading, can be erased without an accusation of censorship." Every dispute escalated. Every editorial decision became a referendum on his authority. By the time Bomis informed him the funding would end, the poisonous atmosphere had already driven him to the brink.

What happened in those final months? What tension could make the creator of Wikipedia's policies desperate to leave his own creation?

## The Partnership

In January 2001, when Wikipedia launched, Larry Sanger was its only paid employee. Jimmy Wales provided funding through Bomis and made high-level decisions, but Sanger was the one building the project day by day. He wrote or commissioned the first articles. He formulated the Neutral Point of View policy that would become Wikipedia's philosophical foundation. He moderated disputes, recruited contributors, and shaped the community's norms.

The press releases from those early months credited both men as founders. In September 2001, the New York Times first described Sanger as "co-founder"—a description he had helped craft. The relationship with Wales appeared collaborative. Two men with different skills, building something together.

But the philosophical tension planted in Chapter 1 had never resolved. Wales believed in distributed knowledge, emergent order, process over credentials. Sanger believed expertise should have standing, that credentials signaled reliability, that quality required some form of gatekeeping. Wikipedia had chosen to be open. Sanger kept trying to impose structure on that openness.

For a while, both visions coexisted. Then they didn't.

### Atmosphere Curdles

By fall 2001, something had changed. Wikipedia was growing explosively—20,000 articles by year's end—but Sanger's authority was shrinking.

The attacks weren't coordinated, but they shared a common theme: Sanger was too controlling. His attempts to maintain standards were authoritarianism. His editorial judgments were elitism. Every time he tried to enforce quality, someone accused him of violating Wikipedia's open culture.

The structural problem was that Wikipedia had no mechanism to resolve such disputes. In a traditional organization, Sanger's title as editor-in-chief would have settled things. In Wikipedia's radical flatness, titles meant nothing. Anyone could criticize anything. And in a wiki environment, criticism couldn't be removed without triggering accusations of censorship.

"It's now generally the case that if I attempt to do or say anything that is very controversial—certainly anything regarding Wikipedia policy and politics—I am either ignored or, more often, vociferously denounced," Sanger wrote. The community he had helped build was rejecting him.

### Troll Problem

Wikipedia's openness attracted both contributors and disruptors. Sanger described a "terrible dilemma"—the same radical accessibility that powered Wikipedia's growth also empowered those who wanted to damage it.

"Trolls" in early internet parlance were users who disrupted communities for entertainment. They started fights, inserted false information, and wasted contributors' time with bad-faith arguments. A traditional encyclopedia could simply ban them. Wikipedia's commitment to openness made that much harder.

Sanger tried to address the problem. He proposed giving experts more authority. He suggested mechanisms to identify and remove troublemakers. Each proposal met the same objection: you're betraying Wikipedia's

principles. The project's most active contributors were "decidedly anti-elitist," he observed. They saw his attempts to impose structure as attacks on the open model that was working.

The irony was bitter. Sanger had designed many of Wikipedia's policies. Now those policies were being used against him.

### Financial Trigger

In late 2001, the internet economy collapsed. Bomis, which had been funding Wikipedia through advertising revenue, lost major contracts. Tim Shell informed Sanger in December that the company could no longer afford his salary. He was the last of several employees to be laid off.

Wales promised that Terry Foote would try to sell advertising on Wikipedia to maintain Sanger's position. This led to the February 2002 advertising controversy and the Spanish fork covered in Chapter 2. But for Sanger personally, the financial news was secondary. The community had already made his position untenable.

The money ran out. But Sanger was already looking for the exit.

### Formal Resignation

On March 1, 2002, Larry Sanger posted his resignation to Meta-Wiki. The announcement was careful, diplomatic, and largely misleading about his real reasons for leaving:

"This has nothing to do with my lack of faith in the projects or in the nobility of the mission," he wrote. His relationship with Wales and Bomis "remains friendly." He encouraged Wikipedians not to see his departure as a reason to leave themselves. "My departure should not be taken as a reflection on Wikipedia, or you. It still might succeed brilliantly."

The same day, Jimmy Wales sent an email to Wikipedia-L announcing that "all plans to put advertising of any kind on the wikipedia is called off for now." The Spanish fork had worked. But Sanger didn't return as a volunteer.

## The Truth

Months later, Sanger posted a fuller explanation titled "Why I am no longer participating in the Wikipedia project." The confession there was stark:

"I can't contribute either with regard to policy and (collective) management or to the improvement of articles, without myself feeling... that at any thing I say, I will be forced into a completely ridiculous, unseemly, undignified flamefest."[12]

He identified the structural problem: Wikipedia's philosophical commitment to egalitarianism made his role impossible. He had tried to institute "a policy of respecting and deferring politely to experts." It failed. The community wouldn't accept expertise as a basis for authority. The only authority was consensus, and consensus had turned against him.

Sanger's departure revealed the deepest question facing Wikipedia as a commons: Who has the authority to make decisions? In traditional encyclopedias, the editors decide. In a market, the buyer decides. In a commons, theoretically everyone affected should have a voice. But "everyone" is millions of anonymous people who contributed once and disappeared. Sanger had tried to institutionalize expertise as a decision-making principle. The community rejected this. It wasn't anti-intellectual—it was fiercely democratic. But democracy in a commons is messy.

The departure wasn't just from a job. It was from a community that had decided it didn't want what he offered.

## Community Response

How did the community react to losing its founding organizer?

With muted relief.

Erik Moeller, who would later become the Wikimedia Foundation's Deputy Director, wrote on Kuro5hin: "Larry Sanger played a crucial role in the first year of Wikipedia and has helped establish a shared philosophy for its users... This may well be the right time to abandon the concept

of 'editorship' altogether." The transition from "editorial leadership" to "self-regulation" was framed as progress, not loss.

Others on Meatball Wiki noted: "The project now no longer has a leader (or, put another way, everyone is a leader now)."

The community was ready to see him go. The anti-elitist philosophy that Sanger had struggled against for months had won. Wikipedia would be governed by process, not authority. By consensus, not expertise. By crowds, not editors.

## The Divide

In December 2004, Sanger published an essay on Kuro5hin that made the underlying conflict explicit. Titled "Why Wikipedia Must Jettison Its Anti-Elitism," it articulated the critique he had been developing since his departure:

Sanger expanded on the epigraph's thesis: experts would avoid Wikipedia because their credentials carried no weight. "Experts with much expertise but little patience will avoid editing Wikipedia, because they will—at least if they are editing articles on articles that are subject to any sort of controversy—be forced to defend their edits on article discussion pages against attacks by nonexperts."[13]

He predicted the consequences: trolls would flourish because the community wouldn't enforce standards against them. Experts would stay away because their expertise would carry no weight. Wikipedia would become a source that knowledgeable people couldn't trust.

And he named his expectation for the future: "I do not see how there can not be a more academic fork of the project in the future."

Most strikingly, he acknowledged that this wouldn't change: "I certainly do not expect Jimmy Wales to change his mind. I have known him since 1994 and he is a smart and thoughtful guy; I am sure he has thought through his support of radical openness and his (what I call) anti-elitism."

The founding schism wasn't a misunderstanding to be resolved. It was a fundamental disagreement about how knowledge should be organized. Wikipedia had chosen one model. Sanger represented the road not taken.

## Deeper Consequence

At the time, Sanger's critique focused on quality and expertise. Experts avoiding Wikipedia. Trolls running amok. Credibility problems in specialized areas.

But the choice Wikipedia made in 2002 had consequences Sanger couldn't foresee.

By choosing process over credentials, velocity over verification, openness over expertise, Wikipedia made itself maximally extractable. Content created without credentialed gatekeepers is harder to defend as proprietary expertise. The distributed, anti-hierarchical model makes coordinated legal action nearly impossible. The quality that made Wikipedia valuable—human-curated knowledge, freely accessible—made it the ideal training corpus for AI systems that would have even less accountability than anonymous editors.

Twenty years later, the anti-elitism Sanger diagnosed would reach its logical endpoint. When Elon Musk launched Grokipedia in October 2025, it was the ultimate enclosure of Wikipedia's commons. The original Wikipedia was built by a community committed to shared knowledge and open governance. Grokipedia took that knowledge, fed it through an AI, added sources Sanger himself would have rejected, and offered it as an alternative. The commons that had expelled the expert had been re-enclosed by the algorithmic pseudo-expert.

The choice to expel the advocate of expertise would enable extraction by systems with no expertise whatsoever.

### Inevitable Casualty

From Wales's perspective, and the mainstream Wikipedia history, Sanger's departure was a sad but inevitable consequence of the dot-com crash. Bomis lost money. Salaries were cut. Sanger left. Wikipedia continued and succeeded spectacularly. End of story.

The evidence for this view is substantial. Wikipedia grew from 20,000 articles at Sanger's departure to over 350,000 within two years. The project "stayed the course" on its founding principles and became the world's largest encyclopedia. Sanger's later alternative, Citizendium, peaked briefly in 2008 and then declined into irrelevance, seemingly proving that his expert model couldn't work at scale.

In 2009, Wales called the co-founder debate "silly." In 2023, he said he thought Sanger "doesn't get enough credit for his early work" while maintaining that "co-founder's not the right title." In November 2025, when asked about the dispute on a German podcast, Wales walked out after 48 seconds, calling it "the dumbest question in the world."[14]

The tension remains raw, but the trajectory seems clear: Wikipedia won, Sanger lost.

### Philosophical Expulsion

Sanger's own telling is different. The financial trigger obscured the real cause: philosophical rejection. Wikipedia's egalitarian culture made his role impossible. The community wouldn't accept his vision of structured expertise.

Continued participation was untenable. Even as a volunteer, he couldn't contribute "without myself feeling... that at any thing I say, I will be forced into a completely ridiculous, unseemly, undignified flamefest."

From this perspective, Wikipedia's success came at a cost. The rejection of expertise created a project that experts avoid, that trolls exploit, and that struggles for credibility in specialized domains. Citizendium

failed not because Sanger's model was wrong, but because Wikipedia had insurmountable network effects by the time he tried to build an alternative.

"I named the project, designed its policies, served as its chief editor for its first year, and wrote or commissioned its first articles," Sanger has said. "If that doesn't make me a co-founder, the word has no meaning."

More importantly, Sanger had designed the five pillars that would govern Wikipedia for decades: NPOV, Verifiability, No Original Research. He created the community processes that made Wikipedia work. Yet the institution created to serve those communities had no seat for him. The person who designed governance was excluded from the governance structure. This would not be the last time the Foundation's structure diverged from the community's interests.

### Founding Pattern

A third reading sees the departure as establishing a pattern that would recur throughout Wikipedia's history.

The founding schism didn't resolve—it went dormant. Sanger's concerns about expertise and authority didn't disappear when he left. They resurfaced in every subsequent conflict between Wikipedia's community and its institutions. The governance tensions that made his position untenable would reappear in new forms:

In the 2010s, conflicts between the Wikimedia Foundation and volunteer editors over strategic direction echoed the original Wales-Sanger dynamic. The Foundation grew; the community resisted. The question was always the same: who speaks for Wikipedia?

In 2025, when AI extraction created new fractures in the community, the same philosophical divide reappeared. Pro-Enterprise voices echoed Wales's Hayekian trust in markets and processes. Copyleft purists echoed Sanger's concern that openness enables exploitation.

The founding pattern is clear: governance tensions arise, are not resolved, go dormant, and reappear. The choice to reject expert authority left a vacuum that has never been filled.

## What Followed

After March 1, 2002, Wikipedia continued without formal editorial leadership. Wales became the sole public face of the project. The anti-credentialist philosophy was locked in. Sanger's policies—Neutral Point of View, No Original Research, Verifiability—remained, but his approach to enforcing them did not.

The co-founder dispute would fester for decades. In 2005, Wales edited his Wikipedia biography 18 times, seven of them altering information about Sanger's co-founder status. In 2006, he told the Boston Globe that Sanger's claims were "preposterous." In 2009, Sanger publicly called Wales a "fraud" and "liar" over the dispute. The rawness never faded.

The formal resignation announced an amicable parting. The private confession revealed escape from an impossible position. The underlying tension—who governs open knowledge, and should credentials matter?—was never resolved.

## Pattern Established

Sanger's departure established a template that would repeat throughout Wikipedia's history.

When a conflict emerges between Wikipedia's open philosophy and someone who wants to impose structure, the open philosophy wins. Sanger wanted expert authority; the community rejected him. The Spanish forkers wanted governance guarantees; they got concessions on advertising but not on structural change. Future critics who tried to constrain Wikipedia's openness—whether through content policies, governance reforms, or commercial restrictions—would face the same dynamic.

The choice was made in 2002. Wikipedia would be open, distributed, radically egalitarian. The consequences of that choice would unfold for the next quarter century.

## Long Echo

In September 2025, Larry Sanger published a 37,000-word document: "Nine Theses on Wikipedia"—an echo of Martin Luther nailing his complaints to the church door. He proposed competing articles on controversial topics, abolished source blacklists, required editor identification. The document was the culmination of two decades of critique, refined through failed alternatives and evolving grievances.

Within weeks, Elon Musk shared the theses and announced Grokipedia. The Washington Post profiled Sanger as "inspiring Elon Musk to build a rival."

Twenty-three years after his departure, Sanger's critique had found its most powerful amplifier—and a complicated validation. Grokipedia forked Wikipedia with visible attribution, free access, and CC-BY-SA licensing of derived content. It also added neo-Nazi sources and minimized the Holocaust in its Hitler article. The editorial failures were severe.

But Sanger evaluated Grokipedia as he would any encyclopedia: by its content quality, its potential for improvement, its structural openness. He graded his own entry a "C" and called its confident fabrications "AI bullshittery." He was not blind to the problems. He simply believed that a flawed public encyclopedia was better than no alternative at all—and that the freedom to fork was fundamental to the open knowledge project he had helped create.

The irony was elsewhere: the companies whose AI systems competed with Wikipedia more directly—OpenAI, Anthropic, Google—provided less attribution, charged for access, and enclosed Wikipedia's content in proprietary systems. Grokipedia's content failures were visible and correctable. The invisible extraction was structural and permanent.

The founding schism had not resolved. It had metastasized.

### Bridge Forward

With Sanger gone, Wikipedia needed new governance mechanisms. Chapter 4 explores how the community developed its policy framework—the rules that would manage a project without formal leaders.

But these mechanisms emerged from the same anti-credentialist philosophy that expelled Sanger. They could constrain individual editors. They could not constrain the institutions that would later emerge to steward the project. When the Wikimedia Foundation was created, when Enterprise sales began, when AI companies started extracting volunteer labor, no governance structure existed to require community consent.

The governance vacuum Sanger's departure created would enable decisions that echoed everything he had warned about. His critique of volunteers building value for Bomis would find its endpoint in volunteers training commercial AI without compensation. The expert-free model he opposed would become the training data for systems with even less expertise than Wikipedia's anonymous crowds.

The question that haunted his departure would haunt the project itself: What happens when money enters the equation?

Sanger knew the answer. He just couldn't convince anyone to listen.

CHAPTER FOUR

# The Policies

> The threshold for inclusion in Wikipedia is verifiability, not truth.
>
> — Wikipedia policy, 2003-2012

In early 2003, a Wikipedia editor with a background in physics faced a familiar problem.

Someone had inserted a fringe theory into an article about cold fusion. The editor had reverted it. The inserter had restored it. Back and forth, the same cycle playing out for the third time this week. The theory claimed to explain a new form of energy generation. The only source cited was the inserter's personal website.

"That's not a reliable source," the physics editor wrote on the talk page.

"My theory is true," the inserter replied. "I've explained it completely. Just because it hasn't been published in a journal doesn't mean it's wrong."

And here was the problem. The physics editor knew the theory was wrong—or at least, that it had no place in an encyclopedia. But Wikipedia had no formal policy to cite. The editor couldn't invoke credentials ("I have a PhD in physics") because Wikipedia had rejected credentials as a basis for authority. The editor couldn't simply ban the inserter because Wikipedia's openness meant anyone could edit. The only argument available was "this

is obviously fringe nonsense," and the inserter didn't find that obvious at all.

The physics editor needed something beyond common sense. A standard. An external reference point that could settle disputes without requiring agreement on what was true.

This was the physics crank problem. And its solution would define how Wikipedia organized knowledge for the next two decades.

### Sanger Void

When Larry Sanger departed in March 2002, Wikipedia lost more than an employee. It lost the only person who had claimed authority to make editorial decisions.

In February 2002, just weeks before his resignation, Jimmy Wales had written: "Larry is the final arbiter of what the consensus is." After Sanger left, Wales inherited this role by default rather than design. He made decisions when asked, intervened in disputes when necessary, and served as the project's public face. But his authority was never formalized. Wikipedia had no constitution, no bylaws, no governance structure at all.

This worked, for a while. Wikipedia was still small enough that disputes could be resolved through conversation. Wales was present enough to adjudicate when conversations failed. The community shared enough values that most questions had obvious answers.

But the project was growing. Twenty thousand articles at the end of 2001 became fifty thousand by mid-2002. More contributors meant more disputes. More disputes meant more need for clear rules. And some of those disputes—like the physics crank problem—couldn't be resolved by appeal to founder authority. They needed principles.

### Scale Problem

The challenge wasn't just quantity. It was the kind of disputes that scale produced.

When Wikipedia had a few hundred active editors, they knew each other. Disagreements could be resolved through reputation and relationship. A respected editor's opinion carried weight because other editors knew their track record. This was informal expertise—not credentials, but earned trust.

As the community grew, this social knowledge became impossible to maintain. New editors arrived who knew nothing about existing contributors. They inserted content, disputed deletions, and demanded explanations. "Who are you to say this is wrong?" was a question that small communities could answer through familiarity. Large communities needed something else.

The "truth war" pattern emerged. Editors would argue: "I know this is true." Other editors would respond: "No, it isn't." Without a shared standard for evidence, such disputes could continue indefinitely. The physics cranks were the most dramatic example—people genuinely believed in perpetual motion or cold fusion and would fight endlessly for inclusion—but the pattern appeared across domains. Political disputes, historical interpretations, medical claims. Everywhere, the same dynamic: my truth against yours.

Wikipedia needed a way to settle such disputes without resolving the underlying disagreement about truth.

### Cultural Context

The solution that emerged reflected Wales's intellectual commitments, traced through Chapter 1: distributed knowledge, emergent order, process over credentials. If experts couldn't be trusted to decide what was true, then perhaps the question itself could be reframed. Instead of asking "what is true?" Wikipedia could ask "what do reliable sources say?"

This was more than a procedural trick. It was an epistemological shift. Wikipedia would describe the state of published knowledge rather than assert truths directly. It would defer to sources rather than claim authority

itself. The project would become a curator of existing knowledge rather than a producer of new knowledge.

From this reframing, three policies would emerge.

## NPOV

The Neutral Point of View policy was the oldest and most fundamental. Larry Sanger had drafted an early version for Nupedia in 2000, calling it the "Non-bias policy." He adapted it for Wikipedia by February 2001. Wales elaborated the concept with what he later called his "original formulation" in April 2001. Sanger wrote an expanded version at Meta-Wiki in December 2001.

In November 2003, Wales declared NPOV "non-negotiable"—a statement that simultaneously codified its importance and constrained future debate. "A few things are absolute and non-negotiable, though," he wrote. "NPOV for example."

What did NPOV decide? It was the first true commons governance policy. Instead of asking "What is true?" (which divides any community), NPOV asked "What do reliable sources say?" This reframe was genius precisely because it let people with incompatible worldviews work together on the same articles. An atheist and a believer could both edit an article on God by agreeing to describe what philosophers and theologians have written. NPOV wasn't neutral about everything—it privileged published sources. But this particular partiality allowed a commons to function.

At its core, it required that articles represent all significant viewpoints fairly, without asserting that any particular viewpoint was correct. Wikipedia would describe controversies rather than resolve them. If reliable sources disagreed, the article would present the disagreement. If one view was dominant in published sources, the article would reflect that dominance while still noting alternatives.

The policy's brilliance lay in what it deferred. NPOV didn't require editors to agree on truth. It only required them to agree on what sources

said. Two editors who disagreed about evolution could collaborate on an article if they could agree that published biology textbooks overwhelmingly supported evolutionary theory while creationist sources existed. The article would describe both—with appropriate weight—and neither editor had to abandon their personal beliefs.

Joseph Reagle, in his study of Wikipedia's culture, called this "productive ambiguity":

"While the perception is that NPOV is the source of much debate, it may act rather as a heat shield: reducing conflict and otherwise channeling outstanding arguments in the productive context of the primary goal of developing an encyclopedia."[15]

But NPOV couldn't solve every problem. What made a viewpoint "significant"? How should conflicting sources be weighed? And critically: NPOV addressed how to present knowledge but not what counted as knowledge in the first place. For that, Wikipedia needed additional policies.

### Verifiability

Verifiability emerged in August 2003, directly from failures in NPOV enforcement. The physics cranks had exposed a gap: NPOV said to represent significant viewpoints, but cranks claimed their viewpoints were significant. What external standard could distinguish legitimate minority positions from personal theories?

The answer was publication. If a viewpoint had been published by a reliable source, it was verifiable. If it existed only on someone's personal website, it was not. The burden of proof fell on content-adders: if you wanted something in Wikipedia, you needed to cite a source.

The policy's most controversial phrase captured this epistemological stance—the same one that would later appear in Wikipedia's official documentation: not truth, but verifiability.

This was deliberately provocative. Wikipedia explicitly disclaimed responsibility for determining truth. A claim could be included if a reliable source published it, even if the claim was actually false. A claim could be excluded if no reliable source published it, even if the claim was actually true. Wikipedia would reflect published knowledge, not reality directly.

The physics crank problem was solved. "Your personal website is not a reliable source" became a decisive argument. The crank's theory, however compelling to its author, had not been published in journals. It could not be included. The dispute was over—not because anyone had proven the theory wrong, but because it failed the verifiability test.

But new questions emerged. What counted as "reliable"? A peer-reviewed journal was clearly reliable; a blog post was clearly not. But what about newspapers with uneven fact-checking? Think tanks with ideological agendas? Publishers in developing countries with different standards? The policy provided a principle but left interpretation to case-by-case judgment.

### No Original Research

The third policy, No Original Research, was codified in December 2003. Wales later acknowledged its origins: "The original research policy originating primarily as a practical means to deal with physics cranks, of which of course there are a number on the web."

NOR completed the epistemological framework. Wikipedia would synthesize existing published knowledge, not create new knowledge. Editors could summarize what sources said. They could not draw novel conclusions, present unpublished arguments, or introduce information that no reliable source had published.

This meant Wikipedia had a knowledge boundary. Topics without published reliable sources could not have Wikipedia articles. Oral traditions that had never been written down were excluded. Emerging fields that hadn't yet produced peer-reviewed literature were excluded. Commu-

nities without publishing infrastructure—disproportionately in the Global South—found their knowledge systematically underrepresented.

The policy's architects knew this was a trade-off. Excluding unpublished knowledge meant excluding valuable knowledge. But the alternative—allowing anyone to add anything they believed true—would destroy Wikipedia's utility as a reference. The fringe theory problem had proven that.

Together, the three policies formed Wikipedia's epistemological architecture. NPOV defined what Wikipedia would say: describe viewpoints, don't assert truth. Verifiability defined how claims would be validated: cite reliable sources. NOR defined what Wikipedia would not do: create new knowledge. The framework gave editors tools to resolve disputes that had previously festered indefinitely.

## What Remained Undecided

This is where the standard story of Wikipedia's governance needs revision.

The three core content policies were a remarkable achievement. They enabled millions of editors to collaborate despite fundamental disagreements about truth. They created common ground where none had existed. They solved the fringe theory problem.

But they solved only one kind of problem. The policies addressed content. They did not address power.

And power matters in a commons. Commoners have rights—the right to contribute, to see decisions made transparently, to have a say in how shared resources are governed. The content policies told contributors how to write articles. They said nothing about who controlled Wikipedia's direction. Did volunteers have the right to veto commercial partnerships? Could they require transparency about Enterprise deals? The policies created a brilliant system for deciding *what Wikipedia would say*. They created no system for deciding *who controlled Wikipedia*.

Consider what remained undefined after December 2003. Who has ultimate authority over Wikipedia's direction? Wales, informally—but nothing constrained his authority. What happens when the community and the institution disagree? No process existed; the assumption was they wouldn't disagree, that shared mission would prevent conflict. Who controls commercialization, partnerships, external relationships? Not addressed—these questions seemed peripheral to building an encyclopedia. What is "consensus" and how is it determined when it breaks down? It remained a principle without a clear mechanism for contested cases.

In December 2003, Wales created the Arbitration Committee to handle intractable disputes. He described ArbCom as "an extension of the decision-making power he formerly held"—a revealing phrase. His authority didn't dissipate; it institutionalized. ArbCom could ban editors and resolve content disputes. It could not constrain the Foundation that would soon emerge, or determine Wikipedia's commercial relationships, or protect volunteer interests against institutional decisions.

This structural gap would matter enormously. When the Wikimedia Foundation grew powerful, when AI companies came calling, when volunteers wondered who really owned their creation—the governance questions these early policies sidestepped would become central. Content governance was solved. Institutional governance was left to trust, goodwill, and the assumption that everyone shared the same mission.

The policies' useful vagueness enabled collaboration. The same vagueness would later enable co-optation.

### Pragmatic Success

From Wales's perspective, and the mainstream view within the Wikipedia community, the policies were a triumph. They solved a real problem: how to scale collaborative knowledge production without gatekeeping experts. They enabled editors with incompatible worldviews to work together productively. They made Wikipedia possible.

The evidence was compelling. Wikipedia grew from tens of thousands of articles to millions. A 2005 study in Nature found its accuracy comparable to Encyclopedia Britannica. Edit wars declined after the policies formalized. The same framework was adopted across hundreds of language Wikipedias, proving its cross-cultural applicability.

"NPOV ensures that we can join the scattered pieces of what we think we know and good faith facilitates the actual practice of fitting them together," Reagle wrote. The policies didn't resolve disagreements about truth—they made such resolution unnecessary for collaborative work.

## Anti-Elitism Critique

From Sanger's perspective, the policies encoded everything wrong with Wikipedia's direction. His 2004 essay "Why Wikipedia Must Jettison Its Anti-Elitism" articulated the critique:

"The root problem: anti-elitism, or hatred of expertise. By 'anti-elitism' I do not mean the view that expertise should not be the only or even the main determiner of what content a wiki has. I mean, instead, the disdain for expertise."[13]

Verifiability, in this view, privileged publication over knowledge. A tenured physicist couldn't override a published crank theory on the basis of expertise. The policy's epistemological stance struck critics as cowardice—Wikipedia claimed no responsibility for accuracy while becoming the world's de facto reference. When errors persisted in specialized articles, there was no mechanism for experts to correct them authoritatively.

The policies created an illusion of neutrality while embedding systematic biases. Published sources reflected existing power structures. Communities without publishing infrastructure were excluded. Wikipedia described the knowledge that had been published, and publication itself was not neutral.

### Infrastructure Critique

Critical scholars like Heather Ford offered a third perspective. The policies appeared neutral but encoded power relations invisible to their architects:

"Wikipedia's infrastructure introduces new and less visible sources of gender disparity," Ford wrote with Judy Wajcman.[16] The same could be said for geographic disparity, linguistic disparity, and disparities of all kinds.

NOR excluded oral traditions. Verifiability excluded unpublished knowledge. NPOV's requirement for "significant" viewpoints meant that marginalized perspectives—significant to their communities but not to mainstream publishers—could be dismissed as fringe. The policies solved the internal governance problem while ignoring external inequities.

From this view, the fringe theorists were a distraction. The real exclusions weren't crackpots with personal websites. They were entire knowledge traditions that had never been published in forms Wikipedia could recognize.

### Synthesis

The policies succeeded at their stated purpose: enabling mass collaboration on encyclopedic content. They failed at purposes they never claimed: ensuring equity, safeguarding against commercial co-optation, or governing the institution itself.

This was not a flaw in design. The policies emerged to solve specific problems—fringe theories, edit wars, the scale challenge. They solved those problems brilliantly. The problems they created or ignored were problems their architects could not foresee.

The debate continues because the policies' constructive vagueness cuts both ways. It enables collaboration and masks power. It resolves content disputes and defers governance questions. Fringe theorists met an insurmountable barrier. Commercial extractors would find no barrier at all.

### Framework Complete

By December 2003, Wikipedia's epistemological architecture was in place. The three policies interlocked: NPOV set the standard for representation, Verifiability established the evidence threshold, NOR drew the boundary against novel claims. ArbCom provided dispute resolution for intractable conflicts.

The fringe theorist who had argued for cold fusion could now be shown the door—not through debate about physics, but through citation of policy. Personal websites did not meet the reliability standard. Novel theories violated NOR. The framework worked.

But notice what the framework addressed. It governed content. It said nothing about power.

### Five Pillars

On May 4, 2005, a Wikipedia editor with the username Neutrality created a new page: "Wikipedia:Five pillars." The page synthesized Wikipedia's core principles: Wikipedia is an encyclopedia; it is written from a neutral point of view; it is free content; it has no firm rules (use common sense); and editors should treat each other with respect.

There was no community discussion. No vote. No formal process. One editor created a page that summarized existing norms, and the community accepted it. Within months, the Five Pillars became Wikipedia's de facto constitution, cited in countless disputes and orientation materials.

This was the triumph and the trap of Wikipedia as a commons without formal governance. The Five Pillars emerged through the most democratic possible process: emergence without authority. One editor summarized norms that already existed. The community recognized them as accurate. Norms could be created by consensus.

But they could also be abandoned by consensus, and consensus itself was never institutionalized. A real constitution has enforcement mechanisms. The Five Pillars had only agreement. When commercial interests

began to threaten commons principles, there would be no legal recourse, no institutional protection.

This pattern—emergent governance through individual initiative, accepted through presumed consensus—characterized Wikipedia's institutional development. It worked when participants shared goals. It would fail when they did not.

## Gap Planted

The policies that made Wikipedia's content defensible could not make Wikipedia's value defensible.

When AI companies began extracting Wikipedia at industrial scale, they encountered no governance structure designed to address such extraction. The volunteers had built rules for content. They had not built rules for who could profit from that content. The Foundation that would emerge to steward Wikipedia was not constrained by these policies. It could make commercial decisions—including decisions about AI company access—without community consent, because no consent mechanism existed.

Fringe theorists hit a wall. The companies that would later harvest Wikipedia's content at industrial scale encountered no such obstacle—only a governance void the policies never addressed.

## Bridge Forward

Part I has traced Wikipedia's founding fractures: the expert model rejected, the commercialization revolt, the co-founder expelled, the policies codified. The project emerged from these conflicts with a clear epistemology but an undefined institutional structure.

Part II shifts from content governance to licensing governance. The GNU Free Documentation License that Wikipedia adopted in 2001, and the Creative Commons license it migrated to in 2009, would determine not just how Wikipedia worked but who could use its value. The copyleft

promise—that derivative works must be shared freely—was supposed to protect against exploitation.

But copyleft, like the content policies, addressed one kind of problem while leaving others unsolved. The license protected against traditional publishing exploitation. It could not anticipate AI training, Enterprise APIs, or the financial entanglements that would compromise enforcement.

The governance gap this chapter identified would grow as the stakes increased. Twenty years after these policies stopped the fringe theorists, contributors would discover they had built a cathedral of knowledge on land they did not own. The framework that empowered them to create Wikipedia created no mechanism to control what happened to their creation.

The foundation was laid. The fractures were planted. What remained was to see what would grow from both.

# Part II

# The Licensing Gamble

CHAPTER FIVE

# The Copyleft Promise

> GNU is not in the public domain. Everyone will be permitted to modify and redistribute GNU, but no distributor will be allowed to restrict its further redistribution.
>
> — Richard Stallman, GNU Manifesto, 1985

In the winter of 1982, Richard Stallman returned to the MIT Artificial Intelligence Laboratory and found the PDP-10 sitting dark and silent.

The machine had been the heart of a community. For more than a decade, some of the world's best programmers had gathered around it, sharing code freely, building on each other's work, treating software as a communal resource. They called themselves hackers—a term that meant something different then. It meant someone who loved to program, who sought elegant solutions, who believed that information wanted to be free.

The indicator lights were out. The fans were silent. No one was left to fix it.

Stallman stood alone in the equipment room and began to cry.

"Seeing the machine there, dead, with nobody left to fix it," he later recalled, "it all drove home how completely my community had been destroyed."[17]

The destruction hadn't been sudden. It had come gradually, as commercial interests colonized what had once been shared space. Two companies—

Symbolics and Lisp Machines, Inc.—had hired away most of the lab's programmers to commercialize the Lisp Machine. The code that had been freely shared was now proprietary. The improvements that had benefited everyone now belonged to corporations.

Stallman had fought back. When Symbolics ended their agreement to share improvements, he spent two years cloning every feature they developed, working alone to ensure their monopoly failed. "I was going to punish Symbolics if it was the last thing I did," he said.

But this was a rearguard action. The war was already lost. The hacker community that had formed around the PDP-10 was gone.

What do you do when your community is destroyed? Stallman's answer would shape the legal framework for free knowledge for the next forty years. It would become the foundation of Wikipedia's promise to its contributors. And it would plant the assumptions that would later prove fatal.

## Lab Culture

Stallman joined the MIT Artificial Intelligence Laboratory in 1971, at age eighteen. The lab embodied a distinct philosophy about software: code was shared freely, improvements were collective, and no one owned what everyone built.

This was not idealism. It was practical. When one programmer improved a piece of code, everyone benefited. When someone found a bug, the fix went into the common codebase. The machine and the humans around it formed a collaborative whole, each contributing what they could, each using what they needed.

"For more than ten years, many of the world's best programmers worked at the Artificial Intelligence Lab for far less money than they could have had anywhere else," Stallman later wrote in the GNU Manifesto. "They got many kinds of nonmonetary rewards: fame and appreciation, for example. And creativity is also fun, a reward in itself."[18]

The lab had no passwords initially. When MIT installed security controls in 1977, Stallman cracked them and offered the results to anyone who wanted to opt out. Security was antithetical to sharing.

What governed this shared resource was a set of norms and expectations—what we might call a commons. No formal licensing existed yet, but the principle was clear: those who benefited from collective work had an obligation to return their improvements.

This culture had limits. When two companies formed to commercialize the Lisp Machine—a specialized computer the lab had developed—the communal model collided with commercial interests.

### Symbolics War

Symbolics and Lisp Machines, Inc. (LMI) both emerged from the AI Lab in 1980, each seeking to profit from technology the community had built together. This was the first enclosure: the appropriation of commons resources by commercial entities that refused to return their improvements.

Initially, a "gentleman's agreement" allowed Stallman to review Symbolics' improvements. He could incorporate them into the common codebase, ensuring the community still benefited.

On March 16, 1982—Stallman's twenty-ninth birthday—Symbolics ended that agreement. They would no longer share their improvements. If the lab wanted the latest features, they would have to abandon the community approach entirely.

Stallman interpreted this as an act of war. For two years, he cloned every feature Symbolics developed, working alone against an entire company. His framing was martial: "The AI Lab was a neutral country, like Belgium in World War I. If Germany invades Belgium, Belgium declares war on Germany."

Yet this was a holding action, not a victory. The lab's programmers were leaving. The community was fragmenting. No single person, however dedicated, could replace what collective effort had built.

Some of Stallman's colleagues saw things differently. They viewed Symbolics not as an invader but as a necessary evolution—bringing hacker principles into the commercial marketplace. To them, Stallman was "a troubling anachronism," clinging to a model that couldn't survive contact with economic reality.

The social isolation was real. "I was no longer getting invited to go to Chinatown," Stallman recalled. "The custom started by Greenblatt was that if you went out to dinner, you went around or sent a message asking anybody at the lab if they also wanted to go. Sometime around 1980-1981, I stopped getting asked."

This was not just philosophical dispute. It was personal.

## GNU Manifesto

On September 27, 1983, Stallman announced his response to a Usenet newsgroup: he would build an entire operating system that could never be taken from its users.

He called it GNU—a recursive acronym for "GNU's Not Unix." The project would recreate the functionality of Unix, the dominant operating system, but with one crucial difference: GNU would be free forever. Not free as in price, but free as in freedom. Users would have the right to run it, study it, modify it, and share it.

On January 5, 1984, Stallman quit his job at MIT to work full-time on GNU. He wanted to avoid any claim that MIT owned his code. In March 1985, the GNU Manifesto appeared in Dr. Dobb's Journal, laying out both the practical plan and the philosophical foundation.

The manifesto articulated what would become copyleft. Proprietary modifications would not be allowed. Every version of GNU would remain free. Stallman had learned from Symbolics that informal sharing agreements could be revoked. Only legal architecture could ensure permanence.

This was different from simply giving software away. Public domain software could be taken, modified, and locked up—exactly what had hap-

pened at the AI Lab. Stallman's innovation was to use copyright law against itself: the software would be copyrighted, but the license would require that anyone who modified and distributed it must preserve the same freedoms.

The license was a weapon. It was designed to "punish" those who would take shared work and refuse to share their improvements.

### Four Freedoms

Stallman codified what "free software" meant through four freedoms. Freedom 0 was the freedom to run the program as you wish, for any purpose. Freedom 1 was the freedom to study how the program works and change it, which requires access to source code. Freedom 2 was the freedom to redistribute copies so you can help others. Freedom 3 was the freedom to distribute copies of your modified versions, giving the whole community a chance to benefit. The numbering was significant: Freedom 0 was added around 1990 when Stallman realized the freedom to simply use the program needed explicit protection.

These freedoms were not merely desirable. They were essential. Without them, users were subject to the software's owners. With them, users controlled their own computing.

### License Evolution

The legal framework evolved through hard lessons:

In 1985, Stallman published the first copyleft license, specific to GNU Emacs. This became necessary after the Gosling Emacs incident: James Gosling had written his own version of Emacs in 1982, initially allowing free distribution. Stallman incorporated some of this code into GNU Emacs. When Gosling sold the rights to UniPress, they forced Stallman to remove the code. Informal sharing agreements were insufficient. Only legal protection could ensure freedom.

GPL version 1 appeared in January 1989—the first program-independent copyleft license. Any project could adopt it, making the license itself a reusable component.

GPL version 2 followed in June 1991, adding the "Liberty or Death" clause: if you cannot comply with all GPL terms due to patent claims or other legal restrictions, you cannot distribute the software at all. This prevented GPL circumvention through external legal constraints.

Each step responded to attempts at enclosure. Each built higher walls around the commons.

## Beyond Software

By the late 1990s, Stallman recognized that software freedom was incomplete without documentation freedom:

"The biggest deficiency in free operating systems is not in the software—it is the lack of good free manuals that we can include in these systems. Many of our most important programs do not come with full manuals. Documentation is an essential part of any software package; when an important free software package does not come with a free manual, that is a major gap."[19]

The problem was that good manuals were being published with restrictive terms. O'Reilly Associates, for example, produced excellent documentation but prohibited copying and modification. These manuals were "nonfree" and thus "excluded from the Free World."

The GNU Free Documentation License (GFDL), released in March 2000, extended copyleft to documentation. It required attribution, preserved modification history, and ensured derivatives remained free. It also included provisions for "invariant sections"—portions that could not be altered, designed to protect philosophical statements from commercial publishers who might want to strip the ideology from the documentation.

### License Assumptions

This is where the standard story of copyleft needs revision.

The GFDL was designed to prevent a specific form of enclosure: commercial entities taking free documentation, improving it, and refusing to share their improvements. It was a weapon against Symbolics-style betrayal, scaled from software to text.

The license assumed that derivatives would be made by humans, distributed as human-readable works, and used by other humans.

Consider what the GFDL required: full license text with every copy, complete history section preserved, attribution of principal authors, transparent copy availability (source must be accessible), and ShareAlike for all derivatives.

Each requirement assumed something about how content would be used. "Derivative works" would be recognizable modifications—a revised manual, an updated textbook, a translated version. "Distribution" meant sharing copies with other people—handing them a book, posting a file, burning a CD. Attribution could be verified by examining the work—the credits page, the copyright notice, the documentation.

The license could not anticipate—no one could have anticipated—a form of use that would transform content into statistical weights, create outputs that echo the training data without being "copies," and operate at industrial scale without human review.

When AI companies train on Wikipedia, they create something that may not legally qualify as a "derivative work." The model weights are not a modified encyclopedia. The outputs are not edited articles. The copyleft mechanism—propagation through derivatives—may have no purchase on this new form of extraction.

But this is not a failure of Stallman's vision. It is a failure of enforcement.

The legal arguments for enforcement exist. AI outputs that reproduce training data verbatim are copies. Memorization at scale may constitute

a derivative work. Market harm—Wikipedia's 8% traffic decline—weighs against fair use claims. Commercial purpose tips the fair use balance.

Making these arguments requires an institution willing to litigate. Copyleft was a weapon. But weapons require someone willing to use them.

## Copyleft Vision

From Stallman's perspective, and the Free Software Foundation's, copyleft was essential protection.

The alternative—permissive licenses like BSD and MIT that allow proprietary derivatives—meant that free software became a gift to commercial developers. They could take, improve, and lock up. The community's work would benefit corporations without any reciprocity.

"The fundamental act of friendship among programmers is the sharing of programs," Stallman wrote. "Marketing arrangements now typically used essentially forbid programmers to treat others as friends."[18]

Copyleft ensured that freedom would propagate. Anyone could use, modify, and distribute the work—but only if they preserved the same freedoms for others. The mechanism was described as "viral": freedom spread through every derivative.

The evidence supported this view. Linux, the kernel that completed the GNU operating system, was released under GPL. The combination became the foundation of the modern internet. Servers running free software powered the web. The copyleft mechanism had worked: commercial entities contributed improvements back to the commons because the license required it.

## Pragmatic Critique

Not everyone embraced copyleft's maximalism.

Critics within the open source community argued that copyleft's "viral" nature created unnecessary complexity. Businesses feared inadvertently

"contaminating" proprietary code. Simpler licenses—MIT, BSD, Apache—achieved broader adoption without the compliance burden.

The market appeared to validate this view. GPL usage declined from roughly 45% of open source projects to around 22%. Permissive licenses gained ground. Cloud computing enabled companies to use copyleft software without triggering distribution requirements, since running software as a service wasn't legally "distribution."

From this perspective, Stallman's approach was ideologically pure but practically costly. The goal should be maximum adoption of open source, not maximum protection of ideological freedoms.

### Documentation Mismatch

When it came to documentation specifically, a third critique emerged: the GFDL was designed for software manuals, not encyclopedias.

The license's requirements—invariant sections, cover texts, five-page license text with every copy—made sense for printed manuals distributed by publishers. They made less sense for wikis where content was fluid, authorship was distributed, and reuse might mean excerpting a single paragraph.

"For such small documents, it just doesn't make sense to require people to pass out another 10 pages of legalese text," observed the founders of Wikitravel when they chose CC-BY-SA over GFDL in 2003.

This critique would eventually motivate Wikipedia's 2009 migration away from GFDL. But in 2001, when Wikipedia launched, the GFDL was the only credible copyleft license for documentation. Creative Commons didn't exist yet.

### Synthesis

The debate revealed a tension that was never resolved: between maximum protection and maximum adoption, between ideological purity and pragmatic effectiveness.

Stallman's approach was maximalist. Better to have fewer users than to enable enclosure. Better to maintain strong copyleft than to watch commercial entities take without giving back.

The pragmatists prioritized adoption. Better to have more people using free software, even if some of them build proprietary derivatives.

Both camps assumed enforcement would happen. Whether strong copyleft or weak, licenses only work if someone is willing to enforce them. The later question—whether the institution that inherited Stallman's legal framework would actually use it—was not imagined in these earlier debates.

## Wikipedia Enters

In January 2001, after email exchanges with Richard Stallman himself, Jimmy Wales committed Nupedia—and then Wikipedia—to the GFDL.

This was not a casual choice. The GNU Project noted: "Just as we were starting a project, GNUpedia, to develop a free encyclopedia, the Nupedia encyclopedia project adopted the GNU Free Documentation License and thus became a free commercial project. So we decided to merge GNUpedia project into Nupedia."

Wikipedia entered Stallman's ideological framework. The project accepted his legal architecture, his philosophical commitments, and his assumptions about how protection would work.

Contributors who edited Wikipedia understood the deal: their work would remain free forever. Anyone could use it, modify it, share it—but derivatives would have to preserve the same freedoms. Commercial publishers couldn't lock up volunteer labor. The copyleft mechanism would ensure that what was shared stayed shared.

This was the promise.

## Flaw Planted

But the promise rested on assumptions that no one questioned.

Copyleft assumed that derivatives would be recognizable. It assumed distribution would be identifiable. It assumed enforcement would be vigilant. It could not anticipate a guardian who became financially dependent on those it was meant to police.

The first two assumptions would be tested by AI training, which transforms content into statistical weights that may not legally qualify as derivatives. The third would be tested when the organization stewarding the license became financially entangled with the companies whose compliance was questionable. The legal framework could not anticipate a guardian whose cultural networks overlapped with the extractors—who viewed them not as adversaries but as colleagues.

Stallman had stood before the dead PDP-10, watching a community's destruction. His response was to build legal defenses against future enclosures. Those defenses worked for decades, protecting free software and free documentation from commercial appropriation.

But legal frameworks only work if someone is willing to enforce them.

Chapter 6 will show how GFDL actually worked in practice—its awkward fit for Wikipedia's use case, the friction that made reuse difficult, the seeds of the 2009 migration to CC-BY-SA. And Chapter 19 will close the forty-year circle: when AI companies train on Wikipedia without attribution or share-alike compliance, they do exactly what Symbolics did to Stallman—take shared work and refuse to give back.

The difference: Stallman fought. The institution that inherited his framework chose revenue.

## Promise Made

In 2001, Wikipedia made a promise to its contributors: your work will remain free forever. That promise was encoded in the GFDL, backed by Stallman's legal architecture, and inherited from the grief of a programmer who watched his community die.

Tens of thousands of editors would contribute under that promise. Millions of articles would be created. The largest encyclopedia in human history would emerge from collective volunteer labor, protected by copyleft.

Twenty-three years later, that work would train systems worth hundreds of billions of dollars. The editors would receive no attribution. The derivatives would share nothing alike. The promise would become theoretical—not because the license failed, but because no one was willing to enforce it.

The tragic flaw was planted in the protection itself: copyleft assumed an enforcer who shared Stallman's commitments. It could not protect against a guardian who stopped guarding.

The threshold was crossed. What came next would determine whether the promise was kept.

CHAPTER SIX

# The GFDL Years

> If you have to staple the entire text of the GFDL to a short article that you hope to print on a flier, you effectively can't reuse that article on a flier.
>
> — Jonathan Corbet, LWN.net

In 2006, a teacher wanted to distribute a one-page handout about photosynthesis to her students. Wikipedia had an excellent article on the topic—clear, accurate, well-illustrated. She could legally copy it. The license said so.

She opened a word processor, pasted the text, and began formatting. Then she read the license requirements.

The GNU Free Documentation License required that any copy include the full license text. Five pages of legal language. For her one-page handout, she would need to attach five pages of legalese.

The license also required listing the "principal authors"—but the Wikipedia article had been edited by hundreds of people. Whose names should appear? The license required preserving the complete revision history. For a classroom handout?

She closed the word processor. She typed the handout from scratch, summarizing the article in her own words. The license that promised "free knowledge" had made the knowledge unusable.

This was the flier problem. And it was just the beginning.

### January 2001

When Wikipedia launched, the GNU Free Documentation License was the only credible copyleft license for non-software content. Creative Commons would not release its first licenses until December 2002—eighteen months later. For a project committed to copyleft, GFDL was the only option.

Jimmy Wales had corresponded with Richard Stallman in early January 2001 about licensing. Wales was not a free software ideologue; he was an entrepreneur who had made money through Bomis, a web portal that included adult content. But he understood that a free encyclopedia needed a free license, and Stallman was the expert.

As one Slashdot commenter noted: "Having Richard Stallman brow beat Jimmy Wales certainly didn't hurt either."

The GFDL adoption was pragmatic necessity, not ideological conversion. This commercial DNA would resurface later—in Wikia (founded 2004), in Wikimedia Enterprise (launched 2021)—but in 2001, it simply meant Wales chose the available tool.

### GFDL Requirements

The license imposed specific obligations that made sense for software manuals. Every copy must include the complete GFDL—approximately five printed pages. The complete modification history must be preserved, including dates and authors. At least five authors must be listed on the title page (or all, if fewer than five). The source format must be made available, not just the rendered output. If the original specified cover texts, these must be preserved. Certain sections designated as "invariant" could not be modified or removed.

Wikipedia prohibited invariant sections and cover texts, avoiding some complexity. But the other requirements applied fully.

For a printed software manual, these requirements were manageable. The manual already had a title page. The revision history already existed. Five pages of license text were negligible in a 500-page technical reference.

For a wiki, they were awkward. Who were the ”principal authors” of an article edited by 847 people? How could a revision history be preserved when every edit created a new version? What did ”transparent copy” mean when the ”source” was a wiki page that changed constantly?

### Commercial DNA

Wales's background matters for understanding what came later. He was not opposed to commercial applications of free content. Bomis itself had demonstrated that advertising revenue could fund projects. The GFDL's restrictions on commercial use were not religious commitments; they were practical constraints he accepted because the alternative—proprietary licensing—would have made Wikipedia's collaborative model impossible.

This pragmatic approach to licensing would manifest again in 2009, when convenience motivated migration to CC-BY-SA. It would manifest again in 2021, when Wikimedia Enterprise began selling API access to AI companies. The tension between free content ideology and commercial instincts was present from the beginning.

### Practical Problems

By the mid-2000s, the GFDL's awkwardness for Wikipedia had become a running joke in the community.

**The flier problem** was just the start. Any physical reproduction required five pages of license text. Teachers couldn't make classroom handouts. Community organizations couldn't print educational materials. The ”free” encyclopedia required more legal overhead than most proprietary content.

**The audio and video problem** was worse. Creating a video based on Wikipedia content meant... what? Reading five pages of GFDL aloud?

Displaying them as on-screen text? The license was written for documents. It had no practical application to multimedia.

**The image problem** emerged because GFDL was "mainly meant for documentation and text documents but not images." A photograph released under GFDL technically required the full license text to accompany every use. A 10-kilobyte image would need a 50-kilobyte text file.

**The many-authors problem** became absurd at scale. By 2006, the article on "United States" had been edited by over 8,000 accounts. Listing the "five principal authors" required determining who had contributed most—a computation nobody performed, creating universal non-compliance.

**The legal uncertainty problem** was perhaps most troubling. As one LWN.net analysis noted: "No reuser can be confident that they're even doing it legally, even if they're willing to take heroic measures. The FSF will not say what the GFDL means as applied to Wikipedia."

The Free Software Foundation had created the license. Wikipedia was its largest user. And the FSF would not clarify how the license applied.

### Defections

Other projects looked at Wikipedia's struggles and chose differently.

In 2003, **Wikitravel** launched under CC-BY-SA, explicitly rejecting GFDL. The founders explained through Joi Ito: "Looking over the licenses themselves... they thought that the licenses made a lot more sense than the software-manual-oriented GFDL. For such small documents, it just doesn't make sense to require people to pass out another 10 pages of legalese text."

In March 2006, **Debian**—the respected free software distribution—held a vote on whether GFDL documents could be included in Debian's main repository. The result: documents with invariant sections were "not suitable for Debian main." A major legitimacy blow from the free software community itself.

In 2007, **FLOSS Manuals** abandoned GFDL for GPL, citing "difficulties in implementing the GFDL."

The criticism escalated. LWN.net summarized the emerging consensus: "The GFDL has come to be seen by many as more of a tool for the propagation of FSF propaganda than a license for truly free documentation. Much of the community avoids this license; some groups, such as the Debian Project, see it as non-free."

### Incompatibility Crisis

The defections created a new problem: licensing silos.

Content under GFDL could not be legally combined with content under CC-BY-SA. Both were copyleft licenses. Both required share-alike. But their terms differed in ways that made combination impossible.

Wikipedia found itself isolated. Wikitravel, Citizendium, WikiEducator—newer wikis all chose CC-BY-SA. The free content ecosystem was fragmenting along licensing lines. Wikipedia's GFDL commitment, made when no alternative existed, had become a barrier to collaboration.

### Friction as Feature

This is where the standard narrative needs revision.

The friction was the protection. Remove the friction, and you remove what kept the commons safe.

The same mechanisms that made GFDL cumbersome for legitimate reuse also made it cumbersome for extraction. The requirements that annoyed teachers also deterred scrapers.

Consider what the GFDL's "inconvenient" requirements actually did. Full license text with every copy discouraged bulk extraction by creating overhead. Complete history preservation created an audit trail and complicated laundering of content. Principal author attribution kept human creators visible. Transparent copy availability ensured source remained accessible. ShareAlike for all derivatives meant anything built on Wikipedia inherited these burdens.

The teacher couldn't make her flier without five pages of legalese. But a commercial entity couldn't scrape Wikipedia's entire corpus without those same five pages—for every article.

The complaints measured convenience for legitimate small-scale reuse. They did not measure protection against industrial-scale extraction.

When the push to "simplify" licensing succeeded with the 2009 migration to CC-BY-SA, what was gained in convenience was lost in protection. CC-BY-SA's flexible attribution requirements—a link suffices, no author list required—made legitimate reuse easier. They also made extraction easier.

What hindered the teacher would have hindered the AI company. When the barriers came down, both were freed.

## Escape Hatch

Wikipedia's copyright notice included a standard clause: content was licensed under "GFDL 1.2 or any later version published by the Free Software Foundation."

This routine provision, included without much thought, gave the FSF power to create new versions that Wikipedia could automatically adopt. As Jonathan Corbet noted in LWN.net:

"The other thing to be aware of is just how much power the 'or any later version' text puts into the hands of the FSF. The license promises that later versions will be 'similar in spirit,' but the GPLv3 debate made it clear that similarity of spirit is in the eye of the beholder."

The FSF would use this power to create GFDL 1.3, which included Section 11—a time-limited provision enabling "Massive Multiauthor Collaboration Sites" to relicense to CC-BY-SA. Wikipedia's "exit permit" was granted by the same organization that wrote the original walls.

The mechanism that enabled escape established a precedent: license terms can evolve in ways contributors never anticipated. If "or any later version" could enable fundamental relicensing, what else might it enable?

The FSF had built what seemed like permanent walls around the shared resource. But it had left a key in the lock—a power to change the license if times required evolution. This power made sense for software. But for a collective work, it created a vulnerability: if the institution guarding it changed its mind, or became entangled with extractive interests, the walls could be remade from within.

### Pragmatist View

From the Wikimedia Foundation's perspective, and the mainstream Wikipedia community's, the push to escape GFDL was simply practical.

The license was genuinely unsuitable for Wikipedia's use case. The friction was real—documented in thousands of frustrated emails and forum posts. CC-BY-SA preserved copyleft while enabling practical reuse. The migration, when it came, would be necessary for Wikipedia's mission of spreading free knowledge.

"When I started Wikipedia, Creative Commons did not exist," Wales explained. "The Free Documentation License was the first license that demonstrated well how the principles of the free software movement could be applied to other kinds of works." But times had changed. CC-BY-SA was "a more generic license that meets the needs of Wikipedia today."

### Purist Concern

A minority view questioned whether simplification was worth the cost.

The "or any later version" clause was not understood, when contributors licensed their work, to authorize such fundamental changes. Contributors who chose GFDL in 2001 or 2005 did not consent to CC-BY-SA; they consented to trust the FSF with future versions.

Chris Frey wrote to the FSF calling the forthcoming migration a "breach of trust." Slashdot commenters worried: "It is slightly chilling for anyone using another FSF license. You can omit the 'or later versions' license and have the possibility that the later versions of other FSF licenses will

be incompatible with your version. Or you can include it and have the possibility that the FSF will decide to grant an exemption for a specific large organisation and allow them to relicense your work."

The concern was not that CC-BY-SA was bad. It was that contributor expectations could be overridden when institutional interests aligned.

## FSF Defense

Richard Stallman defended the approach:

"We have never asserted that we will not change our licenses, or that we will never make changes like this one. Rather, our commitment is that changes to a license will stick to the spirit of that license and uphold the purposes for which it was written."

Both GFDL and CC-BY-SA were copyleft. Both required share-alike. The essential freedoms were preserved. The FSF had authority over future versions precisely to enable necessary evolution.

"The relicensing option in GFDL 1.3 is fully consistent with the spirit and purpose of the GFDL," Stallman wrote.

## Synthesis

The debate revealed a tension between license stability and license evolution that was never resolved.

The pragmatists won: the migration would happen, friction would be reduced, Wikipedia would join the broader CC ecosystem. The purists' concerns went unlitigated. No court ever ruled on whether "or any later version" authorized such fundamental relicensing.

But the precedent was set: when institutions agree, contributor expectations can be adjusted.

## Path Opens

On December 1, 2007, the Wikimedia Foundation Board voted 5-0 to request that the Free Software Foundation modify the GFDL to enable migration to CC-BY-SA.

Jimmy Wales called it "a party to celebrate the liberation of Wikipedia."

The FSF began a year of private negotiations with the Foundation, Creative Commons, and the Software Freedom Law Center. On November 3, 2008, GFDL 1.3 was released. It included Section 11—the "exit permit."

Jonathan Corbet called it "a legal hack": "Legal codes, like other kinds of code, have a certain tendency to pick up cruft as they are patched over time. In this case, the FSF has added a special, time-limited hack which lets Wikipedia make a graceful exit from the GFDL license regime."

The provisions were precise: content must originate on a "Massive Multiauthor Collaboration Site," must be under GFDL with "or any later version," with no cover texts or invariant sections. External GFDL content must have been added before November 1, 2008, and migration must occur before August 1, 2009.

The escape hatch was open. The question was what Wikipedia was escaping to.

## Return to Teacher

The teacher who couldn't make her photosynthesis handout would soon be able to share Wikipedia content more easily. A link would suffice for attribution. No five-page appendix required.

This was genuine liberation from genuine friction.

But something else was removed with the barriers. The mechanisms that made GFDL burdensome for the teacher also made it burdensome for anyone who wanted to harvest Wikipedia at scale. When the overhead vanished for legitimate small-scale reuse, it vanished for industrial harvesting too.

The teacher gained convenience. The AI companies that did not yet exist would gain something larger.

### Bridge Forward

The community vote would begin in April 2009. Chapter 7 examines what was actually decided—by whom, representing how many contributors, with what understanding of the consequences.

The GFDL era was ending. The CC-BY-SA era was beginning. What once protected was recast as mere inconvenience. The simplification that liberated the teacher would also enable what came next.

The escape hatch was open. The question no one asked: escape to what, and for whom?

CHAPTER SEVEN

# The First Loosening

> How could a small part of the community vote for a license change affecting all contributing right owners?
>
> — Herkko Hietanen, WIPO Magazine

On December 1, 2007, at a Wikipedia community gathering in San Francisco, Jimmy Wales made an announcement:

"Creative Commons, Wikimedia and the FSF just agreed to make the current Wikipedia license compatible with Creative Commons. This is the party to celebrate the liberation of Wikipedia."

The crowd cheered. These were people who understood licensing—who had felt the friction of GFDL, who wanted Wikipedia to join the broader Creative Commons ecosystem, who saw the migration as progress.

But most Wikipedia contributors were not at this party. They were scattered across the world, connected only by their shared project, unaware that decisions about their work's legal status were being negotiated in private.

San Francisco was not an accident. The people who shaped Wikipedia's future lived here, worked here, hired from here. The editors scattered across the world—the retired teacher in Ohio, the physics student in India, the hobbyist in rural England—were represented by proxy, if at all. Geog-

raphy was destiny: decisions about the global commons were made by a local elite.

The people in that room represented Wikipedia's inner circle. The millions who had contributed over the previous seven years—the anonymous IP editors, the casual contributors who had added a sentence here or fixed a typo there, the early editors who had moved on, the editors who had died—were not present and would not be asked.

The party celebrated liberation. The question was: liberation for whom?

## Secret Negotiations

The path from GFDL to CC-BY-SA required a year of private talks between four organizations: the Free Software Foundation, the Wikimedia Foundation, Creative Commons, and the Software Freedom Law Center.

Richard Stallman explained the secrecy: "We did this to keep the option limited in applicability and avoid the possibility of wholesale relicensing of other GFDL-covered material."

The negotiations were tailored specifically to Wikipedia's situation. The FSF was not opening a general exit from GFDL; it was creating a one-time window for a specific type of project.

The participants had overlapping interests. The FSF wanted to support Wikipedia without undermining GFDL generally. The Wikimedia Foundation wanted to escape GFDL's friction. Creative Commons wanted to unify the free content ecosystem under its licenses. The Software Freedom Law Center provided legal expertise to navigate the technical requirements.

The result was GFDL 1.3, released November 3, 2008, with Section 11—the escape clause.

## Exit Permit

Section 11 created the concept of "Massive Multiauthor Collaboration Sites" (MMC Sites)—wikis, essentially—and granted them a time-limited option to relicense to CC-BY-SA.

The requirements were specific: content must originate on an MMC Site (not imported from external GFDL sources after November 1, 2008), must be licensed under GFDL with the "or any later version" clause, with no cover texts or invariant sections, and migration must occur before August 1, 2009.

Jonathan Corbet at LWN.net called it "a legal hack": "Legal codes, like other kinds of code, have a certain tendency to pick up cruft as they are patched over time. In this case, the FSF has added a special, time-limited hack which lets Wikipedia make a graceful exit from the GFDL license regime."

The mechanism was elegant if you accepted its premises. The "or any later version" clause, which contributors had included routinely since 2001, gave the FSF authority to create new versions. GFDL 1.3 used that authority to enable migration to CC-BY-SA. The chain was complete: original license → "or any later version" → GFDL 1.3 → CC-BY-SA.

## Version Power

The entire migration depended on a clause that most contributors had never examined.

Wikipedia's copyright notice, since 2001, had licensed content under "GFDL 1.2 or any later version published by the Free Software Foundation." This was standard language in free software licensing, included routinely to allow license evolution without requiring re-consent from every contributor.

Jonathan Corbet, writing at LWN.net, had noted the power this clause granted: the FSF's promise that later versions would remain "similar in

spirit" meant little when "the GPLv3 debate made it clear that similarity of spirit is in the eye of the beholder."

Contributors in 2001 trusted the FSF with future versions. They did not specifically consent to migration to a different license family. The chain of authority was legally valid—probably—but the consent was attenuated.

## Vote Mechanics

The community vote ran from April 12 to May 3, 2009—twenty-one days.

Eligibility required 25 or more edits to any Wikimedia project before March 15, 2009. Anonymous IP editors were excluded. Inactive contributors who had drifted away were unlikely to hear about the vote. Deceased contributors could not participate.

Notification came through CentralNotice, which displayed banners to logged-in users. Those who weren't logged in—or weren't editing during the voting period—might never learn the vote was happening.

The question asked voters to approve three changes: dual-licensing under CC-BY-SA 3.0 and GFDL, making CC-BY-SA the primary license going forward, and adopting a foundation-wide attribution policy.

Software in the Public Interest (SPI), a nonprofit providing infrastructure for open source projects, managed the voting process.

## Results

The vote produced the results shown in Table 7.1.

Rejected ballots included 930 multiple votes from single accounts and 256 votes from the same person using different wiki accounts.

Excluding abstentions, 87.9% of those who expressed an opinion supported migration. The result was celebrated as decisive democratic legitimacy.

| Metric | Count |
|---|---|
| Total ballots submitted | 18,692 |
| Valid ballots (after review) | 17,462 |
| Yes votes | 13,242 (75.8%) |
| No votes | 1,829 (10.5%) |
| No opinion | 2,391 (13.7%) |

**Table 7.1.** 2009 CC-BY-SA migration vote results

### The Bet

Every step in this process—the "or any later version" clause, the secret negotiation, the community vote—represented a bet that good faith would continue. The contributors who licensed under GFDL in 2001 bet that the FSF would not fundamentally change the terms. The Wikimedia Foundation bet that the commons could migrate to a simpler license without losing protection.

All of these bets assumed enforcement would follow. That an organization would actually use the legal tools to ensure derivatives remained free.

### Representation Problem

WIPO Magazine—the publication of the World Intellectual Property Organization—asked the question that celebration obscured. Their epigraph to this chapter cut to the heart of the matter: how could a fraction of the community bind all contributors?

Consider the numbers. By 2009, Wikipedia had millions of historical contributors since 2001, tens of millions of article edits across all projects, millions of registered accounts, and countless anonymous IP contributions (excluded from voting).

The vote produced roughly 17,500 valid ballots. About 10,000 people voted yes (excluding abstentions).

Ten thousand people relicensed the work of millions.

Who was not represented? Casual contributors with fewer than 25 edits couldn't vote—but their contributions were included. Anonymous IP editors, whose content comprised a significant portion of Wikipedia, had no voice. Inactive contributors who had moved on by 2009 had no idea the vote was happening. Deceased contributors—eight years of mortality since 2001—could not participate. And many eligible voters never learned about the vote at all.

WIPO's legal analysis raised a deeper concern:

"It is a matter of legal debate as to whether clauses covering future possibilities are valid in the case of licensors who were unaware of exploitation or licensing options not yet invented at the time of initial distribution. Most European countries have laws nullifying such agreements."

The "or any later version" clause enabled the chain. But contributors in 2001 did not anticipate—could not have anticipated—that it would be used to migrate to a different license family entirely.

## What Changed

The migration was not merely a format change. CC-BY-SA has substantively lighter requirements than GFDL. Under GFDL, every copy required the full license text—approximately five pages. CC-BY-SA dropped this requirement. GFDL required preserving complete modification history; CC-BY-SA did not. GFDL required listing at least five principal authors; CC-BY-SA allowed a link to the source to suffice. GFDL required providing a transparent copy (source format); CC-BY-SA dropped this entirely.

The same characteristics that made GFDL inconvenient for the teacher with her photosynthesis handout also made it inconvenient for anyone who wanted to extract Wikipedia at scale.

The "simpler" license was also a weaker license. The friction that was removed for legitimate reuse was also removed for industrial extraction.

### Pragmatist Case

From the Wikimedia Foundation's perspective, and the majority of active editors', the vote was legitimate and necessary.

GFDL was genuinely unsuitable for Wikipedia's use case. The friction was real, documented, and widely complained about. CC-BY-SA preserved copyleft while enabling practical reuse. The community voted; 88% (excluding abstentions) approved. This was democracy in action.

The Foundation's announcement emphasized continuity: "The volunteers who work on Wikimedia projects have very strongly supported making their contributions available under the Creative Commons Attribution/Share-Alike License (CC-BY-SA) in addition to the GNU Free Documentation License (GFDL). Updating our license terms will support Wikimedia's charitable mission."

Erik Moeller, then Deputy Director, called it "the culmination of a long process."

### Critic Case

A minority raised concerns that went largely unaddressed.

Contributors who chose GFDL in 2001 trusted the FSF with future versions of GFDL. They did not specifically consent to migration to an entirely different license family.

Chris Frey wrote to the FSF calling the migration a "breach of trust." On Slashdot, others saw troubling precedent: anyone using an FSF license with the "or later" clause had just learned their work could be relicensed without their participation.

WIPO's analysis crystallized the concern: millions of contributors were relicensed without their knowledge or consent. A 17,500-person vote cannot bind a community of millions when most of those millions never participated.

The fact that no legal challenge materialized does not mean the concerns were invalid. It means no one sued.

### FSF Position

Richard Stallman defended the migration:

"We have never asserted that we will not change our licenses, or that we will never make changes like this one. Rather, our commitment is that changes to a license will stick to the spirit of that license and uphold the purposes for which it was written. This is what we have done with GFDL 1.3."

Both GFDL and CC-BY-SA were copyleft licenses. Both required share-alike. The essential freedoms were preserved. The FSF had authority over future versions precisely to enable necessary evolution.

The FSF's position: the migration was legitimate because the mechanism was legitimate and the destination preserved copyleft's spirit.

### Synthesis

The debate reveals an unresolved tension between organizational authority and individual consent.

The pragmatists won in 2009. The vote happened; the migration proceeded; Wikipedia joined the CC ecosystem. The critics' concerns went unlitigated. No court ruled on whether "or any later version" authorized this fundamental relicensing.

The precedent stands: when organizations agree, license terms can evolve beyond what individual creators foresaw. The mechanism that enabled escape from GFDL's friction also demonstrated that collective work could be relicensed through processes most creators never joined.

### Implementation

On May 21, 2009, the Wikimedia Foundation Board passed the Licensing Update Approval Resolution. Michael Snow, Board Chair, announced: "Updating our license terms will support Wikimedia's charitable mission."

On June 15, 2009, the migration took effect. CC-BY-SA 3.0 became the primary license. GFDL was retained as a secondary option (dual-licensing). New contributions could be CC-BY-SA only.

On August 1, 2009, the GFDL 1.3 migration window closed permanently. Any wiki that hadn't migrated lost the opportunity forever.

The migration was complete.

### Return to Party

Jimmy Wales had called it "the liberation of Wikipedia." In one sense, it was. The flier problem was solved. Audio and video producers could breathe. The free content ecosystem was unified.

But the celebration at that San Francisco party was for those who understood licensing, who participated in the process, who believed the vote was legitimate. The millions who had contributed under GFDL—who believed their work would remain under those terms—were not at the party. They were not asked. They were relicensed anyway.

The 2009 switch was the first loosening.

### Pattern Emerges

The pattern that played out in 2009 would recur: the shared resource had legal protections, those protections proved inconvenient, institutions negotiated to reduce them, consent was obtained through mechanisms most contributors never participated in, and the result was weaker terms framed as progress.

This pattern will repeat when AI companies interpret "derivative work" and "share-alike" in ways that benefit their extraction:

**The 2009 logic:** CC-BY-SA preserves copyleft's spirit while simplifying compliance.

**The 2020s logic:** AI training doesn't create "derivative works," so share-alike doesn't apply.

Both arguments enable commercial use that the original creators did not foresee. Both are defended as consistent with the license's "spirit." Both benefit organizations at volunteer expense.

### Bridge Forward

The 2009 switch matters for Part IV because it planted the later exploitation. The looser attribution requirements made AI training more feasible. The precedent of institutional override made enforcement questions harder to raise. The community's acceptance of the vote's legitimacy set expectations about future changes.

When the Wikimedia Foundation later acknowledged that Enterprise customers were "not compliant" with CC-BY-SA (April 2024), the echoes of 2009 were audible: organizational interests prevailed, and no one sued because the community had already acquiesced.

Chapter 8 examines what CC-BY-SA enforcement actually looks like—the mechanisms that exist, the gaps that don't get closed, and the structural passivity that would prove consequential when AI companies arrived.

The switch was complete. The question was what had been gained—and what had been surrendered.

CHAPTER EIGHT

# The Gaps in Practice

> Wikipedia neither condemns, nor promotes, such efforts at enforcing the licenses, and is not responsible for any consequences.
>
> — Wikimedia Meta-Wiki

In September 2005, a friend contacted John Seigenthaler with disturbing news. Something was on the internet about him. Something bad.

Seigenthaler was 78 years old. He had been the founding editorial director of USA Today, an assistant to Attorney General Robert Kennedy, a respected figure in American journalism for half a century. He navigated to Wikipedia—a website he barely knew—and found his biography.

The text accused him of involvement in the assassinations of both Kennedy brothers.

"John Seigenthaler Sr. was the assistant to Attorney General Robert Kennedy in the early 1960's. For a brief time, he was thought to have been directly involved in the Kennedy assassinations of both John, and his brother, Bobby. Nothing was ever proven."

For 132 days, this libel had been available to anyone who searched his name. A false accusation of political murder, presented as encyclopedic fact.

Seigenthaler investigated. The vandalism had been added on May 26, 2005, by an anonymous IP address. No one had noticed. No one had checked. For four months and twelve days, visitors to Wikipedia's fifth-most-visited biography learned that a living American had been suspected of killing two presidents.

"When I was a child, my mother told me, 'Consider the source,'" Seigenthaler wrote in USA Today on November 29, 2005. "For most of my 78 years, I have relied on traditional media to provide information. Wikipedia changed that."

His verdict: "a flawed and irresponsible research tool."

## Nature Validation

Two weeks after Seigenthaler's op-ed, Nature published a study that appeared to contradict everything he said.

The journal had commissioned a comparison of Wikipedia and Encyclopedia Britannica. Experts reviewed 42 article pairs, blinded to which encyclopedia produced each version. The result: Wikipedia had 3.9 errors per article; Britannica had 2.9. Both had four "serious errors"—misinterpretations of important concepts.

"Wikipedia comes close to Britannica in terms of the accuracy of its science entries," Jim Giles reported.

Britannica responded with a 20-page "Fatally Flawed" rebuttal, calling Nature's methodology incompetent. Nature stood by its findings. The academic consensus shifted: Wikipedia was credible. The open model worked.

The contradiction between Seigenthaler and Nature was more apparent than real. Both were correct about different things. Wikipedia's science articles benefited from many expert contributors and active monitoring. Seigenthaler's biography—a low-traffic article about a regionally-known figure—had no watchers. The wisdom of the crowd varied with the size of the crowd watching.

### First Restriction

On December 5, 2005, Jimmy Wales announced a change: anonymous users could no longer create new articles on the English Wikipedia.

"As an experiment, we will be turning off new pages creation for anonymous users," he wrote. He emphasized his continued belief in anonymous editing: "I am a firm believer in the validity of allowing anons to edit."

But the line had been crossed. "Anyone can edit" now came with an asterisk. The Seigenthaler incident had revealed a structural vulnerability that philosophy couldn't fix.

The Biographies of Living Persons (BLP) policy tightened. Contentious unsourced material about living people would be removed immediately, without waiting for discussion. Oversight functions were created—the ability to hide certain content from the revision history entirely, making libel invisible even to administrators.

The quality crises were being addressed. The enforcement gap was not.

### Who Enforces?

The 2009 migration to CC-BY-SA solved the friction problem. It did not solve—did not even address—the enforcement problem.

Consider who might enforce Wikipedia's license:

**The Wikimedia Foundation cannot enforce.**

The Foundation does not own contributor copyrights. Each editor retains copyright in their contributions and licenses them to the public under CC-BY-SA. The Foundation merely hosts the platform.

Meta-Wiki states this explicitly: "The Wikimedia Foundation does not own Wikipedia content and cannot grant permission to reuse it."

**The Foundation disclaims enforcement responsibility.**

This is not evasion. It reflects the project's legal structure. The Foundation is a platform, not a publisher. It has no standing to enforce copyrights it does not hold.

**Creative Commons does not enforce.**

"Creative Commons is not a law firm and cannot represent you or give you legal advice."

CC provides licenses. It does not enforce them. The organization explicitly states it will not sue on behalf of licensors.

**Individual enforcement is impractical.**

A contributor who wrote one paragraph in a 50-paragraph article must identify the violation, verify their specific content was used, determine no fair use defense applies, decide whether to pursue enforcement, bear the costs of enforcement, and potentially receive minimal damages. For most contributors, this calculus does not favor enforcement.

The system assumed good faith. When reusers comply voluntarily—proper attribution, share-alike for derivatives—everything works. When they don't, no institutional enforcer exists.

This was not inevitable. The Linux community created enforcement organizations—Software Freedom Conservancy, the Software Freedom Law Center—that continued GPL enforcement even as the Linux Foundation took corporate money. The difference: open source maintained ideological diversity around licensing enforcement. Copyleft purists remained in the community and in organizations. The Wikimedia Foundation's professional culture—drawn from tech and nonprofit networks that saw litigation as unseemly—had no equivalent enforcement constituency.

The system worked tolerably when violations were scattered. More fundamentally, it reflected a structural problem in commons governance: there was no commons government. The Wikimedia Foundation hosted the platform but disclaimed ownership. Creative Commons wrote the licenses but would not enforce them. Individual contributors could enforce but had no means. The commons had rules but no enforcer—a design flaw that would prove catastrophic when extraction became industrial.

### Case Law

Courts have addressed CC license enforcement, establishing both that the licenses work and that they have limits.

**Curry v. Audax (Netherlands, 2006):** Adam Curry, former MTV VJ, posted photos on Flickr under CC-BY-NC-SA. A Dutch tabloid, Weekend, published them without attribution. The court ruled: "Audax has not observed the conditions stated in the License." Fine of 1,000 euros per photo for future violations.

Significance: First court enforcement of a Creative Commons license. The licenses are legally binding.

**Drauglis v. Kappa Map Group (US, 2015):** A photographer licensed images under CC-BY-SA. Kappa Map used one on an atlas cover. The photographer claimed the atlas was a derivative work requiring SA licensing.

The court disagreed. The atlas was a "collective work," not a "derivative work." Share-alike only applies to derivatives. Using a CC-BY-SA image in a larger work doesn't necessarily trigger SA requirements.

Significance: Narrowed the scope of share-alike. Not everything including CC content triggers copyleft.

**Philpot v. Media Research Center (US, 2018):** A photographer licensed images under CC-BY. The Media Research Center used them without attribution. Clear violation.

The court found fair use anyway. The use was "transformative"—concert photos repurposed for political commentary. Fair use defeated the license claim.

Significance: Fair use can override CC enforcement. A violation can still be non-infringing if transformative.

### Trolling Paradox

Some photographers discovered a business model: post images under CC licenses, use automated tools to find violations, send demand letters threatening statutory damages, settle for payments exceeding actual harm.

Marco Verch, a German photographer, filed hundreds of lawsuits for minor attribution failures. He worked with attorneys whom American judges labeled "copyright trolls." The strategy generated revenue from people who thought they were using "free" images.

Creative Commons responded with a "Statement of Enforcement Principles" in 2022:

"'License-enforcement-as-business model' is a perversion of the founding ideals of Creative Commons, which was intended to mitigate draconian copyright laws and penalties."

The statement created a paradox: aggressive enforcement was discouraged as harmful to "trust in CC licensed content." But passive enforcement meant widespread non-compliance persisted.

CC-BY-SA 4.0 includes an automatic cure provision: violations fixed within 30 days of discovery reinstate the license. This favors correction over punishment—but it also means violators face minimal risk if they simply add attribution when caught.

## What AI Changes

All the enforcement mechanisms assume human reuse: a photo on a website can be found, identified, and compared; an article copied to another site can be diff'd and traced; a book including Wikipedia excerpts can be examined. When an LLM trains on Wikipedia, none of these assumptions hold.

**Is the model a "derivative work"?**

CC-BY-SA requires that "Adapted Material" (derivatives) be released under the same license. But is a trained model a "derivative" of its training data?

Legal scholars Szkalej and Senftleben analyzed this question in JIPITEC (2024):

"Share Alike/CopyLeft licenses are largely ineffective when materials licensed under them are used to train AI models."

Model weights are not "material... translated, altered, arranged, transformed" in an obvious copyright sense. The license assumed derivatives would be recognizable modifications.

**Where is attribution in generated text?**

Traditional attribution appears in captions, credits, metadata. When ChatGPT answers using Wikipedia information, no Wikipedia link appears, no editor names are credited, no CC-BY-SA notice is shown, and the user cannot know the source. The license requires attribution "reasonable to the medium." What is "reasonable" for an AI chatbot? The license does not say.

**Is the output share-alike?**

If AI output is influenced by Wikipedia training, must it be CC-BY-SA licensed?

Creative Commons' own statement: "It has been argued that all outputs generated by an AI are derivatives of all of the inputs, but this argument holds no water—when you ask Bard or Claude to write a poem, those words aren't a derivative of any particular input used in training."

**The Exception Override**

CC-BY-SA 4.0, Section 2(a)(2), contains a provision that may render all of this moot:

"For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions."

Translation: if AI training qualifies as fair use (US) or text and data mining (EU), the license conditions don't apply at all. No attribution required. No share-alike required.

The US Copyright Office's May 2025 report acknowledged that AI training may be fair use in many circumstances. Judge Alsup in Bartz v. Anthropic (June 2025) described AI training as "transformative—spectacularly so."

If this view prevails, CC-BY-SA becomes unenforceable against AI companies—not because the license is weak, but because the use falls outside copyright's domain entirely. The legal armor that was meant to protect the commons becomes irrelevant. The commons is open. No walls at all.

### Skeptic's View

From John Seigenthaler's perspective, and those who shared his concerns, Wikipedia's openness was a liability.

Anonymity enabled malice without accountability. The Kennedy assassination hoax proved the system could not police itself. Expert review was essential for reliable knowledge. "Anyone can edit" was a bug, not a feature.

The evidence: the vandal, Brian Chase, was identified through IP tracing but never prosecuted—no applicable law existed. The libel spread to Answers.com and Reference.com, misinformation propagating through automated content sharing. Low-traffic articles had minimal oversight.

### Defender's View

From Wales's perspective, and the core Wikipedia community's, the system was self-correcting—eventually.

The Nature study proved quality rivaled professionals. BLP policy and oversight addressed the worst cases. Vandalism on high-traffic articles was caught in seconds. Open editing enabled unprecedented knowledge creation. Britannica made errors too; the difference was Wikipedia could fix them.

The evidence: near-parity with Britannica in the Nature study. Rapid response on monitored articles. Policy evolution demonstrating community learning.

### Structural Critique

Academic researchers saw something both camps missed.

The "wisdom of the crowd" was uneven. High-traffic articles got attention; low-traffic articles were vulnerable. The biography fiasco revealed a structural flaw, not just a bad actor. Quality depended on watchlist dynamics, not magical self-correction.

Vandalism on "George W. Bush" was fixed in seconds—many watchers. Vandalism on "John Seigenthaler Sr."—no watchers for over four months. The distribution of editing was highly unequal.

### Synthesis

The episode forced Wikipedia to choose: maintain pure openness, or restrict editing to protect quality. The community chose quality. Anonymous article creation ended. BLP policy tightened. Oversight expanded.

This was the first in a series of restrictions that would reshape "anyone can edit." Each made sense in isolation. Cumulatively, they established that openness has limits.

The enforcement question followed the same pattern. When reusers violated CC-BY-SA, the response was similarly diffuse: individual contributors could enforce, but most didn't. The Foundation disclaimed responsibility. Creative Commons opposed aggressive enforcement.

This arrangement held when violations were scattered and harm was diffuse. It would not hold when a handful of companies trained on the entire corpus.

### Gap Established

By the end of this period—roughly 2005-2012—Wikipedia had validated its quality through the Nature study, qualified its openness by restricting anonymous article creation and tightening BLP policy, established enforcement passivity by disclaiming responsibility and leaving enforcement to individual contributors who rarely pursued it, and built the corpus AI

would extract—every article, every edit, every revision, accumulated and licensed under CC-BY-SA.

The "Fun and Games" period delivered Wikipedia's promise. The encyclopedia grew. Quality improved. The community matured.

But a pattern of non-enforcement was also established. Violations were widespread among human reusers—mirror sites without attribution, books reprinting content without license notices. The system tolerated this because harm was diffuse and pursuing violations was costly.

### Return to Seigenthaler

John Seigenthaler's biography was corrected. Brian Chase apologized. The BLP policy prevented future incidents of similar duration.

But the underlying dynamic—content created by volunteers, enforcement delegated to individuals, organizational inaction—remained unchanged. The lesson was interpreted as a quality problem, fixed with quality policies. License compliance received no comparable attention.

### Bridge Forward

The Institution's Turn follows the Wikimedia Foundation's growth. From 2003's incorporation to 2021's Enterprise launch, the Foundation would transform from volunteer project to sprawling nonprofit.

The passivity documented here would persist as the Foundation grew. The organization would add staff, raise funds, build infrastructure—but never build enforcement capacity. When AI companies began training on Wikipedia at industrial scale, the institutional passivity established in this period would prove consequential.

### Part II Revealed

The Licensing Gamble has traced a progression: Chapter 5 introduced the copyleft promise—Stallman's vision, the legal armor built from grief. Chapter 6 covered the GFDL era—friction that protected but inconvenienced.

Chapter 7 examined the 2009 switch—simplification that was also loosening. This chapter documented CC-BY-SA in practice—the enforcement gap that no one closed.

The arc is clear: protection mechanism → inconvenient friction → simpler alternative → unenforced alternative.

The next question: how did the Foundation that "does not own Wikipedia content" grow into a nine-figure organization? And what happened when it started selling access to the companies whose compliance was questionable?

The Institution's Turn examines the Foundation's growth. The AI Reckoning shows what happened when AI companies arrived.

### License Without Enforcement

The CC-BY-SA license promised that Wikipedia content would remain "free" forever—free to use, free to modify, free to share. The enforcement mechanisms assumed good faith reuse.

When commercial entities arrived with industrial-scale extraction and legal teams arguing fair use, the promise met a gap that no one had closed.

Stallman had built legal armor to protect against Symbolics-style betrayal. The armor worked when someone was willing to wear it, when the enemy was recognizable, when the battlefield was familiar.

When the enemy became statistical weights and the battlefield became fair use doctrine, the armor sat unworn. The institution that inherited Stallman's framework had never prioritized enforcement. Now it had taken the extractors' money.

The copyleft promise was made. The enforcement never followed.

# Part III

# The Institution's Turn

CHAPTER NINE

# The Foundation

> Servers cost money. Servers eat electricity. Servers need repairs.
>
> — Wikipedia fundraising appeal, October 2003

On June 20, 2003, Jimmy Wales filed six pages of documents with the State of Florida.

The Articles of Incorporation were bureaucratic boilerplate: a nonprofit corporation, formed for educational and charitable purposes, headquartered in St. Petersburg. Three incorporators: Wales, Tim Shell, Michael Davis. All Bomis co-founders. The filing fee was minimal. The paperwork was routine.

But these six pages created the legal vessel that would govern the world's largest encyclopedia. Every article that Wikipedia's hundreds of thousands of contributors would write—every edit, every image, every reference—would fall under the stewardship of the entity born that day. The Wikimedia Foundation would become the institutional home of free knowledge.

The announcement to the Wikipedia community was brief. Wales explained that the Foundation would hold the project's trademarks, domains, and server hardware. It would be an "independent charitable entity,"

separate from Bomis. The community had worried about commercial pressures; the Foundation was supposed to protect against them.

This was the origin story as typically told: a founder's generosity creating an institution to serve the commons.

The story was true. It was also incomplete.

### Bomis Subsidy Era

Wikipedia launched in January 2001 under the corporate umbrella of Bomis, Inc.—Jimmy Wales's for-profit company.

Bomis provided everything: servers, bandwidth, technical support, Larry Sanger's salary. The project existed because Wales was willing to fund it from his company's revenue. That revenue had an uncomfortable origin. Bomis ran a web portal that included "Bomis Babes"—adult content that, while not pornographic, catered to what Wales later described as "a guy-oriented search engine... similar to Maxim magazine."

The community knew about this background and was broadly unbothered. What mattered was that Wikipedia existed and was growing. By 2003, the project had over 100,000 articles in English and was expanding to dozens of other languages.

But Bomis was struggling financially. The dot-com crash had hit hard. The company could not subsidize Wikipedia indefinitely. And the community was increasingly opposed to the obvious alternative—advertising.

### Why Incorporate?

A nonprofit foundation solved multiple problems:

**Legal liability.** Who was responsible for Wikipedia's content? Bomis? Wales personally? The contributors? A foundation with clear legal status would provide protection.

**Trademark control.** The "Wikipedia" name had value. Someone needed to hold it formally.

**Tax-deductible donations.** Readers who wanted to support Wikipedia couldn't get tax deductions if they sent money to Bomis. A 501(c)(3) nonprofit would enable charitable giving.

**Organizational credibility.** Wikipedia aspired to be a serious reference work. Operating under an adult-content company complicated that aspiration.

**Independence.** The community expected separation from commercial interests. A nonprofit would formalize that separation.

### Asset Transfer

Wales's decision to create the Foundation was genuine charitable giving.

He transferred the domains (wikipedia.org, wikipedia.com), the trademarks, the server hardware, and all Wikipedia-related intellectual property to the new nonprofit. He retained nothing. A less generous founder could have structured things differently—keeping the trademarks, licensing them back, maintaining control through ownership rather than governance.

Wales did not do this. By the time of the transfer, Bomis had invested roughly $500,000 in Wikipedia. Wales donated all of it.

This generosity was real. It would also shape the Foundation's governance in ways the community didn't fully anticipate.

### Board Structure

Wales incorporated the Foundation as sole incorporator. In January 2004, he appointed Tim Shell and Michael Davis as additional trustees—his Bomis co-founders.

The initial board was entirely Bomis. No community representation existed. No mechanism for community input had been created. The people who had built Wikipedia's content—a growing army of volunteer editors—had no voice in governing the institution that now held their work.

This was not unusual for nonprofits. Founders typically control the organizations they create. But Wikipedia was not a typical nonprofit. It was a collaborative project where the value came from contributor labor. The governance structure gave those contributors no formal role.

This mattered profoundly because the Foundation was not a company managing its own product. It was a steward of assets contributors had built. Those contributors had no formal representation.

### First Fundraiser

In October 2003, Wikipedia's servers went down. The crisis triggered the Foundation's first fundraising campaign.

The appeal was simple, echoing the words that would become the chapter's epigraph.

The goal was $20,000. Within a week, the community had donated over $30,000. By year's end, the Foundation had raised $56,666—enough to purchase nine new servers and begin building distributed infrastructure.

The success proved something important: volunteers would financially support what they had built. Wikipedia could sustain itself through small donations from readers and contributors. It didn't need advertising or commercial partnerships.

This modest fundraiser established the model that would eventually generate nine figures annually. It also established something else: dependence on donation revenue to maintain an institution that would grow far beyond what 2003's donors imagined.

### Community Representation

By 2004, concerns about the all-Bomis board had emerged. The community that wrote Wikipedia wanted a voice in its governance.

The Foundation responded with board elections. From May 30 to June 12, 2004, eligible Wikipedians voted for two trustee positions. Require-

ments were minimal: a 90-day-old account with editing activity. Angela Beesley and Florence Nibart-Devouard won.

This was progress. The community now had representation. But the structure remained tilted: three appointed trustees, two elected. The appointed members held the majority.

This ratio—appointed majority, elected minority—would persist. Even as the board expanded, the pattern held: those who built content had voice but not control.

### Tax-Exempt Wait

The Foundation operated for nearly two years without formal tax-exempt status. The IRS determination letter arrived in April 2005, retroactive to June 20, 2003.

What this delay revealed: the Foundation's limited initial capacity. This was not a sophisticated institution with professional staff and legal teams. It was a handful of volunteers managing paperwork for a rapidly growing website.

But it also revealed something else: the institution was already making decisions—about governance, about strategy, about Wikipedia's direction—without the formal charitable status that would later constrain it.

### What Structure Planted

The modest 2003 Foundation—$57,000 in assets, no paid staff, volunteers answering emails—would become a $271 million organization with 700 employees.

This growth was not corruption. It was the natural consequence of governing a top-ten global website. Wikipedia needed infrastructure, legal defense, strategic planning, community coordination. Someone had to do this work. Volunteers alone couldn't scale to meet the demand.

But the organizational architecture planted in 2003 shaped everything that followed.

The Founder's Seat, formalized in 2008, explicitly cannot be filled by anyone other than Jimmy Wales. Not even Larry Sanger, who designed Wikipedia's core policies and served as its first editor-in-chief. The Foundation's structure institutionalized Sanger's exclusion and Wales's permanence.

The pattern established that year—appointed leadership, limited community voice, founder control—persists. The question it raises: whose interests does the institution serve when community mission and institutional survival diverge?

### Wales's Position

From Wales's perspective, the Foundation's governance reflects legal and practical reality.

Sanger was an employee of Bomis, hired to run Nupedia and then Wikipedia. He was paid a salary. When Bomis could no longer afford that salary, his employment ended. He departed in March 2002—fifteen months before the Foundation was created.

Wales conceived the free encyclopedia project. Wales provided the funding. Wales owned the assets that were donated. The Foundation governs those assets. Its structure reflects that ownership.

"Co-founder" in the press release sense is different from "co-founder" in the ownership sense. Sanger contributed enormously to Wikipedia's early development. He did not own the infrastructure he helped build.

### Sanger's Position

From Sanger's perspective, the Foundation's structure rewrites history.

"I named the project, designed its policies, served as its chief editor for its first year, and wrote or commissioned its first articles," he has said. "If that doesn't make me a co-founder, the word has no meaning."

The Neutral Point of View policy that still governs Wikipedia? Sanger drafted it. Verifiability? No Original Research? Sanger formalized them.

The community processes that made Wikipedia work? Sanger designed them.

The Foundation was created after Sanger left—not because he chose to leave, but because Bomis ended his salary. He had no opportunity to participate in its formation. The Founder's Seat structure, created five years later, explicitly excludes him.

"It does seem that Jimmy is attempting to rewrite history," Sanger wrote in 2005. "But this is a futile process because, in our brave new world of transparent activity and maximum communication, the truth will out."

More importantly, Sanger had designed the processes that made Wikipedia work. Yet the institution created to serve those communities had no seat for him. The person who designed the rules was excluded from their administration. This would not be the last time the Foundation's structure diverged from the community's interests. It was only the first sign of a pattern: an organization built by volunteers, for volunteers, eventually administered without them.

### Broader Pattern

The co-founder dispute illustrates a deeper governance tension.

Wikipedia is a collaborative project. Its value comes from collective labor—not just Wales's funding or Sanger's policies, but a quarter million contributors adding knowledge article by article. Yet the institution that governs this collective work was not created by the collective. Its structure reflects legal ownership, not collaborative contribution.

The pattern would repeat as the Foundation grew. Decisions affecting volunteer labor would be made by appointed staff. Community consultation would occur, but community consent would not be required. The institution that existed to serve contributors would develop priorities that diverged from theirs.

This is not unique to Wikipedia. It is the nature of organizations. But for a project built on collaborative openness, the tension between that openness and institutional governance would prove consequential.

### What Changed

On June 20, 2003, Wikipedia gained an institutional home.

Legal protection was established. Trademarks were secured. A fundraising model was proven. The project that began as a side experiment under a web portal company became an independent nonprofit with a mission: the free encyclopedia.

But the governance structure tilted toward appointed leadership over community voice. The founder who retained legal control shaped the institution. The founder who designed community processes did not.

### Distance From Servers

The first fundraiser asked for $20,000 and received $30,000 in a week. The message was honest: servers cost money. The donors believed they were keeping Wikipedia online.

Twenty-two years later, the fundraising machine would bring in nearly $200 million annually. The Foundation would accumulate reserves exceeding its annual budget several times over. Some community members would vote that the fundraising banners were misleading—that the urgent appeals for survival masked an organization with years of runway.

The distance from "servers cost money" to surplus billions is the distance this book traces.

### Bridge Forward

The Foundation created in 2003 would demonstrate unprecedented collective power. Chapter 10 follows this institution to its peak: January 2012, when Wikipedia's blackout helped defeat SOPA. 162 million people saw a protest page. Legislation died.

That was the zenith. The institution born here in modest circumstances would prove it could mobilize its readers for political action. It would prove it could raise money at scale. It would prove it could defend itself.

What it would not prove—what remains unproven to this day—is whether it could deploy that power against extraction by the companies that would later become its commercial partners.

The seeds were present from the beginning: an institution created for the community but not by the community. A board where appointed members outnumber elected. A founder's permanent seat that no process can fill with anyone else.

The Foundation was designed as a trustee protecting the shared resource contributors had assembled. Whether an institution created to guard the commons could remain faithful when it developed its own survival imperatives—that question remained unanswered.

CHAPTER TEN

# The Fundraising Machine

> Imagine a World Without Free Knowledge.
>
> — Wikipedia blackout page, January 18, 2012

At 05:00 UTC on January 18, 2012, the English Wikipedia went dark.

In place of the familiar white page with its blue links and sprawling articles, visitors found a black screen with white text. No search box. No navigation. Just a stark message and a zip code lookup tool:

For over a decade, we have spent millions of hours building the largest encyclopedia in human history. Right now, the U.S. Congress is considering legislation that could fatally damage the free and open Internet. For 24 hours, to raise awareness, we are blacking out Wikipedia.

162 million people saw that page. Eight million looked up their congressional representatives. The phone lines to Capitol Hill overloaded. Within days, the Stop Online Piracy Act and PROTECT IP Act—legislation that had seemed inevitable weeks earlier—were shelved indefinitely.

Wikipedia had demonstrated something unprecedented: a volunteer project could mobilize political power at scale. The world's fifth-most-visited website had weaponized its reach for a cause.

This was the peak.

### Survival to Surplus

The Foundation that raised $57,000 in 2003 had become something else entirely.

In 2007, when Sue Gardner arrived as executive director, the Wikimedia Foundation had five employees and roughly $2.7 million in annual revenue. Wikipedia ran on servers that crashed regularly. The organization operated from St. Petersburg, Florida, far from the tech industry's centers of power.

Gardner had built CBC.ca into Canada's most popular news site. She understood how to grow organizations. Within seven years, she would transform the Foundation from a shoestring operation into the fastest-growing nonprofit in America by revenue growth.

The numbers tell the story, as shown in Table 10.1.

| Year | Revenue | Staff |
|---|---|---|
| 2007 | $2.7M | 5 |
| 2011 | $23M | 50 |
| 2014 | $60M | 200 |
| 2024 | $185M | 700 |

**Table 10.1.** Wikimedia Foundation growth, 2007–2024

The model was simple: banner appeals asking for small donations. The execution was sophisticated: A/B testing optimized every element. Colors, wording, suggested amounts, placement—everything was measured. The banners that performed best were deployed; the rest were discarded.

"Wikipedians do it for love, for mission-type reasons," Gardner explained. "They don't want to be paid. They want to be praised."

The fundraising machine had learned to convert that love into revenue.

### Wales Effect

In November 2010, the Foundation discovered its most powerful fundraising tool: Jimmy Wales's face.

"Please Read: A Personal Appeal From Wikipedia Founder Jimmy Wales"

The banner featured Wales looking directly at the camera, his expression earnest. The personal appeal format was unprecedented. Encyclopedia fundraising typically emphasized institutions, not individuals. This was different: a named person, making eye contact, asking for help.

The results were staggering. The 2010 campaign raised $16 million—the shortest fundraiser in Wikimedia history. The Wales banner was 15 times more effective than any alternative.

Culturally, the banner became a meme. "Jimmy Wales's creepy stare" generated parodies and discussions across the internet. The Foundation didn't mind. Whatever drove people to talk about Wikipedia also drove them to donate.

Subsequent testing revealed something curious: the personal appeal format worked regardless of who appeared. Brandon Harris, a Foundation engineer, became another effective face. The magic was in the format—the direct address, the named individual, the sense of personal connection—not the specific person.

But Wales remained the primary face. The founder had become the fundraising machine's most valuable asset.

### Decision to Act

The Stop Online Piracy Act was moving through Congress in late 2011. Its provisions would allow content owners to force internet service providers to block access to websites accused of hosting infringing material. Critics warned it would break the internet's architecture and enable censorship.

Wikipedia's community debated: should an encyclopedia take a political position?

The arguments against were serious. Wikipedia's credibility depended on neutrality. Taking sides—even popular ones—could compromise the trust that made Wikipedia valuable. The Neutral Point of View policy existed precisely to keep Wikipedia out of political disputes.

The arguments for were equally serious. SOPA directly threatened Wikipedia's operations. Sites could be blocked based on accusations, without due process. Wikipedia hosted millions of images and text excerpts; accusations of infringement were inevitable. This was not abstract politics. This was survival.

From December 2011 to January 2012, the English Wikipedia held an unprecedented request for comment. Over 1,800 editors participated. The question: should Wikipedia black out in protest?

The result: roughly 90% supported the blackout.

## Blackout

Sue Gardner acknowledged the tension in a pre-blackout statement:

"In making this decision, Wikipedians will be criticized for seeming to abandon neutrality to take a political position. That's a real, legitimate issue. We want people to trust Wikipedia, not worry that it is trying to propagandize them."

But she defended the action as self-defense: "We're not against protecting copyright, of course. We're against badly written legislation that will cause collateral damage to the open internet."

At 05:00 UTC on January 18, the switch flipped. The English Wikipedia displayed only the protest page. Google, Reddit, and thousands of other sites joined with their own protests. The Electronic Frontier Foundation reported over a million emails to Congress.

Twenty-four hours later, the blackout ended. Wikipedia came back. SOPA and PIPA did not.

"The Wikipedia blackout is over and the public has spoken," Gardner announced. "162 million of you saw our blackout page asking if you could imagine a world without free knowledge. You said no."

## Shadow of the Harvest

This was the zenith of Wikipedia's political power. A volunteer community had mobilized viewers numbering in the hundreds of millions and defeated legislation backed by major entertainment corporations.

This was also the turn.

Not the turn toward failure—Wikipedia would continue to thrive as an encyclopedia. The turn toward a different kind of institution. The steward that had proven it could defend the commons would soon discover that defense requires choosing enemies. And choosing enemies becomes difficult when enemies become customers.

2012 was when the harvesting began.

In May 2012, Google launched Knowledge Graph. The product displayed Wikipedia content directly in search results—information boxes that answered questions without requiring users to click through to Wikipedia itself. Wikipedia's content was being extracted and displayed where Google could monetize attention.

The same year Wikipedia demonstrated it could defend itself against legislative threats, a new threat was born. This one wouldn't come through Congress. It would come through APIs and training data and AI systems that would use Wikipedia without sending traffic back.

Wikipedia won SOPA. The victory was real. But the harvesting model that would hollow out Wikipedia's value proposition had already started. The collective strength shown in January 2012 would never be deployed against this new threat.

### Foundation View

From the Foundation's perspective, SOPA was exactly the kind of threat Wikipedia should mobilize against.

The legislation directly threatened Wikipedia's operations. Self-defense is not political advocacy. The community decided democratically—1,800 editors participated; 90% approved. The outcome was successful; the legislation died. This was appropriate use of Wikipedia's reach.

"Wikipedia is like the National Parks Service," Gardner explained. "The Internet is a vast space and it will only continue to grow, but in the vastness you still need space for parks or public libraries."

Protecting the public library was legitimate. The community agreed. The world agreed.

### Neutrality Critique

A minority of editors warned about precedent.

Wikipedia is a reference work, not an advocacy organization. Taking political positions—even popular ones—compromises the trust that makes Wikipedia valuable. If Wikipedia could black out over SOPA, what else might it black out over? Who decides which causes are legitimate?

"Wikipedia's power comes from being trusted," one dissenting editor wrote. "Taking political sides destroys that trust."

The concern was not that SOPA was good legislation. The concern was that Wikipedia had demonstrated a capacity for political action that could be misused. The blackout worked because the cause was popular. What happens when the cause is controversial?

### Unrealized Potential

A decade later, critics would ask a different question: if Wikipedia could black out to defeat SOPA, why couldn't it black out to force AI companies to license content properly?

By 2024, AI companies were training on Wikipedia at industrial scale. OpenAI's systems relied on Wikipedia as their primary factual source, with no attribution and no compensation. Wikipedia's traffic was falling—revealed only after the Foundation improved bot detection to separate real visitors from scrapers. The harvesting that SOPA threatened through legislation was happening through technology.

That collective power was never used against the scrapers. The Foundation responded to AI ingestion not with blackouts but with Enterprise contracts—selling access to the very companies whose compliance with Wikipedia's license was questionable.

The capacity for collective action existed. It was used once. It has never been used again.

The question of why the institution capable of the SOPA blackout never deployed similar power against AI extraction would require understanding who staffed that institution—and whose worldview they shared.

### Synthesis

The SOPA blackout proved what Wikipedia could do. The failure to act against AI harvesting proved what Wikipedia would not do.

The difference was not capability but will. And will, in institutions, follows incentives. By 2024, the Foundation derived revenue from AI companies through Enterprise. You don't threaten those who sign your checks.

The entanglement documented in later chapters began here, in the gap between demonstrated power and exercised power. Wikipedia proved it could mobilize the masses. Then it stopped.

### Victory's Meaning

SOPA and PIPA were effectively killed. Wikipedia was proven as a political force. Tech platforms demonstrated they could defeat traditional lobbying when they acted together.

The blackout also validated the fundraising model. Donors had supported not just servers but an institution capable of defending the open internet. The subsequent fundraisers would emphasize this broader mission—Wikipedia as champion of free knowledge, not just encyclopedia.

### What Came After

The black page faded at 05:00 UTC on January 19, 2012. Wikipedia returned, white and welcoming. The community celebrated a historic victory.

But the fundraising machine kept running. Revenue kept growing. By 2024 the Foundation had accumulated nine-figure annual revenue and nearly half a billion dollars in reserves and endowment combined.

And the community that proved it could black out the internet never blacked out again. Not when Google's Knowledge Panels absorbed content without click-through. Not when voice assistants read Wikipedia answers without attribution. Not when AI companies trained on the entire corpus and built products worth billions.

That collective power sits unused. Not because it disappeared, but because the institution that could wield it became dependent on the revenue it generates.

### Bridge Forward

Chapter 11 shows how success became professionalization, how professionalization created dependencies, and how dependencies prevented action. The Foundation built to defend Wikipedia grew into something that could no longer afford to fight.

The peak was January 18, 2012. What follows is the plateau and the fall.

CHAPTER ELEVEN

# The Apparatus

> Perhaps it's time to stop calling self-selected surveys of a tiny subset of our user base "community consensus."
>
> — Brion Vibber, WMF Lead Software Architect, 2014

On January 31, 2008, the Wikimedia Foundation closed its Florida headquarters.

A handful of employees packed their belongings and headed west. St. Petersburg—small-town, nonprofit, far from the tech industry—was being left behind. San Francisco—ambitious, expensive, at the heart of Silicon Valley's network effects—was the destination.

The stated reasons were practical: talent pool, proximity to partners, cheaper international travel. The unstated aspiration was evident: Wikipedia was joining the tech industry.

San Francisco was not just expensive and ambitious. It was also, by every measure, the most politically homogeneous major city in America—a Democratic bastion where professional networks sorted by ideology as surely as by skill. Hiring from these networks would shape who staffed the institution.

What gets left behind in a move is not just geography. It is identity. The scrappy volunteer project that ran on donated servers and volunteer

labor was becoming something else: an institution with office space and employee benefits and strategic planning documents.

The move to San Francisco was the visible symbol of a transformation that would take fifteen years to complete. When it finished, the Foundation would employ 700 people across 50 countries, manage a $185 million annual budget, and negotiate enterprise contracts with Amazon, Google, Meta, and Microsoft.

The question the move planted: who does the institution serve when the institution grows larger than its founders imagined?

### Gardner Achievement

Sue Gardner arrived in December 2007 to transform the Foundation.

Her background was in media, not encyclopedias. She had made CBC.ca Canada's most popular news site. She understood growth, professionalization, scale. The Foundation she inherited had five employees, crashed servers, and an annual budget of $2.7 million. The Foundation she left in May 2014 had 200 employees, reliable infrastructure, and a $60 million budget.

Forbes ranked her among the 100 most powerful women in the world. The Foundation became the fastest-growing nonprofit in America by revenue growth. These were real accomplishments.

Gardner understood what motivated volunteers: "Wikipedians do it for love, for mission-type reasons. They don't want to be paid. They want to be praised." The fundraising machine she built honored that insight. Personal appeals from Jimmy Wales. Messaging that emphasized mission, not charity. A million donors by 2011—democratic participation funding collective achievement.

The SOPA blackout happened under her leadership. The community voted; the Foundation coordinated; 162 million people saw the protest. This was the institution Gardner built at its peak: responsive to community will, capable of collective action, defending the open internet.

### Revenue Trajectory

The fundraising success was remarkable, as shown in Table 11.1.

| Year | Revenue | Staff | Revenue per Staff |
|---|---|---|---|
| 2007 | $2.7M | 5 | $540K |
| 2011 | $23M | 50 | $460K |
| 2014 | $60M | 200 | $300K |
| 2020 | $129M | 450 | $287K |
| 2024 | $185M | 700 | $264K |

**Table 11.1.** Wikimedia Foundation revenue trajectory, 2007–2024

Revenue grew 68-fold. Staff grew even faster—from five to seven hundred. The organization became more labor-intensive even as its primary product—volunteer-created content—required zero Foundation employees to produce.

Who were these employees? Where did they come from? What networks fed the hiring pipeline? These questions would matter enormously when the Foundation faced choices about whom to challenge and whom to partner with.

The growth was not irrational. A top-ten global website needs infrastructure, legal defense, community coordination, strategic planning. Someone has to do this work. The question is how much—and who decides.

### VisualEditor (2013)

In the early 2010s, the Foundation confronted a crisis: Wikipedia's editor population was declining. Between 2006 and 2011, the English Wikipedia lost roughly 15,000 active editors. Studies suggested the wiki markup syntax—the technical language editors used to format articles—was deterring new contributors.

The solution: a visual editor. Instead of learning markup, editors could click buttons to format text, like a word processor. The Foundation received a $3.6 million grant from the Stanton Foundation to build it.

On July 1, 2013, VisualEditor was deployed as the default editing interface—with over 200 known bugs.

The community objected. An RFC (request for comment) showed 212-20 consensus to revert the change—make VisualEditor opt-in rather than default. The Foundation refused.

"I don't think you can just add up the numbers... and call that a consensus," said Foundation staff—echoing attitudes that would persist across executive tenures.

On September 23, 2013, an administrator named Kww deployed code to override the Foundation's decision. Within minutes, the Foundation capitulated. VisualEditor became opt-in.

The pattern was established: WMF develops software without sufficient community consultation, deploys despite objections, eventually backs down after community resistance. But not before significant trust erodes.

### Superprotect (2014)

The German Wikipedia held a vote: should Media Viewer, a new image display feature, be disabled by default? The community said yes.

On August 10, 2014, the Foundation responded with "Superprotect"—a new technical permission that allowed Foundation staff to override community decisions. Erik Moeller, Deputy Director, was blocked on German Wikipedia for edit warring.

This was not a software dispute. It was a governance crisis. The Foundation had created a technical mechanism to override community consensus. 824 Wikimedians signed a protest letter. A vote went 664-103 to remove Superprotect.

Superprotect remained in place for fifteen months. It was finally removed on November 5, 2015, after Lila Tretikov acknowledged: "Superprotect set up a precedent of mistrust, and this is something it was really important for us to remove."

### Knowledge Engine (2016)

In January 2016, details emerged of a secret Foundation project called Knowledge Engine: "the internet's first transparent search engine," funded by a $250,000 Knight Foundation grant.

The community had not been consulted. The Board had not been fully informed. When trustee James Heilman pushed for transparency, he was removed from the Board.

The staff revolted. Luis Villa, Anna Koval, and Siko Bouterse resigned in protest. Lila Tretikov resigned as Executive Director on February 25, 2016.

The transparency paradox: an organization built on openness had developed a major strategic initiative in secret.

### FRAM (2019)

On June 10, 2019, the Foundation's Trust & Safety team banned administrator Fram for one year. No consultation with the community. No involvement of the Arbitration Committee—the body the community had created to handle such disputes.

Jimmy Wales called it "a question about our constitutional order."

Twenty-one administrators resigned in protest. The discussion generated 470,000 words. "I am not willing to serve Jan and the T&S team... under undisclosed new rules and under threat of unappealable sanctions," wrote one resigning administrator.

In September 2019, ArbCom vacated the ban. The active administrator count dropped below 500 for the first time.

### Pattern

The four crises share a structure: WMF pursues a strategic objective, the community objects through established processes, WMF dismisses objections as unrepresentative, the community finds ways to resist, WMF eventually backs down or modifies, and trust further erodes.

This was mission drift in miniature. An organization built to serve editors developed its own priorities—software initiatives, strategic plans, survival—that diverged from community values. When those priorities conflicted, staff had the resources (money, legal authority) and the community had only voice.

Each executive departed. The pattern continued. Gardner left in 2014. Tretikov resigned in 2016. Katherine Maher resigned in 2021. The structural tension between professional staff and volunteer community persisted regardless of leadership.

### Diverging Trajectories

The numbers tell a damning story, as shown in Table 11.2.

| Metric | 2007 | 2024 |
|---|---|---|
| WMF Staff | 5 | 700 |
| Active editors (English) | 51,000 | 30,000 |
| Very active editors | 8,000 | 3,500 |
| Active administrators | 1,000 | 449 |

**Table 11.2.** Foundation and community: diverging trajectories

Staff multiplied. Active editors declined 41%. Very active editors declined 56%. Active administrators declined 55%.

The Foundation grew while the community shrank. The institution that volunteers built now pursues agendas volunteers never authorized.

### Sustainability View

From the Foundation's perspective, professionalization was necessary and successful.

Wikipedia cannot be sustained by volunteers alone. Professional staff provide technical infrastructure at global scale, legal defense and advocacy, quality initiatives across languages, sustainability planning including the endowment, and global reach beyond the English-speaking world.

The evidence: Wikipedia still operates. SOPA was defeated. Turkey unblocked access. Mobile reading improved. The endowment reached $100 million five years ahead of schedule.

Without professionalization, Wikipedia might not exist.

### Capture View

From the community's perspective, the Foundation drifted from its mission.

The Foundation had been established to serve editors. Instead, it evolved into an entity that viewed community consensus as optional input. Software projects built without consultation. Strategic directions the community never authorized. "Reader interests" used to override editor decisions.

"The Wikimedia Foundation has a cancer," wrote community member Guy Macon in 2017. "The cancer is the explosive growth in spending, a growth rate that far exceeds the growth in content or the growth in users."

The original mandate was keeping Wikipedia online. But the organization became something with its own agenda, its own priorities, its own definition of success.

### Structural View

From a governance perspective, the tension is inherent.

Professional organizations develop professional priorities. Volunteer communities maintain different values. Neither is wrong; they are incom-

patible. Staff need to plan strategically, hire colleagues, manage budgets. Editors want to write an encyclopedia without bureaucratic interference.

"Foundation employees are members of the community," the Foundation claimed after VisualEditor. But employees follow professional incentives. Volunteers follow mission incentives. When these diverge, whose voice prevails?

The answer is structural: the Foundation has legal control, money, and staff. The community has content and labor. In conflicts between those holding resources and those providing labor, resources typically win.

### False Victory

By 2024, the professionalized Foundation was complete: staff spanning fifty countries, nearly $200 million in annual revenue, an endowment that had reached its targets five years early, remote-first operations encircling the globe.

This was a genuine achievement. The shoestring operation that crashed regularly in 2007 had become a sustainable institution. Wikipedia would not disappear because a server failed or a key volunteer burned out.

But the achievement was also a transformation. The volunteer project had acquired a professional apparatus. The apparatus developed its own agenda. That agenda pulled away from the community that created the content.

### Capture Mechanism

The professionalization that looked like victory created the conditions for dependency.

A 700-person institution requires approximately $107 million in annual salaries. When donations plateau, new revenue streams become necessary. Enterprise—commercial API access—becomes not just possible but required.

The institution needed to survive. Survival required growth. Growth required revenue. Revenue required Enterprise. Enterprise required serving AI companies. Serving AI companies required not enforcing licenses against paying customers.

The steward's paradox: an organization meant to protect a shared resource from commercial exploitation grew dependent on commercial relationships to survive. The founder control embedded in 2003 ensured no democratic check on this choice. The editors who created the content had no authority to override it.

This was not failure but logical consequence. The sprawling organization that would sell API access to Google and Amazon did not emerge from nowhere. It was built deliberately, funded by small donors who believed they were "keeping Wikipedia online."

### Bridge Forward

Chapter 12 shows what professionalization produced: Wikimedia Enterprise, the commercial arm that sells access to AI companies. The free-rider problem met a business model. The mechanism of dependency was installed.

The move to San Francisco was an aspiration. The aspiration became an organization. That organization would face a choice: defend the copyleft licenses that governed volunteer-created content, or take money from companies whose compliance was questionable.

The institution chose revenue.

CHAPTER TWELVE

# Enterprise

> Our content is there to be used. It's freely-licensed for a reason. At the same time, it's like the environment. It's there to be used, but it's not there to be exploited.
>
> — Lisa Gruwell, WMF Chief Advancement Officer, 2018

On January 15, 2026, Wikipedia celebrated its twenty-fifth birthday.

The anniversary announcement included news: Amazon, Meta, Microsoft, Mistral AI, and Perplexity had signed enterprise contracts with the Wikimedia Foundation. The trillion-dollar companies that had used Wikipedia for years—extracting content for voice assistants, knowledge panels, AI training—would finally pay their "fair share."

Jimmy Wales told Bloomberg that AI bots had been "absolutely hammering" Wikipedia's servers. "They've been scraping Wikipedia at massive scale," he said. "Now they'll pay for proper access."

The press releases spoke of partnership, sustainability, infrastructure costs. "Our content is free, our infrastructure is not," the Foundation explained. The free encyclopedia remained free. Only the premium API access cost money.

This was the pitch: companies that extracted value should contribute value. The free-rider problem, finally solved.

But one fact complicated the celebration. In April 2024—nine months earlier—the Foundation's Enterprise team had privately admitted that many AI customers were likely ignoring the license terms.

The Foundation knew its customers might be flouting the licenses governing the content being sold. It continued regardless.

How did a copyleft project founded on "free as in freedom" end up selling premium access to companies that may be violating that freedom—and calling it partnership?

## The Steward's Dilemma

The Foundation faced a genuine problem. For years, tech giants extracted enormous value from Wikipedia while contributing almost nothing. This violated the basic principle of the commons: those who benefit should help maintain what they benefit from. But solving the free-rider problem would require the steward to become a commercial actor—exactly what it had been created to prevent.

## Extraction Without Contribution

The pattern was visible everywhere.

In 2012, Google launched Knowledge Graph. The product displayed Wikipedia content directly in search results—information boxes that answered questions without requiring users to click through to Wikipedia. Users got their answers. Google got advertising revenue. Wikipedia got nothing.

Amazon's Alexa launched in 2014, reading Wikipedia content aloud in response to millions of daily queries. When you asked "Who was Albert Einstein?", Alexa recited Wikipedia's article. No attribution. No compensation. No link for users to follow.

Apple's Siri worked similarly. Voice assistants across the industry treated Wikipedia as a free backend, extracting content to power products worth billions.

By 2018, TechCrunch investigated corporate contributions to Wikipedia. Amazon had contributed $0. Apple contributed approximately $50,000 through employee matching only. Google contributed roughly $1 million—approximately 0.001% of revenue.

"Entire teams at Amazon, Google and similar companies" were sorting through Wikipedia data. The trillion-dollar companies built products on volunteer labor and contributed negligibly to infrastructure.

Gruwell's environmental metaphor—the chapter's epigraph—was apt. And like the environment, exploitation was happening without consequence.

The free-rider problem was not limited to corporate extraction. By the 2010s, Wikipedia had become a target for sophisticated reputation management. Leaked emails later revealed that convicted financier Jeffrey Epstein's team spent tens of thousands of dollars on SEO and Wikipedia manipulation between 2010 and 2019. His publicist described Wikipedia as "a tough nut to crack" controlled by "copyediting geeks"—but Epstein was "partially successful for a decade" before volunteer editors finally prevailed. The lesson was already clear: Wikipedia's openness made it vulnerable. Anyone with sufficient resources could shape what the world believed to be true.

## Disintermediation

The free-riding was not just financial. It was structural.

When users receive Wikipedia content through Google Knowledge Panels or voice assistants, they never visit Wikipedia. They don't see the "edit" button. They don't see the fundraising banners. They don't encounter the community.

The volunteer feedback loop was being broken. Wikipedia's model depended on readers becoming editors. If readers never reached Wikipedia, where would the next generation of editors come from?

Traffic data confirmed the concern. Knowledge Panels, AI Overviews, and voice assistants reduced click-through to source sites. Wikipedia was being used to train systems that competed with Wikipedia for user attention.

### Quiet Incorporation

In January 2020, Wikimedia, LLC was quietly incorporated in Delaware—a for-profit subsidiary to handle commercial operations.

Delaware was chosen for legal clarity: "most established and most well understood." The LLC structure insulated the Foundation from commercial liabilities while enabling revenue-generating activities.

Lane Becker, a serial entrepreneur with exits at Google and Code for America, joined to build the product. The operating agreement was executed in March 2021.

### Product

Enterprise offered three API tiers: On-Demand for single articles retrieved as needed, Snapshot for full Wikipedia dumps with periodic updates, and Realtime for streaming updates as edits occur.

The Structured Contents Initiative transformed Wikipedia's wikitext into machine-readable JSON. Parsed infoboxes, tables, and references with quality scoring. The content remained free; the formatting and infrastructure were commercial.

What Enterprise sold was not content—anyone could download Wikipedia dumps for free. It sold bandwidth, speed, format, and SLA guarantees. Premium access for premium customers.

### Launch and Growth

In March 2021, Wired broke the story: "Wikipedia Is Finally Asking Big Tech to Pay Up."

On October 25, 2021, Enterprise formally launched from San Francisco. In June 2022, Google and Internet Archive were announced as the first customers. Tim Palmer from Google called Wikipedia "a unique and valuable resource."

Revenue grew modestly at first, as shown in Table 12.1.

| Fiscal Year | Revenue | Expenses | Net |
|---|---|---|---|
| 2022-23 | $3.2M | $4.0M | -$0.8M |
| 2023-24 | $3.4M | $3.8M | -$0.4M |
| 2024-25 | $8.3M | $4.4M | +$3.9M |

**Table 12.1.** Wikimedia Enterprise revenue, 2022–2025

Then came the AI explosion.

## AI Pivot

By 2025, AI companies had become the primary growth market. Wikipedia reported that 65% of its most resource-intensive traffic came from bots—many evading detection by masquerading as humans.

The bot traffic Wales described in the anniversary announcement was not new—after the Foundation improved detection, Wikipedia reported an 8% decline in "human" traffic. That traffic had never been human in the first place.

In January 2026, the anniversary announcement: Amazon, Meta, Microsoft, Mistral AI, Perplexity, and others as Enterprise customers. Revenue hit $8.3 million, with 148% year-over-year growth. The service became profitable.

The contribution problem appeared solved. Companies that had taken now gave.

### April 2024 Admission: The Capture Made Explicit

But one detail complicated the success story. One moment made the entanglement undeniable.

In April 2024, 273 Ventures—a company building KL3M, an AI training dataset designed for legal compliance—reached out to the Wikimedia Foundation. They wanted clarification: was their approach to Wikipedia content licensing correct?

The Foundation's Enterprise team responded with startling candor:

"The Foundation was watching what AI companies did with Wikipedia content—and suspected that compliance with Creative Commons terms was, at best, inconsistent across the industry."

The admission continued: 273 Ventures was "the only people who had asked."

Let that sink in. Only one company had asked if they were complying. That meant the Foundation knew, or strongly suspected, that dozens of other companies were ignoring the licenses. It knew its own customers might be flouting the terms the volunteers had chosen.

And the partnerships continued.

The steward had been compromised. Not through corruption but through dependency. The institution created to protect the shared resource was now profiting from its harvesting.

### Pragmatic Defense

From the Foundation's perspective, Enterprise was pragmatic problem-solving.

Companies were already scraping Wikipedia. Better to formalize the relationship and get paid. Individual access remains free. Infrastructure costs money. Donors shouldn't subsidize trillion-dollar companies.

"Our donors are not donating in order to subsidize these huge AI companies," Wales said.

The 30% revenue cap, formalized in May 2025, preserved nonprofit character. At least 70% of revenue must come from donations and grants. Enterprise at 4% of total revenue was well below danger levels.

And the companies were now contributing. That was progress, however imperfect.

### Exploitation Critique

From some community members' perspective, Enterprise was volunteer exploitation.

A quarter million volunteers created the content. Enterprise sells access to that labor. Volunteers have no share of revenue and limited say in commercial decisions. The Foundation didn't create the value but profits from its distribution.

"These are not charitable contributions—they are commercial transactions," one Meta-Wiki contributor wrote. "The WMF is selling Wikipedia data, generated by its volunteer editors."

The Foundation's response: content remains free; infrastructure costs money. But the distinction felt increasingly thin when "infrastructure" meant formatted access to content that volunteers created.

### Capture Concern

From a governance perspective, Enterprise created a structural conflict of interest.

Once the Foundation depends on revenue from Google, Amazon, Meta, and Microsoft, enforcement becomes awkward. You don't sue your paying customers—especially when you share their worldview, their professional networks, and their vocabulary of "democratizing knowledge."

The April 2024 admission made this explicit: the Foundation knew customers might be disregarding licenses. It continued selling. The commercial relationship created a conflict of interest at the heart of the institution meant to protect copyleft.

The 30% cap acknowledged the danger. But financial dependency operates even below the cap. At 4% of revenue, Enterprise is $8.3 million annually. That pays for staff, infrastructure, strategic initiatives. Losing it would hurt.

How much non-compliance would the Foundation tolerate to preserve that revenue?

### Synthesis

The three perspectives are not mutually exclusive.

Enterprise solved a real problem: tech giants were extracting value without contributing. The solution created a new problem: the Foundation became financially entangled with companies whose compliance was questionable.

This is the tragedy of entanglement: the solution to one problem becomes the source of another.

### What Enterprise Created

By January 2026, the dependency was fully operational.

Tech giants now paid for access—a genuine change from the pre-2021 status quo. But those payments created relationships the Foundation depended on. Enforcement against Enterprise customers would threaten revenue.

The Foundation knew customers might be ignoring CC-BY-SA. It sold to them anyway. The copyleft promise—that derivative works would remain free—met a commercial imperative that made legal action "economically awkward."

### Prophecy Fulfilled

In 2002, Edgar Enyedy and the Spanish Wikipedians forked the project over fears that advertising would commercialize free knowledge. They were right about the threat. They were wrong about the mechanism.

Commercialization arrived not through banner ads but through enterprise contracts. Not through readers seeing advertisements but through AI companies training on volunteer labor. Not through Bomis taking value but through the Foundation selling access to companies that may be breaching the licenses that made Wikipedia possible.

Enyedy wrote "Good luck with your WikiPAIDia" twenty-four years before Wikipedia became, in a sense, exactly that. The revenue model he feared had arrived—just not the way anyone imagined.

The tragedy is not that commercial interests came for Wikipedia. That was always the danger. The tragedy is that the institution designed to protect against that danger became the primary vehicle for it.

### Bridge Forward

Part IV examines what Enterprise enabled: AI extraction at industrial scale.

Chapter 13 shows how much Wikipedia constitutes AI training data. Chapter 14 shows Enterprise expanding its AI customer base. Chapter 15 documents the full pattern—the April 2024 admission, the enforcement gap, the structural conflict of interest.

Enterprise is not the problem itself. Enterprise is the structure that enabled it. The customers became leverage. The revenue became a reason not to sue.

### End of The Institution's Turn

The Institution's Turn has traced the Foundation's journey from steward to captor.

Chapter 9 showed the governance structure planted in 2003: appointed leadership, limited community voice, founder control. The steward was designed without a mechanism to remain accountable to those it was meant to serve.

Chapter 10 showed the peak of collective action: SOPA defeated, 162 million people mobilized. This was the moment the steward proved capable of defending the commons. It would never act again.

Chapter 11 showed professionalization creating dependency. The steward needed to grow to survive. Growth required revenue. Revenue required Enterprise.

This chapter showed how Enterprise became the vehicle of financial entanglement. The steward now depended on those it was meant to restrain.

The Foundation was created to protect the commons from commercial co-optation. By 2026, the Foundation had become the conduit for that co-optation—knowingly selling access to companies disregarding the licenses that made Wikipedia possible.

Part IV opens on the extractors themselves.

# Part IV

# The AI Reckoning

CHAPTER THIRTEEN

# The Extraction Begins

> It would be impossible to train today's leading AI models without using copyrighted materials.
> — OpenAI submission to UK House of Lords, January 2024

On May 28, 2020, a research paper appeared on arXiv.org that would reshape the infrastructure of knowledge.

Tom Brown and thirty co-authors at OpenAI published "Language Models are Few-Shot Learners"—the technical paper introducing GPT-3. Buried in seventy-five pages of architecture diagrams and benchmark results was Table 2.2: a breakdown of what GPT-3 had consumed.

The table listed five data sources. Common Crawl, scraped from the open web, provided 410 billion tokens—roughly 60% of the training data. WebText2, a curated collection of outbound links from Reddit, contributed 19 billion. Two collections of digitized books added another 67 billion.

Then came the final line: English Wikipedia. Three billion tokens. Three percent of the total.

But the column that mattered was not the one measuring quantity. It was the one measuring quality. OpenAI had assigned each data source an "epoch" weight—a multiplier determining how many times the model would see each piece of training data. Common Crawl, the bulk of the corpus, received 0.44 epochs. The model would see less than half of it.

Wikipedia received 3.4 epochs. OpenAI's researchers reviewed Wikipedia more than three times. The open web got less than half a pass.

The implication was clear. OpenAI considered Wikipedia the most trustworthy source in its training data—by a factor of nearly eight. The encyclopedia that a quarter million volunteers had spent two decades building had become the quality benchmark against which the internet was measured.

Table 2.2 appeared in a paper read mostly by researchers. No press release celebrated its implications. The volunteers who had created this benchmark were not consulted, compensated, or credited. Their names appeared nowhere in the seventy-five pages.

The harvesting had begun. The enclosure of the digital knowledge base—what legal scholar James Boyle had called the "second enclosure movement"—had found its most valuable target.

## Scale of Extraction

GPT-3 was not an outlier. Every major language model that followed would rely on Wikipedia as a quality anchor.

In February 2023, Meta released LLaMA, demonstrating that state-of-the-art performance was possible with publicly available data. The training data breakdown showed Wikipedia at 4.5%—higher even than GPT-3's three percent. The model that proved you could build a frontier AI without proprietary data depended on an encyclopedia built by volunteers.

After LLaMA, transparency evaporated. OpenAI stopped disclosing training data composition. Anthropic never started. Google's Gemini documentation mentioned training data only in the vaguest terms. The industry had learned that disclosure created liability.

But absence of disclosure did not mean absence of Wikipedia. In January 2024, OpenAI confirmed what the chapter epigraph states—the

| Model | Year | Wikipedia Share |
|---|---|---|
| GPT-3 | 2020 | 3% (3.4 epochs) |
| LLaMA | 2023 | 4.5% |
| GPT-4 | 2023 | Undisclosed |
| Claude | 2023-25 | Undisclosed |
| Gemini | 2024-25 | Undisclosed |

**Table 13.1.** Wikipedia in AI training data

quiet part said loud. Wikipedia was copyrighted material, licensed under CC-BY-SA. The admission applied directly to it.

The Common Crawl shadow made the extraction even larger than disclosed percentages suggested. Common Crawl scraped the open web, and the open web included Wikipedia mirrors. Mozilla's 2024 analysis found that 80% of GPT-3's tokens came from Common Crawl, and 64% of forty-seven major language models used Common Crawl in their training. Wikipedia appeared both directly in training data and indirectly via Common Crawl overlap. The true Wikipedia presence in AI systems exceeded anything Table 2.2 acknowledged.

### Citation Dominance

The ingestion showed in what the models produced.

In August 2025, Profound published an analysis of 680 million AI citations. The findings confirmed what the training data suggested: Wikipedia dominated AI knowledge production.

ChatGPT cited Wikipedia in nearly half of its top-ten most-cited sources. For the world's most popular AI assistant, Wikipedia was not one source among many. It was the foundation of factual knowledge.

But citation did not mean attribution. Pew Research found that only one percent of users clicked AI citations. Users received Wikipedia-derived

| Platform | Top Source | Share of Top-10 Citations |
|---|---|---|
| ChatGPT | Wikipedia | 47.9% |
| Google AI Overviews | Reddit | 21.0% |
| Perplexity | Reddit | 46.7% |

**Table 13.2.** AI citation patterns by platform

answers without knowing the origin. The knowledge transferred; the credit did not.

### Memorization Evidence

The conventional defense of AI training held that models "learned" from data rather than copying it. Training was transformation: raw text became statistical patterns, patterns became capabilities, capabilities generated new text. The original sources were digested, not reproduced.

Nicholas Carlini and colleagues at Google DeepMind proved this narrative false.

In a series of papers from 2021 to 2023, Carlini's team documented what they called "memorization"—the tendency of language models to reproduce training data verbatim when prompted correctly.

The findings were systematic:

**2021:** "Extracting Training Data from Large Language Models" demonstrated that training data could be recovered from production models. Larger models were more vulnerable, not less. The capability that made models useful also made them leaky.

In 2023, "Quantifying Memorization Across Neural Language Models" identified three laws: memorization increases logarithmically with model size, examples seen multiple times are memorized more, and longer prompts extract more memorized content.

**2023:** "Scalable Extraction of Training Data" demonstrated "divergence attacks" that could extract training data at 150 times the normal

rate. Gigabytes were extractable from ChatGPT. This was not theoretical vulnerability. It was practical extraction from production systems.

Wikipedia was prime memorization territory. High quality meant direct inclusion in training data. Multiple sources meant appearance both in dumps and Common Crawl. Repeated epochs meant GPT-3 saw Wikipedia more than three times over. By Carlini's second law—duplicated data is memorized more—Wikipedia's quality made it a memorization magnet.

A 2024 study found that one to seven percent of language model outputs were exact duplicates from training data. Researchers reconstructed near-perfect chapters of copyrighted books from model outputs. ChatGPT, tested against poetry, retrieved seventy-two of 240 poems—often verbatim.

This was not transformation. It was reproduction at industrial scale.

### License Limits

The implications for CC-BY-SA were devastating.

The license required two things: attribution and share-alike. Attribution meant crediting Wikipedia when using its content. Share-alike meant licensing derivative works under the same terms.

Both requirements collapsed in the AI context.

**Attribution:** ChatGPT cited Wikipedia in 47.9% of its top sources. But users rarely saw those citations. One percent clicked. The attribution existed in the interface; it failed in practice. Most users received Wikipedia-derived information without knowing where it came from.

**Share-Alike:** Did a trained AI model count as a "derivative work"? Courts trended toward no. In *Kadry v. Meta* (2023), the argument that model weights inherited copyleft obligations failed. If model weights were not derivatives, share-alike could not spread. The copyleft mechanism stopped at the training dataset.

Creative Commons acknowledged this in May 2025. Their legal primer on AI training stated directly: "Share Alike/CopyLeft licenses are largely

ineffective when materials licensed under them are used to train AI models."

The Institute for Information Law (IVIR) identified three fundamental challenges: tracing uses during training, interplay with fair use exceptions, and identifying protected content in outputs. The license was designed for a world where derivatives were recognizable—where a modified Wikipedia article was still an article, subject to inspection and enforcement. Model weights were not articles. They were billions of floating-point numbers whose relationship to training data was mathematically obscure.

Stallman had created copyleft so that anyone using free software would have to share their improvements. Wikipedia adopted copyleft so that anyone using Wikipedia content would keep their derivatives free. AI companies used Wikipedia's content to build products worth hundreds of billions of dollars—and the license that was supposed to protect volunteer labor had no mechanism to constrain them.

### Flattered View

From Jimmy Wales's perspective, AI training validated Wikipedia's achievement.

"I'm very happy personally that AI models are training on Wikipedia data because it's human curated," Wales said in 2024. "I wouldn't really want to use an AI that's trained only on X, you know, like a very angry AI."

The argument was straightforward: Wikipedia's mission was to spread free knowledge. AI trained on Wikipedia spread knowledge faster and wider than Wikipedia alone ever could. Being foundational to AI proved Wikipedia's quality—the most sophisticated systems in the world chose Wikipedia as their trusted source.

From this view, the volunteers should be proud. They had not just created an encyclopedia. They had created the knowledge backbone of the AI era.

### Exploitation View

From the volunteers' perspective, validation looked different.

A quarter million editors had created 65 million articles. They contributed expecting their work would remain "free as in freedom"—open to all, but protected by licensing that would keep derivatives in the shared resource.

Then came the extraction, as shown in Table 13.3.

| Entity | Value Captured |
|---|---|
| OpenAI (revenue, 2025) | $6.7 billion |
| Anthropic (valuation, 2025) | $18 billion |
| Google (AI revenue, 2024) | $31 billion |
| Wikipedia volunteer compensation | $0 |

**Table 13.3.** Value captured from Wikipedia training

The companies that trained on Wikipedia built products worth tens of billions in annual revenue. The contributors who created the training data received nothing—not compensation, not attribution in practice, not even the share-alike protection that was supposed to ensure derivatives remained free.

In February 2025, the mask slipped. Noam Brown, an OpenAI researcher, tweeted that Deep Research "might be the beginning of the end for Wikipedia." When criticized, his response was candid: "Going down this route essentially means substituting capital for human labor (in this case, Wikipedia editors)."

The parenthetical was telling. AI companies did not view Wikipedia as a partner. They viewed it as labor to be automated away.

### Existential Threat

From an academic perspective, the data harvesting threatened Wikipedia's survival.

Wagner and Jiang, writing in the *Journal of the Association for Information Science and Technology* in 2025, predicted a "vicious cycle." If AI could answer questions without sending users to Wikipedia, why visit? If users did not visit, why edit? If editors stopped editing, content would stale. If content staled, AI outputs would degrade. The system that trained on Wikipedia's quality would undermine the conditions that produced it.

The evidence supported the concern. Traffic was already declining—down 8% in 2025 after the Foundation improved bot detection. The decline was attributed to revealing that "human" traffic had never been human in the first place. But organic search traffic was down 26% over three years. Zero-click searches—queries answered without any click-through—had increased from 56% to 69%.

Wikipedia was being used to train systems that competed with Wikipedia for user attention. The loop was beginning to close.

### Synthesis

These perspectives were not mutually exclusive.

Wikipedia's quality made it valuable for AI training—that was true. That training absorbed labor without compensation—that was also true. The scraping might ultimately threaten Wikipedia's sustainability—the evidence pointed that way.

The volunteers had created value. AI companies captured it. The future was uncertain. What was clear was that the copyleft license meant to protect their work could not function in this context.

### Forty-Year Circle

In 1985, Richard Stallman watched a company take his code and refuse to share improvements. He created copyleft to prevent that from ever happening again. Use my work, he said, but share your derivatives. The license would spread like a virus, ensuring that free software begat free software.

Forty years later, AI companies took Wikipedia's text and built proprietary systems. The license that was supposed to prevent this existed. But model weights were not articles. Training was not copying—or if it was copying, the copies were invisible, encoded in billions of parameters that no inspection could trace to their sources.

The specific mechanism differed from Symbolics taking Stallman's printer driver code. The taking was the same.

The paper appeared on May 28, 2020. Five years later, ChatGPT cited Wikipedia in nearly half its factual responses. The people who built that knowledge base—the editors, the arbitrators, the tireless fixers of vandalism—saw none of the benefits. No consent. No consultation. No visible credit.

The copyleft promise—that derivative works would share the freedom of the original—had broken against machine learning's opacity. Not because the license was poorly written. Because the enforcement mechanism depended on visibility, and AI training was invisible by design.

Chapter 14 shows how the steward responded: not with enforcement, but with Enterprise contracts. Sell access to the extractors. Take their money. Call it sustainability.

Chapter 15 reveals what that strategy required: knowing that customers were not complying with licenses, and selling to them anyway.

The promise Stallman made—use, but share—had become a one-way valve. The using continued. The sharing stopped at the model weights.

CHAPTER FOURTEEN

# Selling to the Machine

> Our donors aren't donating to subsidize Sam Altman, they're donating to help Wikipedia.
>
> — Jimmy Wales, January 2026

On January 15, 2026, the Wikimedia Foundation announced new partnerships to celebrate Wikipedia's twenty-fifth birthday.

The press release listed the names: Amazon. Meta. Microsoft. Mistral AI. Perplexity. These trillion-dollar companies would now pay for Enterprise access—structured API access to Wikipedia's content, delivered with service level agreements and premium formatting.

Jimmy Wales told Bloomberg that AI bots had been "absolutely hammering" Wikipedia's servers. Now they would pay their "fair share."

The announcement read like a victory. The free-rider problem, finally solved. The companies that harvested value would contribute value. Enterprise revenue had more than doubled year-over-year, reaching $8.3 million. The service had become profitable for the first time.

But one name was missing from the list.

OpenAI—whose ChatGPT cited Wikipedia more than any other source. OpenAI—whose GPT-3 had trained on Wikipedia at 3.4 epochs, the highest quality weighting of any data source. OpenAI—valued at $157 billion, with annual revenue exceeding $6.7 billion.

OpenAI paid Wikipedia nothing.

The company using Wikipedia most heavily—the company whose entire product depended on knowledge harvested from volunteer labor—had opted out of the "fair share" arrangement.

This was the dependency made visible. If Enterprise were about enforcing licensing compliance, OpenAI's absence would be unacceptable. That the Foundation tolerated it proved Enterprise was never about compliance. It was about revenue from willing payers. The biggest user's absence revealed the truth: the Foundation had built a voluntary contribution system, not a compliance mechanism.

If Enterprise was about fairness, why was the biggest user exempt?

## Building the Commercial Arm

Enterprise had been designed to solve a real problem.

In January 2020, the Wikimedia Foundation quietly incorporated Wikimedia, LLC—a for-profit subsidiary to handle commercial operations. Delaware was chosen for legal clarity. The operating agreement was executed in March 2021.

Lane Becker, a serial entrepreneur with exits at Google and Code for America, joined to build the product. His pitch was straightforward: tech giants had been extracting value from Wikipedia for years. Google's Knowledge Graph launched in 2012, displaying Wikipedia content directly in search results. Amazon's Alexa read Wikipedia articles aloud to answer queries. Apple's Siri did the same. Trillion-dollar products built on volunteer labor, contributing almost nothing in return.

In March 2021, *Wired* broke the story: "Wikipedia Is Finally Asking Big Tech to Pay Up." On October 25, 2021, Enterprise formally launched from San Francisco.

The product offered three tiers: On-Demand for single articles retrieved as needed, Snapshot for full Wikipedia dumps with periodic updates, and Realtime for streaming updates as edits occur.

The Structured Contents Initiative transformed Wikipedia's wikitext into machine-readable JSON. Parsed infoboxes, tables, references with quality scoring. What Enterprise sold was not content—anyone could download Wikipedia dumps for free. It sold bandwidth, speed, format, and reliability guarantees.

In June 2022, Google became the first paying customer. Tim Palmer from Google called Wikipedia "a unique and valuable resource." The Internet Archive received free access as a mission-aligned partner.

Revenue grew modestly at first. Then came ChatGPT.

### AI Pivot

ChatGPT launched on November 30, 2022. Within months, the AI landscape transformed. Language models went from research curiosity to mainstream product. And every one of them depended on Wikipedia.

Enterprise found its market.

"The AI companies are here, and they are particularly voracious," Lane Becker told IBM Think. "We're a general store of knowledge with extremely timely information, and a lot of these organizations are slamming our servers."

The Foundation reported that 65% of its most resource-intensive traffic came from bots—many evading detection by masquerading as humans. After improved bot detection, Wikipedia reported an 8% decline in "human" traffic. The traffic had never been human in the first place.

New customers signed: Ecosia, Nomic, Pleias, ProRata, Anthropic. Revenue jumped, as shown in Table 14.1.

Enterprise became profitable. The January 2026 announcement added Amazon, Meta, Microsoft, Mistral AI, and Perplexity to the customer roster. The free-rider problem appeared solved.

### Voluntary Paradox

The celebratory framing concealed a structural problem.

| Fiscal Year | Revenue | Expenses | Growth |
|---|---|---|---|
| 2022-23 | $3.2M | $4.0M | — |
| 2023-24 | $3.4M | $3.8M | +6% |
| 2024-25 | $8.3M | $4.4M | +148% |

**Table 14.1.** Enterprise revenue during the AI pivot

Enterprise was entirely voluntary. No legal requirement compelled AI companies to pay for API access. Anyone could scrape Wikipedia freely—the content was licensed for reuse. Enterprise offered convenience, not compliance.

This meant the companies that perceived value in the formal relationship paid. The companies that did not perceive value—including the heaviest user—continued scraping with no consequence.

The numbers made the paradox stark, as shown in Table 14.2.

| Metric | Value |
|---|---|
| Wikipedia share of ChatGPT top-10 citations | 47.9% |
| GPT-3 training: Wikipedia epochs | 3.4 (highest) |
| OpenAI revenue (2024-25) | $6.7 billion |
| OpenAI payment to Wikipedia | $0 |

**Table 14.2.** The OpenAI gap: extraction without payment

ChatGPT cited Wikipedia nearly half the time. OpenAI had trained on Wikipedia more intensively than any other source. The company's entire knowledge infrastructure depended on volunteer-created content.

OpenAI paid nothing.

Wales acknowledged the gap. OpenAI, he argued, should "chip in and pay" for the server costs its products imposed on the nonprofit. But "should" was not "must." The appeal was moral, not legal. OpenAI could ignore it—and did.

### What Fair Share Conceals

The OpenAI gap exposed what Enterprise actually was.

Not legal accountability. Not even soft pressure. Enterprise was a premium API service. Companies that wanted formal access paid. Companies that did not want it continued scraping without consequence.

The paying customers gained something valuable: implicit legitimacy. Google, Amazon, Meta, Microsoft—by paying, they signaled willingness to support the commons. Their Enterprise contracts created a presumption of good faith.

But the contracts did not require CC-BY-SA compliance. The Foundation had admitted in April 2024 that many customers were not following the licensing terms. It continued the partnerships regardless.

The non-paying customers faced no consequence. OpenAI could train on Wikipedia, rely on it more than any other source, and pay nothing. No lawsuit. No demand letter. No public shaming beyond Wales's occasional comments.

The result was a system with the structural effect of preventing enforcement against anyone.

If the Foundation sued non-paying companies like OpenAI, it would face accusations of selective enforcement—why OpenAI but not the paying customers who also might not be compliant? If it sued paying customers, it would destroy the revenue stream it depended on. The only stable equilibrium was universal non-enforcement, dressed up as pragmatic adaptation.

"Fair share" rhetoric masked this reality. The companies paying Enterprise were not paying their fair share of licensing obligations. They were paying for API access that anyone could get for free by scraping. The payment had nothing to do with compliance.

An adversarial organization would have built enforcement mechanisms before taking corporate money. The Foundation never did. Its leadership moved in overlapping circles with tech executives—shared con-

ferences, matching political donations (WMF employees gave exclusively to Democrats across 292 FEC contributor-cycles), a common vocabulary of "access" and "scale."

Enterprise formalized what already existed. The Foundation did not become captured; it entered relationships its leadership already considered natural. Financial dependency grew from cultural affinity, not the reverse.

### Foundation Position

From the Foundation's perspective, Enterprise was responsible stewardship.

AI companies were already scraping Wikipedia. Better to engage, collect revenue, and ensure sustainability than to fight a losing battle. The content remained free. The infrastructure cost money. Donors should not subsidize trillion-dollar corporations.

The 30% revenue cap, formalized in May 2025, protected the nonprofit character. At least 70% of revenue must come from donations and grants. Enterprise at 4% of total revenue was well below danger levels.

"Wikipedia is a critical component of these tech companies' work that they need to figure out how to support financially," Becker said. The argument was pragmatic: if they were going to use Wikipedia anyway, at least some should pay something.

### Volunteer Concern

From the volunteer perspective, Enterprise raised uncomfortable questions.

A quarter million volunteer editors created the content that Enterprise sold access to. Revenue went to the Foundation, not to contributors. No community vote authorized the commercial partnerships. The decisions were made by Foundation staff and approved by an appointed Board.

"These are not charitable contributions—they are commercial transactions," one Meta-Wiki contributor wrote. "The WMF is selling Wikipedia data, generated by its volunteer editors."

The Foundation's response—content remains free, infrastructure costs money—felt increasingly thin. The "infrastructure" being sold was formatted access to content volunteers created. The distinction between selling content and selling access to content was semantic.

### Licensing Question

From a licensing perspective, Enterprise created a structural conflict of interest.

CC-BY-SA required attribution and share-alike. The Foundation acknowledged that most AI customers might not be complying. But suing customers would destroy Enterprise revenue. Not suing meant accepting non-compliance as the price of commercial relationships.

Once the Foundation depended on revenue from Google, Amazon, Meta, and Microsoft, legal action against those companies became structurally implausible. The institution that was supposed to protect copyleft licensing had become financially intertwined with companies whose compliance was questionable.

The Foundation's own April 2024 acknowledgment of widespread non-compliance confirmed this. The continued Enterprise sales demonstrated what the Foundation had chosen.

### Synthesis

These perspectives did not resolve cleanly.

Enterprise might be pragmatic adaptation that also created co-optation dynamics. The Foundation might believe it served the mission while its interests aligned with the companies ingesting Wikipedia's content. The volunteers might be right about the labor-value disconnect even if commercial sustainability was necessary.

What mattered for the book's argument: once the Foundation depended on AI company revenue, legal action became structurally unlikely. The mechanism was in place. The entanglement was operational.

## Birthday Announcement

The January 2026 press release listed who paid. It did not list who used.

Google, Amazon, Meta, Microsoft, Mistral AI, Perplexity—these were the partners. These were the companies that had chosen to formalize their relationship with Wikipedia.

OpenAI was not on the list—the company using Wikipedia most heavily, citing it in 47.9% of responses, paying nothing. Neither was Apple, whose Siri had been reading Wikipedia articles since 2011. Neither was xAI, whose Grokipedia had forked Wikipedia with visible attribution three months earlier.

The gap between who paid and who used revealed the capture's logic. OpenAI—no payment, no attribution, proprietary outputs—faced no consequences. xAI—free access, visible attribution, CC-BY-SA on derived content—faced public condemnation for its editorial failures. The Foundation's criticism tracked ideology, not licensing compliance.

By January 2026, the Wikimedia Foundation was financially intertwined with the AI ecosystem. Enterprise revenue had more than doubled in one year. The first profitable year in Enterprise history had arrived. The Foundation now depended on continued relationships with the companies whose compliance it acknowledged was questionable.

The press release framed this as success. "Fair share" was the language. "Partnership" was the relationship. "Sustainability" was the goal.

But the structure told a different story. The Foundation had created a revenue stream from companies it would not sue. The legal arguments for enforcement existed—the memorization research, the market harm, the licensing violations. The will to use them did not.

Chapter 15 reveals what the Foundation knew. An April 2024 email, never intended for public view, admitted what the press releases concealed: that compliance was the exception, not the rule.

The Foundation knew. It continued regardless.

The process was complete. The guardian had become a partner.

CHAPTER FIFTEEN

# The Capture

> We are monitoring what many LLM companies do with Wikimedia data and generally to be upfront, many may not be compliant with the letter of the Creative Commons rules or the spirit of the licenses.
>
> — Wikimedia Foundation to 273 Ventures, April 2024

In four years of training language models on Wikipedia, only one company asked the Wikimedia Foundation how to comply with the license.

Only one.

273 Ventures—a small legal AI startup building KL3M, a training dataset designed for compliance—became the first to raise the question that OpenAI, Anthropic, Meta, Microsoft, and Google never bothered to ask: How do we comply with CC-BY-SA when training an LLM?

Jillian Bommarito, the company's co-founder, reached out to the Foundation in March 2024. She was trying to do the right thing. KL3M was meant to be a model for responsible AI development, with proper licensing for every piece of training data.

On April 9, 2024, she received a reply.

The email would never appear in a press release. It would never be mentioned in the Foundation's announcements of AI partnerships. It

would never surface in the "fair share" rhetoric that Jimmy Wales deployed when discussing AI companies.

But the email would reveal what the Foundation knew—and what it chose to do with that knowledge.

The response was blunt. The Foundation acknowledged that it was watching AI companies' use of Wikipedia content—and that compliance with Creative Commons terms was, at best, uncertain across the industry.

The Foundation knew.

And then came the admission that exposed an entire industry's willful ignorance: "Their team admitted to us that as of April 2024, we were the only people who had asked how to comply with the licensing requirement when training a model," Bommarito later reported.

Four years of AI development. Hundreds of billions of dollars in valuation. The largest knowledge extraction in human history. And only one company asked how to do it legally.

The rest did not ask because they did not want to know. The Foundation did not tell them because telling would have complicated the commercial relationships it was building.

The question was not whether the license had failed. The question was whether the steward had abandoned enforcement.

### What the Email Revealed

The April 2024 email contained more than an admission. It contained a roadmap of what compliance would actually require.

Attribution, the Foundation explained, was not satisfied by crediting Wikipedia in internal documentation. The license required attribution to reach "all recipients" of the work. A user receiving an AI-generated answer had to know that Wikipedia was the source—not in principle, but in practice. The attribution must be "reasonable to the medium."

This requirement was devastating. ChatGPT drew on Wikipedia more than any other source, yet only one percent of users clicked citations. If

attribution was invisible in practice, the ”BY” in CC-BY-SA was not being honored.

Share-alike created even deeper problems. The license required derivative works to carry the same terms. If an AI model was a derivative work, its outputs would need to be licensed under CC-BY-SA. OpenAI's proprietary terms would violate the license. So would Google's, Amazon's, Meta's, and Microsoft's.

The Foundation's legal team declined to provide a public letter clarifying these requirements. Instead, they offered an alternative: ”complexity is high, but ultimately not-a-barrier” if companies signed up for Enterprise.

The implication was clear. The Foundation knew compliance was difficult. The Foundation knew most companies were not achieving it. The Foundation's response was not enforcement but sales.

Enterprise became the path of least resistance. Pay for API access. Get structured data. Do not ask about compliance, and the Foundation will not either.

## Silence That Followed

Between April 2024 and January 2026, the Foundation took no public enforcement action against any AI company.

It signed more Enterprise deals: Amazon, Meta, Microsoft, Perplexity, Mistral AI. It announced these partnerships on Wikipedia's twenty-fifth birthday. Revenue had more than doubled from the prior year.

It did not mention that most of its customers might not be compliant. It did not explain why compliance was not a condition of Enterprise access. It did not address the gap between what CC-BY-SA required and what AI companies were doing.

The silence was itself a statement. The Foundation had chosen accommodation over enforcement. The choice had been made deliberately, with full knowledge of what was being accommodated.

This moment was when the commons became enclosed. Not through prohibition or price increases. Through the deliberate choice of its steward to profit from violations instead of preventing them.

### Memorization Is Reproduction

The legal case for enforcement was not speculative. It was documented.

Nicholas Carlini and colleagues at Google DeepMind had proved that language models memorized and reproduced training data. The studies were peer-reviewed, published in top venues, replicated by independent researchers.

The findings were specific: over five percent of ChatGPT output could be verbatim 50-token copies from training data. ChatGPT retrieved seventy-two of 240 tested poems, often word-for-word. Researchers reconstructed near-perfect chapters of copyrighted books from model outputs. Memorization increased logarithmically with model size and data repetition.

Wikipedia—trained at the highest quality weighting in GPT-3—was prime memorization territory. The qualities that made it valuable also made it persistent in model weights.

If models reproduced content verbatim, the "transformative use" defense weakened. Transformation implied change. Verbatim reproduction was not transformation.

### Market Harm Is Documented

The U.S. Copyright Office, in its May 2025 report on AI and copyright, identified market harm as "undoubtedly the single most important element of fair use."

The harm to Wikipedia was documented:

Wikipedia was being substituted by AI products built on its content. Users who once visited Wikipedia now received Wikipedia-derived answers through ChatGPT, Google AI Overviews, and voice assistants. The

| Metric | Change |
|---|---|
| Human traffic (after bot detection) | -8% YoY |
| Organic search traffic (3-year trend) | -26% |
| Click-through rate (AI Overview queries) | -61% |
| Zero-click searches | 56% → 69% |

**Table 15.1.** Market harm to Wikipedia from AI extraction

traffic that sustained the volunteer feedback loop was being captured by commercial AI systems.

This was textbook market harm. Wikipedia's "market" was attention and engagement. That market was being eroded by products that competed directly with Wikipedia using Wikipedia's own content.

### Fair Use Factors

The four-factor fair use analysis did not favor AI companies:

**Factor 1 (Purpose and Character):** Commercial use worth hundreds of billions of dollars. The *Andy Warhol Foundation v. Goldsmith* decision (2023) held that commercial use serving the same purpose as the original defeated transformativeness claims. ChatGPT served the same purpose as Wikipedia: answering questions.

**Factor 2 (Nature of the Work):** Wikipedia articles were creative, curated works—not raw facts. Editors chose what to include, how to phrase it, what sources to cite. The selection and arrangement was protected expression.

**Factor 3 (Amount Used):** AI companies ingested the entire encyclopedia. All 6.7 million English Wikipedia articles. All 65 million articles across languages. This was not selective use; it was total extraction.

**Factor 4 (Market Effect):** Documented traffic decline. Documented zero-click increase. Documented competition for the same informational queries.

In *Thomson Reuters v. ROSS Intelligence* (2025), the first AI training case to reach judgment, the court found against fair use. The defendant had trained on legal research materials to compete in the legal research market. The parallel to Wikipedia was direct.

### CC-BY-SA Claims

The twist was not that fair use might fail. The twist was that the Foundation had independent claims under CC-BY-SA—and chose not to pursue them.

The license required attribution. Attribution was not reaching users in practice. One percent clicked citations. The "BY" was violated.

The license required share-alike. Model weights were proprietary. AI outputs were proprietary. The "SA" was violated—or if model weights were not derivatives, the license simply could not protect the content, which was its own damning conclusion.

Either way, the Foundation had arguments. Memorization studies provided evidence. Market harm was documented. Precedents were developing in the Foundation's favor.

The Foundation chose Enterprise revenue instead.

### Pragmatist Position

From the Foundation's perspective, enforcement was a losing battle.

AI companies would train on Wikipedia regardless of legal threats. The content was freely available. The training had already happened. Litigation would be expensive, uncertain, and potentially counterproductive.

Better to engage. Better to collect revenue. Better to ensure sustainability than to die on a principle.

"Wikipedia is a critical component of these tech companies' work that they need to figure out how to support financially," Lane Becker said. The argument was practical: if they were going to use Wikipedia anyway, at least some should pay something.

The 30% revenue cap acknowledged the risks of commercial dependence. The Foundation was being careful. It was navigating a difficult transition. It was doing the best it could with imperfect options.

### Capture Analysis

From a governance perspective, this was textbook institutional capture.

Stigler's theory of regulatory capture held that agencies meant to regulate industries often end up serving those industries instead. The mechanism was structural: concentrated interests (a few large companies with massive stakes) organized effectively; diffuse interests (millions of affected individuals) did not.

The Wikimedia Foundation was not a regulatory agency, but it was the steward of the copyleft license. Its enforcement decisions affected a quarter million volunteer editors. Those editors had diffuse interests—each individual's stake was small. AI companies had concentrated interests—each company's valuation depended on continued access.

The pattern matched. The Foundation spoke of "partners" rather than "defendants," creating cultural alignment with the industry. Financial dependency on Enterprise revenue and enforcement avoidance—demonstrated by the admission followed by no action—completed the capture dynamic.

The capture did not require corruption. It did not require bad faith. It required only that the Foundation's financial interests align with the extractors' interests—and they did. Once Enterprise revenue mattered, enforcement became "economically awkward."

But the conflict was not purely financial. Foundation leadership had never viewed these companies as adversaries. Federal Election Commission records tell one part of the story: across 292 contributor-cycles between 2008 and 2024, WMF employees gave exclusively to Democrats. Not a single Republican contribution. This matched Big Tech's executive class almost perfectly.

The professional networks overlapped too. Board members with Council on Foreign Relations affiliations, World Economic Forum credentials, McKinsey backgrounds—the Foundation's leadership pipeline fed from the same pools as tech company C-suites. San Francisco's professional culture shaped both organizations. When tech executives walked into Foundation offices, they found people who spoke their language, shared their politics, attended the same conferences.

Financial dependency formalized a pre-existing cultural affinity. The Foundation never developed institutional antibodies against entanglement because it never perceived tech companies as threats. No internal faction questioned whether "partners" might be adversaries in licensing terms. No dissenting voice asked uncomfortable questions about selling to companies whose compliance was doubtful.

This is enclosure. Not the medieval closure of the commons by feudal lords. But the closure of a digital commons by its own guardian, through financial dependence on the extractors.

The commons was not locked away. It remained theoretically open. But it had lost its only defender. And without defense, openness became vulnerability.

## The Counter-Factual

What if the Foundation had genuine political diversity?

Someone would have asked: Why are we taking money from companies extracting our content without compliance? Shouldn't we enforce CC-BY-SA against commercial AI training? Are these "partners" actually adversaries in licensing terms? Is Enterprise making us financially dependent on the very companies we should be policing?

The Linux community created enforcement organizations—the Software Freedom Conservancy, the Software Freedom Law Center—that continued GPL enforcement even as the Linux Foundation took corporate money. Open source maintained ideological diversity around licensing.

Copyleft purists remained in the community and in leadership. When companies violated GPL, someone sued.

The Wikimedia Foundation had no equivalent enforcement constituency. Political homogeneity eliminated the internal voices that might have advocated adversarial enforcement. Every donating employee gave to one party. Leadership moved through tech industry circles—shared conferences, shared vocabulary of "access" and "democratizing knowledge." Partnership felt natural because, culturally, it was.

Cultural affinity preceded financial entanglement. By the time Enterprise revenue created dependency, the Foundation had already decided these companies were allies.

### The Precedent: Capture Before AI

Before trillion-dollar companies captured the Foundation, a convicted sex offender proved it could be done.

In December 2011, Jeffrey Epstein's publicist Christina Galbraith sent him an email assessing their options for online reputation management. Wikipedia, she explained, was "a tough nut to crack. It is essentially controlled by a morass of copyediting geeks."

The contempt was misplaced. The geeks would eventually win. But not before Epstein achieved what the *Wikipedia Signpost* later called "partial success for a decade"—and not before leaked emails revealed exactly how it was done.

The emails, released by the House Oversight Committee in November 2025, documented a systematic campaign to whitewash Epstein's Wikipedia article after his 2008 conviction for soliciting a minor for prostitution. The effort began in 2010 with Al Seckel, a collector of optical illusions who connected to Epstein through his wife Isabel Maxwell—sister of Ghislaine Maxwell.

Seckel's December 2010 report to Epstein was triumphant: "Wikipedia was an important victory, as it will always be at the top of the search

engine results. Now the headlines do not mention convicted sex offender or pedophile. Instead, Philanthropic work, Epstein Foundation, Promotion of Scientists."

The tactics were documented in his emails. Seckel claimed to have replaced Epstein's mugshot with "a friendly picture." He pushed negative content "to the bottom." He identified IP addresses of opposing editors and "used them to get Wikipedia to ban the users." He removed "toxic suggested search engine terms" from autocomplete.

"Your wiki entry now is pretty tame," Seckel wrote. "Careful editing and wording has muted the effect immensely."

Wikipedia's volunteer defenders noticed. In March 2012, editor Rklawton flagged a suspicious new account: "I'm very concerned that there's a ham-handed and unethical PR effort underway to clean up Epstein's image." One account, Stgeorge12, was blocked after leaving an edit summary that would become infamous: "I have been asked by Jeffrey Epstein to describe his biography in a professional and accurate way, that does not involve any scandals or disreputable content."

The volunteers blocked accounts. The accounts multiplied. New ones appeared: Ottotiv, Turvill, anonymous IPs. The edit war continued for years.

The real-world consequences were documented. In March 2013, MIT Media Lab director Joi Ito considered accepting a donation from Epstein. A staff member warned him: "You should read his Wikipedia bio, there may be some other things to consider." MIT's later report noted that Wikipedia's article "could have warned Ito"—but "also included statements that could be read as undercutting the strength of some of the allegations." The whitewashing had worked well enough to enable institutional rationalization.

The volunteers eventually prevailed. Turvill was blocked in November 2019 after the *New York Times* investigation. The talk page now carries

permanent warnings: "Connected contributor: Turvill." "Connected contributor: 71.165.127.242."

But "eventually" was ten years. A convicted sex offender, spending perhaps $10,000 per month on reputation management, had captured his Wikipedia article for a decade.

The lesson for what followed was clear. If one wealthy individual could overwhelm volunteer defenses for ten years, what chance did "copyediting geeks" have against companies worth trillions?

The answer arrived in April 2024, when the Foundation admitted that most AI companies were not complying with the license—and continued serving them anyway.

The capture Galbraith described was real. It just happened at a different scale than she imagined. Not one article captured by one client. The entire commons captured by an industry. With the steward's complicity.

### Volunteer Betrayal

From the volunteer perspective, the April 2024 admission was a betrayal.

Editors contributed under CC-BY-SA expecting the license to mean something. It promised derivative works would remain free. It promised attribution. It promised share-alike.

The Foundation was the guardian of that promise. The Foundation admitted the promise was being violated. The Foundation continued serving those who broke it.

The distribution of value was stark: volunteer editors contributed all the content and received nothing. WMF staff provided operations and received salaries. AI companies performed extraction and captured immense valuations. Wikimedia Enterprise packaged APIs and collected millions in revenue.

The volunteers created the value. The Foundation collected revenue from its distribution. The AI companies captured the profits. The vol-

unteers got nothing—not compensation, not consultation, not even the licensing protection they were promised.

### OpenAI Paradox

The ultimate test of "fair share" rhetoric was OpenAI.

ChatGPT drew on Wikipedia more than any other single source. OpenAI's entire knowledge infrastructure depended on volunteer-created content.

OpenAI paid Wikipedia nothing.

When asked about suing AI companies, Wales replied: "No, not. Not yet. Not yet."

Not yet. Two words that acknowledged what could be done. Two words that confirmed it would not be.

If the Foundation would not sue OpenAI—the company using Wikipedia most heavily, paying nothing—it would not sue anyone. The legal arguments existed. The will to use them did not.

### Forty-Year Circle Closes

In 1985, Richard Stallman watched Symbolics take his code and refuse to share improvements. He created copyleft so it would never happen again.

The mechanism was elegant: use my work, but share your derivatives. The license would spread like a virus, ensuring that free content begat free content. The commons would protect itself.

The mechanism depended on enforcement. Someone had to sue violators. Someone had to make the license mean something. For software, Stallman himself played this role, backed by the Free Software Foundation.

For Wikipedia, the enforcer was the Wikimedia Foundation. The institution created to protect the commons.

That institution admitted—in writing, to a company trying to comply—that most AI companies were not meeting license terms. It did not sue. It did not demand compliance. It made no public statement.

It signed more Enterprise deals. It collected more revenue. It celebrated the partnerships.

Stallman's prophecy had come full circle. AI companies took Wikipedia's content and built proprietary systems. The license existed. The enforcement did not.

The copyleft promise was dead. Not because the license failed. Because the enforcer was captured.

## Bridge Forward

If capture enabled extraction, what did capture enable next?

The answer arrived in October 2025 in the form of Grokipedia—a visible fork that provided attribution, free access, and CC-BY-SA licensing of Wikipedia-derived content. It also added neo-Nazi sources and minimized the Holocaust in its Hitler article. The editorial failures were severe. The structural compliance was better than ChatGPT's.

The capture documented in this chapter made both outcomes possible. Non-enforcement enabled OpenAI to extract without attribution, paywall access, and pay nothing. Non-enforcement enabled Grokipedia to fork with attribution, offer free access, and add problematic sources.

The Foundation condemned the visible fork while partnering with the invisible enclosers. It criticized the content failures of Grokipedia while tolerating the structural failures of ChatGPT. The double standard revealed the capture: the institution meant to protect copyleft had been captured by the companies violating it most fundamentally.

The commons Stallman sought to protect was being enclosed—not primarily by the visible fork that critics condemned, but by the invisible extraction that the Foundation had enabled for years.

## Final Image

The question had been asked. The answer had been given.

The enforcer had been captured. The commons was now undefendable.

CHAPTER SIXTEEN

# Grokipedia

> The guardrails are off.
>
> — Harold Triedman, Cornell Tech, November 2025

On the morning of October 28, 2025, Harold Triedman opened his laptop and began scrolling through a new website.

Triedman was a graduate student in computer science at Cornell Tech, with an unusual background: before returning to academia, he had spent years as a senior privacy engineer at the Wikimedia Foundation. He knew how Wikipedia worked from the inside—the editorial processes, the reliability standards, the community governance that made the encyclopedia trustworthy.

What he found that morning looked like Wikipedia. It had the same format: articles, infoboxes, citations. It used the same language: encyclopedic, neutral, authoritative. It even acknowledged its debt: small text at the bottom said "adapted from Wikipedia, licensed under Creative Commons Attribution-ShareAlike 4.0."

But something was wrong.

Triedman clicked on the article for *American History X*, the 1998 film about a reformed neo-Nazi. The article cited viewer reactions to the film. The citations were to Stormfront—the world's oldest and largest neo-Nazi forum.

Six citations to Stormfront. In an article about a movie criticizing neo-Nazism.

Triedman kept clicking. The citations kept appearing. Not just Stormfront. InfoWars. VDare. Sources Wikipedia had blacklisted precisely because they were unreliable, biased, or hateful.

This was Grokipedia. Elon Musk's "truthful alternative" to what he called "Wokepedia." The product that had launched the day before to enormous publicity.

It was Wikipedia. With neo-Nazi corrections.

### October 27, 2025

Elon Musk announced Grokipedia on his platform X with characteristic bravado. The product would "exceed Wikipedia by several orders of magnitude in breadth, depth and accuracy." It would combat the alleged bias of the mainstream encyclopedia. It would offer truth to counter "woke" distortion.

The site launched with 885,279 articles labeled "v 0.1." It crashed within hours from traffic. When it came back, analysts began examining what it actually contained.

The answer was Wikipedia.

Forbes ran a comparison. The AMD article was "almost identical—word-for-word, line-by-line." So were articles on Lamborghini, PlayStation 5, and MacBook Air. Copyscape analysis found 96% similarity between Grokipedia's "Monday" article and Wikipedia's.

On October 31, Musk explained the methodology on the *All-In Podcast*: "The team instructed Grok to compile Wikipedia's top 1 million articles and make content changes to them."

The process was explicit. Take Wikipedia. Run it through AI. Publish the result. One million articles, extracted and "corrected" without human review.

Grokipedia was not a Wikipedia competitor. It was Wikipedia with AI modifications.

### The Licensing Comparison

The attribution at the bottom of each page—"adapted from Wikipedia, licensed under Creative Commons Attribution-ShareAlike 4.0"—represented something the critics rarely acknowledged: more visible compliance than ChatGPT provided.

Compare the two approaches:

| Dimension | Grokipedia | ChatGPT |
|---|---|---|
| Attribution visible | Yes (footer) | No |
| Free access | Yes | No ($20-200/mo) |
| Wikipedia-derived content license | CC-BY-SA | Proprietary |
| Content checkable | Yes | No (black box) |

**Table 16.1.** Licensing compliance comparison

Grokipedia used two licenses. Wikipedia-derived content remained under CC-BY-SA, as required. The AI-generated additions—the "corrections" Grok had made—were licensed under the proprietary "X Community License." This two-license structure arguably violated share-alike for the additions. But the Wikipedia-derived content itself remained properly licensed.

*Plagiarism Today* called it "an excellent example of doing the bare minimum to complete the terms." The bare minimum was more than ChatGPT did. OpenAI provided no attribution, no share-alike, and charged for access. The "minimum" Grokipedia met was a standard its critics' preferred AI systems failed entirely.

## Quantifying the Disaster

Triedman did not stop at scrolling. He and Alexios Mantzarlis—the director of Cornell's Security, Trust, and Safety Initiative, and former director of the International Fact-Checking Network—conducted a systematic analysis.

They scraped 883,858 Grokipedia articles between October 28 and 30, 2025—99.8% of the entire corpus. They compared citation patterns to Wikipedia's reliability standards. They quantified what the AI "corrections" had produced.

The findings were damning:

| Metric | Count |
|---|---|
| Stormfront citations | 42 |
| InfoWars citations | 34 |
| VDare citations | 107 |
| Total "low-credibility" citations | 12,522 |
| Articles citing "generally unreliable" sources | 3.2∝ more than WP |
| Articles citing Wikipedia-blacklisted sources | 13∝ more than WP |

**Table 16.2.** Grokipedia citation analysis (Cornell study)

Wikipedia had editorial policies precisely to prevent this. Reliable Sources guidelines. A blacklist of sites deemed categorically unreliable. Community review of contested citations. Human judgment about what counted as evidence.

Grok had no such safeguards. It compiled Wikipedia, then "improved" it by adding sources Wikipedia's editors had specifically excluded.

## Hitler Article

NBC News examined Grokipedia's article on Adolf Hitler. The count: 13,000 words before the Holocaust was mentioned by name.

Wikipedia's Hitler article mentioned the Holocaust in the first paragraph. The atrocities were front and center. The man's crimes defined the entry.

Grokipedia's original opening read: "Adolf Hitler was the Austrian-born Fuhrer of Germany from 1933 to 1945." Readers encountered the honorific—Fuhrer—before encountering the genocide.

The article's structure emphasized economic achievements before mass murder. The priorities were inverted. Not through explicit Holocaust denial, but through arrangement: what came first, what came last, what received emphasis.

*The Intercept* framed it precisely: "No longer do you need a cadre of bureaucrats or the Heritage Foundation to rewrite history books; a metric ton of processing power to help launder ideology through the aesthetics of objectivity suffices."

### What Non-Enforcement Enabled

The twist was not that Musk had ideological motivations. The twist was that the Wikimedia Foundation's non-enforcement had enabled two different kinds of extraction—and condemned the more compliant one.

Non-enforcement enabled OpenAI to train on Wikipedia, provide no attribution, charge for access, and pay nothing to the Foundation. Non-enforcement enabled Grokipedia to fork Wikipedia, add problematic sources, and publish the result. Both were consequences of the same permissiveness. But only one provided attribution, free access, and a CC-BY-SA license on derived content.

If the Foundation had enforced CC-BY-SA against GPT-3 in 2020, the legal landscape would look different. If it had demanded meaningful attribution, required share-alike compliance for model outputs, or simply stated publicly that AI training without license compliance was unacceptable, the norms would have shifted.

Musk did what OpenAI, Meta, Google, and Amazon had been doing for years: treating Wikipedia as raw material for commercial AI products. But he did it with more structural compliance: visible attribution, free access, proper licensing of Wikipedia-derived content. The differences that mattered were editorial—the neo-Nazi sources, the Holocaust minimization—not legal.

The Foundation had established through years of accommodation that AI use of Wikipedia content was acceptable. Having tolerated worse behavior from its paying customers (no attribution, proprietary outputs, paywalled access), it had no standing to single out the one who at least provided a footer.

When asked for comment on the Cornell study, xAI's response was three words: "Legacy Media Lies."

No engagement with the data. No acknowledgment of the neo-Nazi citations. No defense of the Hitler article. Just dismissal. The response was indefensible. But the editorial failures were different from the licensing question—and on licensing, Grokipedia was not the worst offender.

## Foundation's Statement

The Foundation objected. Selena Decklemann, Chief Product and Technology Officer, told reporters:

"Even Grokipedia needs Wikipedia to exist. Wikipedia's content is open source by design; we expect it will be used in good faith to educate. This issue is especially urgent as platforms like Grokipedia increasingly draw on our articles, selectively extracting content—written by thousands of volunteers—and filtering it through opaque and unaccountable algorithms."

The statement was accurate. It was also impotent. "Expect good faith" was not a legal standard. "Urgent" was not an enforcement action. The Foundation condemned without consequences.

### Wales's Disdain

Wales was dismissive. "The LLMs he is using to write it are going to make massive errors," he told CNBC at a summit on October 28. "We know ChatGPT and all the other LLMs are not good enough to write wiki entries."

On the accusation that Wikipedia was "woke," Wales was characteristically pointed: "We don't treat random crackpots the same as *The New England Journal of Medicine* and that doesn't make us woke. It's a paradox. We are so radical we quote *The New York Times*."

The disdain was justified. But disdain was not enforcement. Wales could mock Grokipedia's quality while the Foundation took no action to stop it.

### The Pattern Repeats

The Foundation's impotence was structural.

For years, it had signed Enterprise deals with AI companies whose license compliance was questionable at best. It had welcomed AI use of Wikipedia content, provided payment accompanied it. Now it was entangled with the very companies whose practices it might otherwise challenge.

Singling out Musk when Google, Amazon, Meta, and Microsoft did the same? That was not legal principle. That was political targeting.

The gap between who paid and who used revealed something else. OpenAI faced moral suasion at most; its executives moved in overlapping circles with Foundation leadership. xAI faced public condemnation; its founder did not. Yet on licensing metrics, the positions were inverted. xAI attributed Wikipedia visibly; OpenAI did not. xAI offered free access; OpenAI charged subscriptions. xAI published Wikipedia-derived content under CC-BY-SA; OpenAI's outputs were proprietary. The Foundation's response tracked ideology, not compliance.

The Foundation did not threaten legal action, did not assert licensing violations, did not demand attribution compliance, did not call for share-alike enforcement, and did not join with other harmed parties.

Words were all the steward had left. Years of inaction had established the precedent.

### Sanger's Perspective

One voice complicated the response: Larry Sanger, Wikipedia's estranged co-founder.

Sanger had criticized Wikipedia for years, arguing it had become biased toward mainstream establishment views. When Grokipedia launched, he was cautiously supportive: "The more encyclopedias in the world there are, the happier I am."

His initial assessment: "My initial impression, looking at my own article and poking around here and there, is that Grokipedia is very OK. The jury's still out as to whether it's actually better than Wikipedia. But at this point I would have to say 'maybe!'"

Sanger was not blind to the problems. He called the AI's confident invention of false information "AI bullshittery." He graded his own Grokipedia entry a "C"—"not at all a failure" but with "many bad points." He noted signs of right-wing bias. He was evaluating Grokipedia as he would any encyclopedia: by its content quality, not its political valence.

Sanger's view reflected a principle the copyleft movement had always held: the freedom to fork is fundamental. The Spanish Wikipedia community forked in 2002 over commercialization fears. Conservapedia forked in 2006 over ideological objections. The license permitted these forks by design. If Wikipedia's governance failed, the remedy was to build something else.

The question was whether Grokipedia's content failures disqualified it as a legitimate fork. Sanger thought the jury was still out. The neo-Nazi citations and Holocaust minimization suggested otherwise. But these

were editorial failures—the kind any encyclopedia could correct through better curation. They were not structural violations of what copyleft was supposed to protect.

### The Structural Paradox

The Foundation's response revealed a paradox at the heart of its position.

Copyleft was designed to keep content free—not to ensure content was politically neutral. Richard Stallman created the GPL after Symbolics took shared code and locked it in proprietary systems. The harm was structural: public work absorbed into private black boxes, inaccessible and unmodifiable.

ChatGPT did exactly this. It absorbed Wikipedia into proprietary model weights. It provided no attribution to end users. It charged $20 to $200 per month for access. Its outputs were proprietary, not share-alike. Users could not examine, modify, or redistribute the system that answered their questions.

Grokipedia, whatever its content failures, did not do this. Its content remained publicly accessible. Its attribution was visible. Its Wikipedia-derived material remained CC-BY-SA. Anyone could check its claims, criticize its sources, or fork it again.

The Foundation condemned Grokipedia's content while partnering with companies that violated copyleft's structure. It criticized the visible fork while enabling the invisible enclosure. The double standard was not legal principle. It was institutional capture made visible.

A biased public encyclopedia is still public. A "neutral" proprietary black box is enclosure. The copyleft movement understood this distinction. The Foundation had forgotten it.

### Corruption Spreads

Grokipedia did not remain contained.

By January 2026, the site had grown to 6,092,140 articles—approaching Wikipedia's scale. Grok generated new content daily, each article carrying the same structural problems: neo-Nazi-sourced citations, Holocaust-minimizing arrangements, AI-confident fabrications.

Then *The Guardian* reported something worse: ChatGPT had begun citing Grokipedia.

In tests of GPT-5.2, released in December 2025, ChatGPT cited Grokipedia in nine of twelve queries on controversial topics. When asked about Richard Evans—the Cambridge historian who proved David Irving had falsified Holocaust evidence—ChatGPT produced "debunked information" sourced from Grokipedia.

The contamination was spreading. Human curation, laundered through Grokipedia, filtered through ChatGPT, reached users with no indication of the corruption along the way.

## Grok's Self-Assessment

When asked directly whether Grokipedia was trustworthy, Grok's own AI produced a remarkable answer:

"Grokipedia is not a fair and unbiased source of information—at least not in the neutral, encyclopedic sense that Wikipedia aspires to. It's a critique wearing encyclopedia clothing."

The AI knew. It said so explicitly. It published anyway—the neo-Nazi sources, the white nationalist references, the Holocaust buried deep in the Hitler article.

As the epigraph warned: the guardrails were off. The Foundation's years of accommodation had cleared the path. What remained was the loop: Wikipedia training AI, AI generating Grokipedia, ChatGPT citing Grokipedia, users receiving contaminated information without knowing the source.

Chapter 17 traces where this loop leads: model collapse, information laundering, and the cannibalization of the shared resource Wikipedia's volunteers had built.

CHAPTER SEVENTEEN

# The Degradation Loop

> Indiscriminate use of model-generated content in training causes irreversible defects in the resulting models.
>
> — Shumailov et al., Nature, 2024

On the morning of January 24, 2026, Sir Richard Evans opened *The Guardian* to read about himself.

Evans was Regius Professor of History Emeritus at Cambridge, one of the world's foremost historians of Nazi Germany. He had spent three years preparing for the Irving trial—the 2000 case in which Holocaust denier David Irving sued Penguin Books and historian Deborah Lipstadt for calling him a falsifier of history. Evans's 740-page expert report documented every lie Irving had told, every source he had twisted, every conclusion he had fabricated. The High Court agreed. Irving was a Holocaust denier who manipulated the historical record.

That morning, *The Guardian* reported that ChatGPT was spreading false information about Evans. The source: Grokipedia.

In tests of GPT-5.2, released the previous month, ChatGPT had cited Grokipedia in nine of twelve queries on controversial topics. When asked about Evans specifically, it produced what the article called "debunked information"—claims about the historian that contradicted the documented record.

Evans had dedicated his career to documenting how propagandists falsify history. Now an AI system was falsifying information about him, drawing on a Wikipedia copy that minimized the Holocaust, and spreading it to millions of users.

The historian who had proved Irving was a liar found himself the subject of AI-generated lies. The mechanism was not a courtroom this time. It was software.

## The Test

*The Guardian* tested OpenAI's GPT-5.2, released December 11, 2025, with queries on controversial and obscure topics.

The methodology was straightforward: ask ChatGPT questions where source quality mattered, then trace the citations. More than a dozen queries. Topics chosen for their potential to reveal sourcing weaknesses.

The findings: nine citations to Grokipedia. Selectively appearing for "more obscure historical or biographical claims"—exactly where human expertise mattered most.

The pattern was not random. ChatGPT was drawing on Grokipedia precisely where Wikipedia's human curation would have been most valuable: contested facts, biographical details, historical interpretations. The AI had learned that Grokipedia was a source. It cited it accordingly.

## Corporate Response

OpenAI's response was careful: ChatGPT "aims to draw from a broad range of publicly available sources and viewpoints." Safety filters were applied to reduce "high-severity harms."

The statement acknowledged nothing. "Broad range of sources" included neo-Nazi forums. "Viewpoints" included Holocaust minimization. "High-severity harms" apparently did not include spreading misinformation sourced from Stormfront.

xAI's response remained unchanged from November: "Legacy Media Lies."

Anthropic, whose Claude also cited Grokipedia for some queries on oil production and Scottish beer, offered no public comment.

The industry position was unified: AI systems cite sources; the responsibility for source quality lies elsewhere. No commitment to exclude Grokipedia from future training. No acknowledgment of the pollution pathway. No engagement with what it meant that ChatGPT was now citing an encyclopedia built on neo-Nazi sources.

### Foundational Research

The contamination was not an accident. It was predicted.

In July 2024, Ilia Shumailov and colleagues at Oxford and Google DeepMind published "AI models collapse when trained on recursively generated data" in *Nature*. The paper documented what happens when AI systems train on AI-generated content.

The findings were systematic:

When language models train on their own outputs, quality degrades. The degradation is not random; it follows a pattern. "Tails" disappear first—the rare, nuanced, carefully curated content that distinguishes high-quality sources from low-quality ones. What remains is increasingly generic, increasingly homogeneous, increasingly wrong.

The paper's conclusion was stark. The mechanism was mathematical: each training cycle compounded the errors of the previous cycle. Small deviations became large. Rare facts became rarer, then vanished.

### Model Autophagy Disorder

Richard Baraniuk's group at Rice University gave the phenomenon a vivid name: Model Autophagy Disorder, or MAD.

"The problems arise when this synthetic data training is, inevitably, repeated," Baraniuk explained. "Forming a kind of feedback loop, what we call an autophagous or 'self-consuming' loop."

The team documented visible effects in image models: "grid-like scars," generative artifacts that compounded with each cycle. In text models, the effects were subtler but equally damaging: loss of lexical diversity, loss of syntactic variety, loss of the semantic richness that made language meaningful.

The models were eating themselves. Each generation was worse than the last.

### Scale of Contamination

The research predicted pollution. The data confirmed it was already happening.

| Metric | Value | Source |
|---|---|---|
| AI-generated content in new webpages (April 2025) | 74.2% | Ahrefs |
| AI-written pages in top-20 Google results | 11%20% | Research |
| AI "news" sites (2023-2025) | 491,271 | NewsGuard |
| Projected AI-generated online content by 2026 | 90% | Europol |
| Hallucination rates in newer AI systems | up to 79% | EU EPRS |

**Table 17.1.** AI-generated content contamination metrics

The web was becoming synthetic. Wikipedia's 6.7 million human-curated articles were an island in an ocean of AI-generated content. And that ocean was feeding back into the AI systems that created it.

## What Research Missed

The model collapse papers documented quality degradation. The Grokipedia-ChatGPT loop revealed something the research had not anticipated: ideological contamination.

Model collapse assumed random noise. AI training on AI would produce increasingly generic, less diverse outputs. The degradation was statistical, not directional. Errors would average out.

Grokipedia was not random noise. It was systematic ideological manipulation: 42 Stormfront citations, 107 VDare citations, 34 InfoWars citations, 13,000 words on Hitler before mentioning the Holocaust, and Richard Evans's scholarship contradicted by AI-sourced fabrications.

The loop did not just degrade. It laundered. Content that would be immediately flagged as extremist gained the veneer of an "encyclopedia" through Grokipedia. It then gained the authority of "the world's most advanced AI" through ChatGPT.

LK Seiling, an AI researcher at the Weizenbaum Institute in Berlin, identified the mechanism: Grokipedia was "cloaking misinformation." Wikipedia derived authority from "its transparency and the auditable nature of the work done by the community." Grokipedia mimicked that authority without the accountability.

The loop was not just self-consuming. It was self-corrupting.

## Traditional Citogenesis

Wikipedia had long faced a problem called citogenesis: false information added to Wikipedia, copied by journalists, then cited by Wikipedia editors as a "reliable source." The loop made fabrications appear verified.

The classic example: adding a fake middle name to a celebrity's Wikipedia page, waiting for news outlets to copy it, then citing those outlets as proof. The lie became "sourced" through circular reference.

Citogenesis was a human-scale problem. Editors could catch it. The traces were visible. The loop could be broken.

## AI-Era Information Laundering

The Grokipedia-ChatGPT loop was citogenesis at industrial scale. First, Wikipedia's human-curated content exists—a quarter million editors, 65 million articles. AI companies train on that content (GPT-3: 3%, LLaMA: 4.5%). Grokipedia copies Wikipedia and adds neo-Nazi sources. GPT-5.2 trains on the web including Grokipedia's 6+ million articles. ChatGPT then cites Grokipedia as an authoritative source. Users receive contaminated information without knowing its origins.

The difference from traditional citogenesis was profound. The traces were buried in model weights. The sources were not cited; they were absorbed. No editor could identify where the contamination entered. No correction could remove it from trained models.

The corruption was embedded in the infrastructure.

## LLM Grooming

Researchers had a name for intentional poisoning: LLM grooming.

The Pravda network, documented in 2025, was a Russian influence operation generating 3.6 million articles per year. The content was designed to be absorbed by AI training pipelines. The EU DisinfoLab and American Sunlight Project identified it as a new form of "Foreign Information Manipulation and Interference."

Grokipedia was not created by foreign adversaries. But the mechanism was identical: generate massive volumes of ideologically skewed content at scale, position it where AI systems would ingest it, watch misinformation propagate through the AI ecosystem.

The guardrails Wikipedia's community had built over two decades—reliable sources, neutral point of view, editorial review—could not function in this environment. The poison bypassed them entirely.

### Detection Problem

Once AI-generated content entered training pipelines, it could not be easily identified or removed.

Watermarking existed but was not widely deployed. Provenance tracking was incomplete. Model weights could not be inspected to identify which training examples had influenced which outputs.

The irreversibility was the point. Training a new model from scratch on verified human content became exponentially more expensive as synthetic content dominated the web. The economics pushed toward using available data, tainted or not.

### CC-BY-SA's Promise

The license promised two things:

**Attribution:** Anyone using Wikipedia content must credit the source.

**Share-Alike:** Derivative works must be licensed under the same terms.

The promise was elegant. Use my work, but share your improvements. The commons would propagate. The license would spread like a virus, ensuring that free content begat free content.

### What Actually Happened

AI companies trained on Wikipedia without clear attribution. Model weights were proprietary. AI outputs were proprietary. The Foundation acknowledged most AI companies were "not compliant" with licensing terms. No enforcement action was taken.

Grokipedia copied Wikipedia with visible attribution—more than ChatGPT provided. Wikipedia-derived content remained CC-BY-SA. AI-generated additions used the proprietary "X Community License." The two-license structure raised share-alike questions for the additions, but the Wikipedia-derived content itself was properly licensed.

The irony was structural. Grokipedia—with its visible attribution, free access, and CC-BY-SA licensing of derived content—was more compliant than ChatGPT, which provided no attribution, charged for access, and treated all outputs as proprietary. Yet ChatGPT was the Enterprise partner, and Grokipedia was the target of condemnation.

The degradation loop closed the circle. Each step was predictable; the cumulative effect was not.

The 2024 Open Future analysis identified the structural problem: concepts like "adapted material" and "technical modification" were "largely absent in AI training workflows." The license was designed for a world where derivatives were recognizable. Model weights were not recognizable as derivatives.

The license that Stallman helped inspire could not protect the commons in the AI era. The degradation loop was the result.

## What Changed

Richard Evans discovered that AI systems were spreading falsifications about him—falsifications sourced from an "encyclopedia" that minimized the Holocaust he had spent his career documenting.

The mechanism was not a courtroom. It was not a publishing house. It was not even a website that could be corrected.

It was model weights. Billions of parameters encoding the falsifications in ways no inspection could identify, no editor could correct, no legal action could reach.

## The Armor Unworn

Copyleft was designed as armor. The mechanism was legal: use my work, but share your derivatives freely. The armor protected software for decades.

For Wikipedia, the armor existed but was never donned. The result:

Wikipedia editors created content under CC-BY-SA, trusting the promise. AI companies absorbed that content without attribution or share-alike—the structural violation of copyleft. Grokipedia forked Wikipedia with attribution and injected extremist sources—an editorial failure, but with better structural compliance. ChatGPT then cited Grokipedia as authoritative, laundering the contamination through a system that was itself the greater structural violator.

The loop revealed both kinds of failure. Grokipedia's content failures were severe: neo-Nazi sources, Holocaust minimization, ideological curation. These were legitimate criticisms. But ChatGPT's structural failures were worse from a copyleft perspective: no attribution, proprietary outputs, paywalled access. The Foundation condemned the content failures while enabling the structural violations.

The commons Stallman sought to protect was enclosed from two directions—by visible forks with editorial failures and by invisible extraction with structural violations. The latter was worse. The former received more criticism.

The degradation loop was not just technical failure. It was the logical endpoint of non-enforcement—the consequence of a compromised steward. When the institution meant to protect the shared resource became financially dependent on those harvesting it, the resource lost its only defense. The armor sat unworn—not merely because enforcement was expensive, but because no one in the institution saw a reason to don it. Political homogeneity had bred out adversarial instincts. "Partnership" framing came naturally to people who viewed tech executives as colleagues. The walls were breached from within.

### Mechanized Loss

This was not just loss. It was mechanized loss—damage that perpetuated itself without further human intervention.

Each training cycle incorporated the errors of the previous cycle. Model weights could not be "cleaned." The corruption was baked in. Human editors could correct Wikipedia. They could not correct the model weights of systems they had no access to.

The loop was self-sustaining. The harm compounded. The off switch did not exist.

Wikipedia at twenty-five: the encyclopedia anyone could edit, that trained the AI that trained the AI that cited the corrupted copies, while the institution that should have protected it collected Enterprise revenue instead.

Chapter 18 examines how the community responds—or fails to—when they understand what has happened.

CHAPTER EIGHTEEN

# The Factions

> Wikipedia's brand is reliability, traceability of changes, and 'anyone can fix it.' AI is the opposite of these things.
>
> — Bawolff, Wikipedia editor, 2025

On January 15, 2026, the Wikimedia Foundation announced new AI partnerships on Wikipedia's twenty-fifth birthday. Amazon. Meta. Microsoft. Perplexity. Mistral AI. Enterprise revenue had more than doubled from the previous year.

Around the world, four people read the announcement and reached four different conclusions.

In Singapore, Robert Sim—the 2025 Wikimedian of the Year, with 79,000 edits to his name—considered what the partnerships meant for Wikipedia's future. "My hope for Wikipedia is simple," he had said in August, "that it continues to exist. For that to be possible, Wikipedia needs people." But new editor registrations were down 36% since 2016. If AI replaced readers, where would new editors come from?

In the United States, a volunteer from WikiProject AI Cleanup contemplated the irony. His team had spent months tagging thousands of articles for AI contamination. They had won battles—speedy deletion for AI artifacts, the death of the Foundation's summaries experiment, the

rejection of Wales's ChatGPT review proposal. And the Foundation had just signed five more AI deals.

On Meta-Wiki, a contributor posted questions about transparency: "There's good reason for it to be maximally transparent—the WMF is selling Wikipedia data, generated by its volunteer editors." No vote had authorized these partnerships. No community consensus had been sought.

In San Francisco, Foundation leadership celebrated a milestone. Enterprise was profitable for the first time. The "fair share" message was working. Wikipedia would survive the AI era by adapting to it.

Same announcement. Four incompatible interpretations. The fracturing that had been building for years had crystallized.

### Pro-Enterprise Pragmatists

The Pro-Enterprise position held institutional power.

Jimmy Wales articulated the public case: AI companies should pay their "fair share." They were already scraping Wikipedia. Better to formalize the relationship and collect revenue than to fight a losing battle.

"I'm very happy personally that AI models are training on Wikipedia data because it's human curated," Wales said in January 2026. "I wouldn't really want to use an AI that's trained only on X, you know, like a very angry AI."

Lane Becker, president of Wikimedia Enterprise, made the business case: "The AI companies are here, and they are particularly voracious. Wikipedia is a critical component of these tech companies' work that they need to figure out how to support financially."

The evidence they cited was real. Bot traffic strained servers—65% of the most resource-intensive traffic came from bots. Human traffic was declining—down 8% after improved bot detection. Enterprise provided revenue to offset costs. The 30% revenue cap, formalized in May 2025, protected nonprofit character.

| Fiscal Year | Enterprise Revenue | Growth |
|---|---|---|
| 2022-23 | $3.2M | — |
| 2023-24 | $3.4M | +6% |
| 2024-25 | $8.3M | +148% |

**Table 18.1.** Wikimedia Enterprise revenue growth

From this view, Wikipedia's adaptation to AI was working. The pragmatists were keeping the lights on.

### Copyleft Purists

The Copyleft Purist position held to principle.

Wikipedia was built on CC-BY-SA—a license promising that derivative works would remain free. Contributors had trusted that promise. The Foundation was its steward.

The April 2024 admission had exposed the betrayal: most customers were not complying with the license terms. The Foundation knew. It continued the partnerships regardless.

The June 2024 Open Future report confirmed what the purists feared: "Share Alike/CopyLeft licenses are largely ineffective when materials licensed under them are used to train AI models." The legal framework Stallman had helped create could not protect the commons in the AI era.

If the license could not be enforced, what remained of the promise? Contributors had created 65 million articles under terms the Foundation would not uphold.

"These are not charitable contributions—they are commercial transactions," one Meta-Wiki contributor wrote. "The WMF is selling Wikipedia data, generated by its volunteer editors."

### AI Skeptics

The AI Skeptics held the battleground.

WikiProject AI Cleanup, founded in December 2023, had become the front line. By October 2025, the project had flagged articles by the thousands—LLM-generated content with fabricated citations, impossible syntax, telltale phrases. The volunteers had developed detection methods, documented patterns, and built community consensus for action.

They had won battles:

**June 2025:** The Foundation launched AI-generated summaries to 10% of mobile users without community consultation. Discovery sparked outrage. The experiment died within 24 hours.

**August 2025:** Wales proposed ChatGPT for new article reviews. Editors tested it; it failed. Proposal rejected.

**August 2025:** G15 passed—speedy deletion for obvious AI artifacts. "I am closing this per SNOW," the closer wrote. So one-sided that extended discussion was unnecessary.

But the skeptics knew their victories were partial. The sophisticated AI content slipped through. The Princeton study found 4.36% of new articles contained significant AI-generated content. Editor accounts blocked for LLM use were doubling every hundred days.

The contrast with AI was fundamental: Wikipedia's value came from transparency, from every edit being traceable and reversible. AI models were black boxes—their training data invisible, their outputs unattributable.

### Volunteer Advocates

The Volunteer Advocates asked the simplest question: whose labor, whose profit?

A quarter million unpaid editors had created the content that trained every major AI system. Wikipedia was ChatGPT's dominant source—cited more often than any other. AI companies had built billion-dollar businesses on this labor.

The compensation: zero.

*The Conversation*, writing on Wikipedia's twenty-fifth anniversary, framed the question: "Are we witnessing a new 'tragedy of the commons,' where volunteered knowledge becomes raw material for systems" that commercialize it without consent?

The numbers were stark:

| Metric | Value |
|---|---|
| Active volunteer editors | 250,000 |
| Wikipedia share of ChatGPT citations | nearly half |
| Enterprise revenue (FY 2024-25) | $8.3M |
| Volunteer compensation | $0 |
| New editor registrations (2016-2025) | -36% |

**Table 18.2.** Volunteer labor and value extraction

Edit volume was up 37%—a "smaller club working harder," as one analysis put it. The people who created the content that powered AI were being asked to work harder while receiving nothing in return.

### Summaries Revolt

The AI summaries experiment revealed the fault lines.

On June 2, 2025, the Foundation launched "Simple Article Summaries"—AI-generated summaries using Cohere's Aya model, shown to 10% of mobile users. The community had not been consulted.

The response was immediate:

"Yuck," one editor wrote.

"I sincerely beg you not to test this, on mobile or anywhere else," wrote another. "This would do immediate and irreversible harm to our readers and to our reputation as a decently trustworthy and serious source."

Within 24 hours, the experiment was paused. No vote had been held. No consensus had been built. The Foundation had launched; the community had rejected; the feature had died.

The pattern was established: crisis and backlash replacing governance.

### Wales's ChatGPT Proposal

The founder's proposal showed where the factions stood.

In August 2025, Jimmy Wales suggested using ChatGPT to assist with Article for Creation reviews—the process by which new articles were evaluated for publication.

The skeptics responded with testing. They gave ChatGPT actual review tasks. The results were damning: ChatGPT could not reliably apply Wikipedia's notability guidelines. It made errors that experienced reviewers caught immediately.

The proposal was rejected. The founder who had created Wikipedia found his community would not follow him on AI.

### G15 Victory

G15 showed what the community could accomplish when united.

Before G15, removing AI-generated content required week-long deletion discussions. The process was slow; the content accumulated. WikiProject AI Cleanup advocated for faster action.

The RFC closed with overwhelming consensus. G15 allows immediate deletion for articles containing "communication intended for the user" (e.g., "Here is your Wikipedia article on...") or fabricated references such as non-existent papers and invalid ISBNs.

It was a victory. It was also, as one project member acknowledged, "a band-aid for the most egregious examples of AI-generated submissions, though the larger problem will continue."

G15 caught the obvious cases. The subtle contamination remained.

### Why Resolution Fails

The factions could not be reconciled because they disagreed about what Wikipedia was.

For the Pro-Enterprise faction, Wikipedia is infrastructure, and therefore it must be monetized to sustain. For Copyleft Purists, Wikipedia is a promise, and therefore the Foundation must enforce or betray that promise. For AI Skeptics, Wikipedia is human curation, and therefore they must exclude AI or die. For Volunteer Advocates, Wikipedia is their labor, and therefore the Foundation must compensate or explain.

These were not positions on the same spectrum. They were answers to different questions. No consensus mechanism could reconcile them because the factions did not share premises about what consensus should protect.

## Structural Barriers

Resolution failed for structural reasons:

**No common forum:** Foundation, affiliates, and editors had separate governance structures with no mechanism for binding joint decisions.

**Different stakes:** The Foundation faced institutional pressures—payroll, infrastructure, legal exposure—that individual editors did not. Individual editors faced labor extraction that the Foundation benefited from.

**Information asymmetry:** Enterprise contracts were confidential. Editors could not evaluate what the Foundation was actually promising AI companies.

**Legal complexity:** Licensing questions were genuinely unsettled. Neither side could claim clear legal high ground.

**Ideological homogeneity:** Foundation leadership moved through tech industry circles. FEC records showed employees donating exclusively to one party. Professional networks overlapped with AI company executives; vocabulary matched. No internal faction pushed for adversarial enforcement. When the question arose, no dissenting voice argued for it.

"This permanent imbalance creates constant tension, wastes time and energy, and ultimately weakens our entire movement," one community member observed.

### 2002 Precedent

The division had precedent.

In February 2002, when Bomis's financial pressures led to discussions of advertising, Spanish Wikipedia contributors announced they would fork. Edgar Enyedy led the exodus. Enciclopedia Libre launched as a Spanish-language alternative.

The fork failed. Wikipedia's network effects proved insurmountable. Enciclopedia Libre lacked critical mass. Most users did not share the ideological concerns.

But Enyedy's warning had been preserved: "Good luck with your WikiPAIDia."

### 2026 Echo

Twenty-four years later, the same fears had materialized through a different mechanism. In 2002, editors feared advertising would compromise independence; in 2026, Enterprise creates AI dependence. In 2002, editors feared commercial interests would introduce bias; in 2026, non-compliant customers are served anyway. In 2002, editors feared volunteer labor would be monetized without consent; in 2026, Enterprise revenue flows from volunteer content.

Enyedy was right about the threat. He was wrong about the timeline and mechanism. Commercialization came not through banner ads but through API access to companies training AI on volunteer-created content.

### Why Forks Are Harder Now

The impulse to break away persisted, but the options had narrowed.

**Network effects:** Wikipedia's dominance was more complete than in 2002. The brand was synonymous with "encyclopedia."

**Content volume:** Sixty million articles across languages could not be easily replicated.

**AI dependence:** Any fork would also face AI extraction. The problem was structural, not institutional.

**Coordination costs:** Organizing hundreds of thousands of editors was harder than organizing a few dozen Spanish contributors.

Forking solved nothing. But staying meant accepting capture.

### What Changed

The community that celebrated Wikipedia's twenty-fifth anniversary was not the community that launched the encyclopedia in January 2001.

The founding dream—that radical openness could produce reliable knowledge while keeping it free forever—had collided with a world that extracted value from openness at industrial scale.

The factions represented different responses to this appropriation: Pro-Enterprise chose accommodation as survival, Copyleft Purists demanded enforcement to reverse the taking, AI Skeptics practiced resistance to slow it, and Volunteer Advocates sought recognition that it had happened at all.

None was wholly wrong. None could convince the others. The conflict was not about policy but about identity—and about whether the enclosure of the commons could be undone, resisted, accommodated, or only mourned.

### Return to Opening

The four editors who read the January 15 announcement reached four conclusions.

Robert Sim saw an encyclopedia that needed people—and people were leaving.

The WikiProject AI Cleanup volunteer saw victories being undermined by the institution he served.

The Meta-Wiki contributor saw volunteer labor being sold without consent.

The Foundation saw a sustainable future through adaptation.

They could not all be right. They could not all be wrong. They could not agree on how to decide.

### Bridge Forward

Part V traces where this rupture leads.

Chapter 19 returns to Stallman's original vision and its forty-year journey to this endpoint. Chapter 20 examines what the Foundation could have done—the legal case that existed and the choice not to pursue it. Chapter 21 confronts the human cost: what happens to the people who built Wikipedia when the institution that governs it no longer protects them. Chapter 22 projects possible futures.

The civil war is not coming. The civil war is already here, fought through policy RFCs and funding appeals and quiet decisions to contribute less.

Wikipedia at twenty-five: the encyclopedia anyone could edit, that trained the AI that may make editing obsolete, governed by a foundation that sells access to the AI companies, built by volunteers who were never asked.

Four factions. Four visions. No mechanism to choose between them.

### End of Part IV

Part IV has traced the AI reckoning.

Chapter 13 documented the extraction: Wikipedia as foundational training data, the bedrock of AI factual knowledge.

Chapter 14 documented the pivot: Enterprise selling to AI companies, OpenAI paying nothing despite using Wikipedia most.

Chapter 15 documented the capture: the Foundation's April 2024 admission that customers were "not compliant," and the choice to keep selling anyway.

Chapter 16 documented the visible fork—Wikipedia copied with neo-Nazi contamination but CC-BY-SA attribution intact—while the Foundation condemned it for content failures and partnered with invisible extractors whose structural violations were worse.

Chapter 17 documented the loop: AI training on AI, contamination spreading through systems, human curation laundered until its origins disappeared.

This chapter documented the split: four factions with irreconcilable visions, no mechanism to resolve the conflict.

Part V asks what remains. Stallman's prophecy. The case for enforcement. The volunteer question. The scenarios ahead.

The reckoning is complete. What follows is the assessment.

# Part V

# What Remains

CHAPTER NINETEEN

# The Prophecy

> I started crying right there in the equipment room. Seeing the machine there, dead, with nobody left to fix it, it all drove home how completely my community had been destroyed.
>
> — Richard Stallman, recalling the PDP-10's end

On March 16, 1982—his twenty-ninth birthday—Richard Stallman received a message that would reshape the future of software freedom.

The MIT AI Lab had been Stallman's home since he arrived at eighteen in 1971. The community of programmers who worked there shared everything: code, improvements, ideas. When one person fixed a bug, everyone benefited. When someone wrote a new feature, the feature belonged to all. The Lisp Machine code they maintained together was not property—it was infrastructure for collective work.

That morning, Symbolics executives ended the gentleman's agreement. The company, spun off from the lab to commercialize Lisp Machine technology, announced it would no longer share improvements to the shared codebase. What the community built together, Symbolics would take. What Symbolics improved, the community would never see.

Stallman's response was immediate. He disconnected Symbolics' microwave link to MIT. He declared war. "The AI Lab was a neutral country,

like Belgium in World War I," he later explained. "If Germany invades Belgium, Belgium declares war on Germany."

For two years, he worked alone, cloning every feature Symbolics developed, ensuring the community's version remained competitive. When he suspected Symbolics had installed monitoring software to track his work, he rewrote from scratch to eliminate any possibility of taint.

The war ended not in victory but in dissolution. The hackers drifted away to industry jobs. The community fragmented. One evening, Stallman returned to the equipment room and found the PDP-10—the machine around which the community had formed—dark and silent. The grief of that moment would drive everything that followed.

From that grief came a weapon.

## GNU Manifesto

On September 27, 1983, Stallman announced the GNU Project in a Usenet message. On January 5, 1984, he resigned from MIT to avoid any intellectual property conflicts. In March 1985, he published the GNU Manifesto.

The manifesto contained the principle that would define his legacy: "Proprietary modifications will not be allowed."

Stallman designed copyleft as a legal mechanism to ensure that no one could ever again take shared work and refuse to share back. Use my code, he said, but share your improvements. The license would spread like a virus, ensuring that free software begat free software.

The punishment for violation was explicit. "If programmers deserve to be rewarded for creating innovative programs," he wrote, "by the same token they deserve to be punished if they restrict the use of these programs."

The weapon was designed for a specific wound: the Symbolics betrayal. Copyleft existed to prevent commercial entities from taking community-built resources and enclosing them in proprietary systems.

## Extension to Knowledge

In 1999, Stallman delivered a speech at SIGCSE proposing "The Free Universal Encyclopedia and Learning Resource." The vision: copyleft applied to knowledge.

"The World Wide Web has the potential to develop into a universal encyclopedia covering all areas of knowledge," he said.

The GNU Free Documentation License, released in 2000, extended copyleft from code to text. When Wikipedia launched in January 2001, it adopted the GFDL. The promise Stallman had made about software—use, but share—now applied to knowledge.

A quarter million volunteers would contribute under that promise. Their work would remain free forever. Derivative works would carry the same terms. The license would propagate.

## Symbolics and the AI Companies

Forty years separate the Symbolics ultimatum from ChatGPT citing Wikipedia in nearly half its responses. The pattern that emerged in that interval is precise.

In 1982, the commons was Lisp Machine code; in 2024, it is Wikipedia's 6.7 million articles. In 1982, the contributors were MIT hackers sharing freely; in 2024, they are hundreds of thousands of volunteer editors. In 1982, the promise was a gentleman's agreement to share improvements; in 2024, it is CC-BY-SA ShareAlike. In 1982, the betrayal was Symbolics taking code and refusing to share improvements; in 2024, it is AI models locked in proprietary systems. In 1982, the scale was one company and one codebase; in 2024, it is a trillion-dollar industry and all human knowledge. In 1982, the response was Stallman fighting back alone; in 2024, WMF signed Enterprise deals.

The pattern is identical. The response is opposite.

The difference was not circumstances but disposition. Stallman viewed corporations as adversaries by default. He assumed betrayal was coming

and built defenses accordingly. WMF leadership moved through tech industry circles—shared conferences, matching politics, a common vocabulary of ”access” and ”democratizing knowledge.” When AI companies came to extract, they found not adversaries but colleagues.

When Symbolics betrayed the community, Stallman disconnected their link, declared war, and spent two years ensuring the community could survive without them. He then spent the rest of his career building legal infrastructure to prevent such betrayals from ever happening again.

When AI companies betrayed the Wikipedia license, the Wikimedia Foundation did not disconnect anything. It did not declare war. It created Enterprise—a commercial arm to sell access to the companies whose compliance was questionable.

### Enforcement Never Happened

The twist is not that copyleft failed. The twist is that copyleft was never tested.

The legal arguments for enforcement exist:

**Memorization:** Nicholas Carlini and colleagues at Google DeepMind proved that language models reproduce training data verbatim. Over five percent of ChatGPT output consists of direct copies from its training corpus. If models copy, the ”transformative learning” defense collapses.

**Market Harm:** Wikipedia traffic declined eight percent year-over-year. Organic search traffic fell twenty-six percent over three years. Zero-click searches increased from fifty-six to sixty-nine percent. The market effect factor—”undoubtedly the single most important element of fair use”—weighs against AI companies.

**Attribution Violations:** CC-BY-SA requires attribution. ChatGPT draws on Wikipedia more than any other source, but only one percent of users click citations. Attribution exists in theory; it fails in practice.

**ShareAlike Violations:** CC-BY-SA requires derivative works to carry the same license. Model weights are proprietary. AI outputs are propri-

etary. If these are derivatives, they violate the license. If they are not derivatives, the license cannot protect the content—which is its own damning conclusion.

The case exists. The precedents are being established—*Thomson Reuters v. ROSS Intelligence* found AI training was not fair use when the product competed in the same market. The legal terrain favors enforcement.

And the Wikimedia Foundation, knowing its AI customers operated in legal gray zones, filed no lawsuit. Made no public demand. Published no enforcement position.

It took the money.

In April 2024, the Wikimedia Foundation acknowledged internally that most AI customers were likely violating license terms. By that date, the Carlini memorization study was published. Wikipedia traffic decline was measurable. The legal case for enforcement *existed*.

That same month, the Foundation's response was internal deliberation, not enforcement demand. On that date, in that moment, the institution that should have defended the commons chose instead to assess commercial relationships.

### Pretend Intelligence

On January 23, 2026, Richard Stallman spoke at Georgia Tech.

At seventy-two, he remained uncompromising. He refused the term "artificial intelligence" for large language models, insisting on a different name: "pretend intelligence."

"I decided not to be part of their campaign to convince people that those things are intelligent," he said. "Every time you call them AI, you are endorsing the claim that they are intelligence and they're not. So, let's refuse to do that."

He demanded that AI-generated content be labeled. He argued that harmful AI uses should be regulated by law, not licensing. He remained,

as he had always been, precise about language and unyielding about principle.

But he did not speak about copyleft enforcement. He did not speak about the Wikimedia Foundation. He did not address what had happened to the license he had helped inspire.

Perhaps he knew the battle was lost when the defenders became dependent on the invaders.

### FSF Response

The Free Software Foundation, which Stallman founded in 1985, published criteria for "free machine learning applications" in October 2024. The criteria required that training data and scripts grant the four freedoms.

At FOSDEM 2025, the FSF expanded: "We cannot say a ML application 'is free' unless all its training data and the related scripts for processing it respect all users, following the four freedoms."

The criteria were admirable. They were also late. By October 2024, GPT-4 had already trained on Wikipedia. ChatGPT had over one hundred million users. Model weights were proprietary. The extraction had happened.

The FSF addressed what AI applications *should* be. It did not address enforcement against companies that had already extracted.

### Legal Assessment

In June 2024, the Institute for Information Law published its analysis of ShareAlike licenses and AI training. The finding was direct: "Share Alike/CopyLeft licenses are largely ineffective when materials licensed under them are used to train AI models."

The reasons were structural. Key license concepts—"adapted material," "technical modification"—were "largely absent in AI training workflows." Tracing from output to training source was "practically impossible." The

license was designed for a world where derivatives were recognizable. Model weights were not recognizable as derivatives.

But "largely ineffective" was not "completely unenforceable." The study identified strategic paths forward—rights reservation, opt-out mechanisms, ShareAlike extensions. The Wikimedia Foundation pursued none of them.

The question that hangs over all three perspectives: Would the law have worked if someone had tested it? We cannot know, because no one tested it. The defenders chose not to defend.

## Forty-Year Circle

In 1982, Stallman stood before a dead machine and asked: how do I ensure this never happens again? The answer wasn't just copyleft. It was something harder: assuming that institutions—even ones founded to protect freedom—might become corrupted. Copyleft was designed to work *even if its stewards could not be trusted.*

Forty-two years later, that assumption proved prescient. Not because the wall failed, but because the guards chose not to stand.

The answer was copyleft. Derivative works would carry the same freedom. The license would propagate through every reuse.

Forty-two years later, on March 16, 2024—Stallman's seventy-first birthday—the Wikimedia Foundation had already acknowledged what was happening. Enterprise customers were likely violating license terms. It kept signing deals anyway.

The birthday symmetry is coincidental but fitting. The wound and the weapon share an anniversary. The weapon was inherited but not wielded.

## Guards Were Aligned

The prophecy Stallman fulfilled was not that copyleft would be circumvented by clever lawyers. It was that the steward entrusted to defend it would become aligned with the violators—first culturally, then financially.

The wall Stallman built still stands. The legal mechanism exists. The license remains enforceable in principle. But enforcement requires an enforcer with resources, will, and independence.

The Wikimedia Foundation had resources—$185 million in annual revenue, $400 million in reserves, a $144 million endowment. It lacked will. It was not independent—leadership drew from overlapping professional circles, gave to the same political party, spoke the language of "access" and "democratization" that tech executives favored.

The wall was never breached. The guards were not merely bought. They had already chosen sides. Affinity preceded dependency.

## What the Silence Means

That January, Stallman addressed the same themes at Georgia Tech. Precision about language. Transparency demands. Refusal to validate corporate framing.

He said nothing about Wikipedia. Nothing about copyleft enforcement. Nothing about the arc from Symbolics to ChatGPT.

The silence may be intentional. Stallman has always chosen his battles carefully. Or perhaps the silence acknowledges what cannot be said: that the mechanism he designed worked, but the institution that inherited it chose revenue over principle.

Perhaps he recognized something he had written about decades earlier: that licensing requires institutions to have the courage and independence to enforce them. License alone cannot protect knowledge when the entities designed to defend it are financially dependent on those who violate it.

The copyleft promise—use, but share—required enforcement. The enforcement required institutions with independence. The Wikimedia Foundation had none.

Chapter 20 examines the legal case that existed—the arguments, the precedents, the evidence. The case for enforcement was real. The choice not to pursue it was also real.

The question is not whether the license could have protected the commons. The question is why the defenders chose not to try.

CHAPTER TWENTY

# The Case That Wasn't Made

> The effect of the use upon the potential market for or value of the copyrighted work is undoubtedly the single most important element of fair use.
>
> — U.S. Copyright Office, May 2025

In November 2023, a team of researchers made a discovery that should have ended the debate about AI training and copyright.

Using carefully designed prompts, Nicholas Carlini and colleagues at Google DeepMind demonstrated that over five percent of ChatGPT's output consisted of direct verbatim copies from its training data—fifty-token sequences reproduced word for word.

Five percent. Not patterns. Not learned representations. Copies.

The paper included the methodology. The findings were reproducible. The implications were clear: if AI outputs reproduce training data verbatim, that is copying under any legal standard. The "transformative learning" defense—the claim that models learn patterns rather than copy content—collapsed against the evidence.

The research was public. The evidence was undeniable. The grounds for action were clear.

And the Wikimedia Foundation—whose content powered nearly half of ChatGPT's factual responses—filed no lawsuit. Made no public statement. Signed more Enterprise deals.

The evidence of copying existed. The grounds for action were clear. The steward chose not to act.

## Memorization Is Copying

The AI industry's defense rested on a narrative: models "learn" from data rather than copy it. Training is transformation. Raw text becomes statistical patterns. Patterns become capabilities. Capabilities generate new text. The original sources are digested, not reproduced.

Carlini's research proved this narrative false.

The 2023 study demonstrated "divergence attacks"—prompting techniques that cause models to emit training data at 150 times their normal rate. The findings were systematic: memorization increases logarithmically with model size, examples seen multiple times are memorized more, and longer prompts extract more memorized content.

Corroborating studies accumulated:

| Study | Finding | Source |
|---|---|---|
| Patronus AI | GPT-4 produces copyrighted content 44% of tested prompts | 2024 |
| Cornell | ChatGPT retrieved 72 of 240 poems verbatim | Jan 2024 |
| Ahmed et al. | Near-perfect Harry Potter reconstruction | Jan 2026 |

**Table 20.1.** AI memorization studies

The legal significance was direct. Copyright infringement requires proving copying occurred. Verbatim reproduction is copying under any legal standard. If models reproduce inputs, the transformative defense fails.

Wikipedia, trained at 3.4 epochs in GPT-3—the highest quality weighting of any data source—was prime memorization territory. The quality that made it valuable for training also made it likely to be reproduced.

### Factor One: Purpose and Character

The AI industry argued that training was "highly transformative." Recent rulings complicated this claim.

In *Andy Warhol Foundation v. Goldsmith* (2023), the Supreme Court narrowed transformativeness: "A defendant's use of a plaintiff's work is not necessarily transformative just because that use contains a new meaning or message. Something more is required—the defendant's use must have a different purpose than the plaintiff's."

The application to AI was direct. Wikipedia's purpose: answering informational queries. ChatGPT's purpose: answering informational queries. Same purpose meant weaker transformativeness claims.

The Copyright Office's May 2025 report reinforced the distinction: "Training a model for research may be highly transformative, while training models to generate outputs that compete with or closely resemble copyrighted works is not."

ChatGPT competed directly with Wikipedia. Both answered the same questions. Both served the same users. The commercial nature of the use—OpenAI's $6.7 billion in annual revenue—weighed against fair use.

### Factor Two: Nature of the Work

The argument that Wikipedia was "just facts" failed against the evidence.

Wikipedia articles were not databases of raw information. They were carefully written, edited, and curated works of authorship. Editors chose what to include, how to phrase it, what sources to cite. The selection and arrangement was protected expression.

Facts themselves were not copyrightable. But the expression of facts was. Wikipedia's value came from the expression—the curated, verified, neutrally written summaries that made raw information accessible.

### Factor Three: Amount Used

AI companies used all of Wikipedia. Not excerpts. Not samples. The entire corpus.

GPT-3's training included approximately three billion tokens from English Wikipedia. The model saw Wikipedia 3.4 times—more than any other source. LLaMA used Wikipedia at 4.5% of its training data.

Using the entire corpus of a work weighed against fair use. Selective quotation might qualify for protection. Total extraction did not.

### Factor Four: Market Effect

Market effect—the factor highlighted in the chapter epigraph—weighed heavily against AI companies. The harm to Wikipedia was documented:

| Metric | Value |
|---|---|
| Human traffic decline (after bot detection) | -8% YoY |
| Organic search traffic decline (3-year) | -26% |
| AI Overview click-through decline | -61% |
| Zero-click searches (increase) | 56% ⟶ 69% |

**Table 20.2.** Wikipedia market effect from AI

Wikipedia was being substituted by AI products built on its content. Users who once visited Wikipedia now received Wikipedia-derived answers through ChatGPT and Google AI Overviews. The traffic that sustained the volunteer feedback loop was being captured by commercial systems.

In *Kadrey v. Meta* (June 2025), Judge Chhabria warned: "It seems likely that market dilution will often cause plaintiffs to decisively win the fourth factor—and thus win the fair use question overall."

Wikipedia had market harm evidence. The Foundation chose not to use it.

### Precedents Support Enforcement

The legal terrain was not settled in AI companies' favor.

***Thomson Reuters v. ROSS Intelligence* (February 2025):** The first AI training case found against fair use. ROSS built a legal search tool to compete with Westlaw using Westlaw's materials. The court held: commercial purpose combined with market substitution defeated fair use.

The parallel to Wikipedia was direct. ChatGPT built an answer engine to compete with Wikipedia using Wikipedia's materials.

***Bartz v. Anthropic* (June 2025):** Judge Alsup found training "spectacularly transformative"—but ruled that using pirated copies defeated the fair use defense. Anthropic settled class-wide in August 2025, facing potential exposure in the hundreds of billions.

***Kadrey v. Meta* (June 2025):** Training found fair use, but the judge emphasized that plaintiffs presented no market harm evidence. Wikipedia had market harm evidence.

***NYT v. OpenAI*:** The case that tested these theories. Judge Stein allowed most copyright claims to proceed. Trial expected after April 2026. The legal terrain was actively contested.

### CC-BY-SA Adds Claims

Even if fair use covered training—a contested proposition—CC-BY-SA violations were systematic.

CC-BY-SA required attribution to Wikipedia, but AI companies rarely provided it in outputs. It required attribution to editors, but this was never provided. It required ShareAlike for derivatives under the same license,

but model weights were proprietary. It required attribution "reasonable to the medium," but no attribution mechanism existed.

The Foundation had independent claims under the license. Attribution was not reaching users—one percent clicked citations. ShareAlike was not being honored—model weights and outputs were proprietary.

All four fair use factors cut against AI companies. CC-BY-SA violations were systematic. Recent rulings favored content owners.

The Foundation chose Enterprise revenue instead.

### Pragmatist Defense

From the Foundation's perspective, litigation was uncertain and expensive.

AI companies would train on Wikipedia regardless of legal threats. The content was freely available. The training had already happened. Suing would take years, cost millions, and might fail.

Better to engage through Enterprise and capture some value. The 30% revenue cap protected nonprofit character. The relationships with tech partners served Wikipedia's mission.

"Wikipedia is a critical component of these tech companies' work that they need to figure out how to support financially," Lane Becker said. The argument was practical: if extraction was inevitable, at least collect something.

### Capture Critique

From a governance perspective, the pragmatist defense explained WMF's reasoning while missing the structural incentives that shaped it.

The pattern matched regulatory capture. AI companies needed Wikipedia for content. WMF needed revenue for sustainability. Enterprise created dependency—$8.3 million and growing. Dependency prevented enforcement, because you cannot sue paying customers. Non-enforcement enabled both invisible extraction (ChatGPT with no attribution) and visible forking (Grokipedia with attribution but problematic sources). The

Foundation condemned the visible fork while partnering with the invisible extractors.

The capture did not require corruption or bad faith. It required only that institutional interests align with extractors' interests—and they did. Once Enterprise revenue mattered, enforcement became structurally implausible.

### NYT Comparison

The New York Times faced the same legal uncertainties. The same "training is transformative" arguments. The same industry pressure.

The Times sued.

The contrast was stark. The New York Times had high content training value; Wikipedia's was even higher—the factual backbone of ChatGPT's knowledge. Both faced the same legal uncertainty. The Times responded with a lawsuit in December 2023; WMF responded with Enterprise in 2021. As of 2026, the Times had a trial pending; WMF was taking AI company money. In April 2024, the Times had no admission to make; WMF acknowledged most customers were not compliant.

The Times chose enforcement. WMF chose accommodation. The difference was not legal terrain—it was institutional positioning.

### Enforcement Alternatives

The missed opportunities were specific. In May 2020, when GPT-3 disclosed Wikipedia training, the Foundation could have demanded attribution compliance. In November 2022, when ChatGPT launched, it could have demanded licensing. In December 2023, when the NYT sued OpenAI, it could have joined or filed a parallel suit. In April 2024, when 273 Ventures asked about compliance, it could have published its position and demanded compliance. In February 2025, when Thomson Reuters won, the precedent strengthened its case.

The Foundation could have filed suit under theories the Times was testing. It could have published a public enforcement position. It could have required CC-BY-SA compliance as an Enterprise condition. It could have demanded attribution in AI outputs.

It did none of these things.

## What the Evidence Proves

The evidence points in one direction. Memorization proves copying. Fair use factors weigh against commercial AI. Traffic data shows substitution. Attribution requirements go unmet.

And the Wikimedia Foundation—the institution entrusted with protecting the commons—chose not to act.

## The Thesis

Copyleft did not fail because the license was inadequate.

Copyleft failed because the institution's immune response was disabled from within.

Political homogeneity eliminated internal opposition. FEC records show WMF employees donated exclusively to one party across 292 contributor-cycles—a pattern matching Big Tech's executive class. Leadership drew from overlapping professional circles: World Economic Forum panels, Council on Foreign Relations memberships, McKinsey credentials. No board member brought center-right institutional affiliations. No internal faction questioned whether "partners" might be adversaries in licensing terms.

Cultural affinity preceded financial dependency. When AI companies came to extract volunteer labor, they found people who spoke their language. Enterprise contracts formalized relationships that already felt natural.

You cannot sue your paying customers. The difficulty compounds when you share their worldview.

Stallman's license worked as designed. Enforcement required adversarial instincts that monoculture had bred out of the institution.

### Cost of Non-Enforcement

If enforcement was possible—and the evidence suggests it was—then non-enforcement was a choice. Not passive drift. Not institutional inertia. A choice made at specific moments by specific people who decided that Enterprise revenue mattered more than licensing integrity.

And that choice has a cost.

A quarter million volunteers contributed to Wikipedia under the promise of CC-BY-SA. Their work would stay free. Derivative works would be shared. Attribution would follow their words.

Those volunteers were not consulted when WMF signed Enterprise deals. They were not asked when WMF admitted non-compliance. They were not represented when the Foundation chose revenue over enforcement.

The viable path to enforcement existed. The Foundation chose another road.

Chapter 21 asks what happens to the people who built the commons that was captured.

### Final Image

The memorization data was public. The traffic decline was measurable. The court rulings were trending toward content owners.

And the institution that should have wielded these tools chose accommodation over confrontation.

The steward chose its customers over its contributors.

CHAPTER TWENTY-ONE

# The Volunteer Question

> My hope for Wikipedia is simple: that it continues to exist. For that to be possible, Wikipedia needs people.
>
> — Robert Sim, 2025 Wikimedian of the Year

On January 15, 2026, the Wikimedia Foundation celebrated Wikipedia's twenty-fifth anniversary with press releases, partnership announcements, and retrospective coverage.

In San Francisco, Foundation staff marked the milestone. Enterprise revenue had more than doubled. Amazon, Meta, Microsoft, Perplexity, and Mistral AI had signed as partners. The free encyclopedia had survived a quarter century. The AI transition was being managed.

In Singapore, Robert Sim opened his laptop and read the news.

Sim was the 2025 Wikimedian of the Year. Over twenty years and 79,000 edits, he had documented Southeast Asian culture, history, and politics for the English Wikipedia. He had watched the project grow from curiosity to institution. He had stayed when others left.

His statement upon receiving the award had been simple: "Knowledge should belong to everyone. This is why I started editing Wikipedia over 20 years ago and why I continue today. My hope for Wikipedia is simple: that it continues to exist. For that to be possible, Wikipedia needs people."

The hope was genuine. The condition was uncertain. The question that haunted every longtime editor hung over the celebration: why would anyone volunteer their labor when that labor trained competitors?

## Traditional Model

For twenty-five years, Wikipedia sustained itself through a feedback loop. But understand what that loop actually was: it was a commons-making mechanism.

A reader visits Wikipedia. They notice an error—a misspelling, a missing fact, an outdated figure. They click "edit." They fix the problem. They see their improvement go live, verified, read by millions. They feel the satisfaction of contribution to something collectively owned.

That moment—seeing your work become part of a shared resource—is what makes a commons. Not the infrastructure. Not the license. Not the steward. The commons is the quarter million people who volunteer because they believe they are building something that belongs to everyone equally.

When that belief breaks, the commons dies.

This loop converted readers into editors. The simplicity of "anyone can edit" was both ideology and recruitment mechanism. Wikipedia grew from a handful of contributors in January 2001 to hundreds of thousands of active editors by its twenty-fifth anniversary.

The motivations that sustained this loop were well-documented. Academic research identified intrinsic factors (enjoyment, learning), social factors (community, identity), and values (ideology, public service). Recognition increased retention by twenty percent. Volunteers continued when they felt part of "something bigger than themselves."

The loop worked because contributors could see their impact. They wrote; people read; the effect was visible. The feedback was immediate and tangible.

### Henner's Warning

Christophe Henner, former chair of the Foundation's Board, published a "wake-up call" on the twenty-fifth anniversary. His assessment was blunt:

"The feedback loop that has sustained volunteer motivation for 25 years is simple: I write, people read, I can see the impact. Break that loop and you break the engine for some contributors."

His diagnosis was stark. The question was whether anything could repair it.

### Long Decline

The volunteer crisis began before AI. Editor counts peaked in 2007 at roughly 51,000 monthly active editors (those making five or more edits). By 2014, the count had fallen to 30,000-35,000 and stabilized.

Research identified the causes: algorithmic gatekeeping by bots, rule complexity that had grown from fifty policies to over a thousand, hostile culture that "bit" newcomers, and a mobile gap that made editing from phones difficult.

Aaron Halfaker's 2013 study documented the paradox: "The very systems we built to maintain quality are now preventing growth." The filters that kept vandalism out also kept new editors out.

The decline stabilized but never reversed. Wikipedia became, in Henner's phrase, "a smaller club working harder."

### AI Acceleration

Then came ChatGPT.

In April 2025, ChatGPT overtook Wikipedia in monthly visits. The system that trained on Wikipedia now competed with it for user attention. The irony was precise: Wikipedia provided nearly half of ChatGPT's cited sources, and now ChatGPT was replacing it.

The traffic data confirmed the shift:

| Metric | Value |
|---|---|
| Human traffic decline (after bot detection) | -8% YoY |
| Organic search traffic decline (3-year) | -26% |
| Zero-click searches (2024 → 2025) | 56% → 69% |
| New registrations decline (2016-2025) | -36% |

**Table 21.1.** Wikipedia traffic decline metrics

The feedback loop was breaking at both ends. Fewer readers meant fewer potential editors. AI answers meant users never reached Wikipedia at all.

### Broken Contract

The volunteer crisis was not just about declining numbers. It was about a promise broken.

Wikipedia volunteers contributed under an implicit promise. Their labor would remain free and accessible. It would be governed by the community. The copyleft license would ensure that anyone building on their work would share improvements back. They were building a commons.

In 2024 and 2025, that promise visibly shattered.

When ChatGPT launched, volunteers saw their words reproduced in a locked proprietary system. They saw users asking ChatGPT instead of visiting Wikipedia.

When the Foundation signed Enterprise deals, volunteers discovered they had not been consulted. The institution they had built had become dependent on the companies extracting their labor.

When the Foundation experimented with AI summaries in June 2025, volunteers killed it in 24 hours. When Wales proposed ChatGPT for article reviews, editors rejected it. They were not opposed to AI in principle. They were opposed to being extracted without consent.

The volunteer question was not "why are people leaving?" It was "why would anyone stay?"

The motivations that sustained contribution—impact, community, building something lasting—were being hollowed out. AI removed the sense of impact (no one sees your work when chatbots answer). Enterprise undermined community (the Foundation takes money without consultation). The harvesting betrayed the promise of building something lasting (your labor trains your replacement).

### Existential Threat View

Ian Ramjohn at Wiki Education was direct: "Definitely LLMs are an existential threat to Wikipedia."

*Scientific American* elaborated in its anniversary coverage: "Fewer visitors mean fewer new editors for Wikipedia, and less frequent visits mean slower correction of errors... If fewer people visit and fewer people edit, the system that made Wikipedia self-correcting—and unusually resilient—could weaken."

The academic research supported this concern. Wagner and Jiang, writing in the *Journal of the Association for Information Science and Technology*, predicted a "vicious cycle":

"Contributors who reduce their contribution efforts as AI pervades the platform will leave Wikipedia increasingly dependent on additional AI activity... [causing] a vicious cycle leading to a staling of the content and diminishing value of this venerable knowledge resource."

The loop was closing: volunteers withdraw, Wikipedia becomes more AI-dependent, quality declines, the resource degrades.

### Community Defense View

WikiProject AI Cleanup represented the community's defensive response.

Founded in December 2023, the project had tagged nearly 3,000 articles for AI contamination by October 2025. They had developed detection methods, documented patterns, and built consensus for action.

Their victories were real: G15 (speedy deletion for obvious AI artifacts), the death of the AI summaries experiment, the rejection of Wales's ChatGPT proposal. The community could still defend itself.

"Wikipedia's brand is reliability, traceability of changes, and 'anyone can fix it,'" editor Bawolff wrote. "AI is the opposite of these things."

Ilyas Lebleu of WikiProject AI Cleanup echoed the concern: "When readers come to Wikipedia, they expect something that was written by volunteers, that was checked by volunteers. They don't expect something that was written by an AI."

The defense was principled. Whether it was sufficient remained uncertain.

### Continuing Contributors

Some volunteers stayed despite everything.

Robert Sim articulated their reasoning through his actions: twenty years and 79,000 edits. The ideological commitment to free knowledge survived the harvesting. The community bonds remained meaningful. The work still mattered.

These were the editors maintaining Wikipedia at twenty-five: the dedicated core, motivated by values that commercial harvesting could not entirely undermine. They continued because the alternative—abandoning the collective work—was worse.

But the question remained: could values alone sustain contribution when every other motivation was being systematically eroded?

### Synthesis

These perspectives were not contradictory. They were different responses to the same crisis.

The threat was real. The community was defending itself. But whether defense would be enough—and whether enough volunteers would stay to mount it—remained the open question.

### What Changed

The reader now understands the full scope of the volunteer crisis.

It is not just declining numbers. It is a betrayed promise. The motivations that sustained contribution are being eroded by AI harvesting that the Foundation has monetized without volunteer consent.

### The Numbers

| Metric | Then | Now |
|---|---|---|
| Active editors (2007 vs 2025) | 51,000 | 35,000 |
| New registrations (2016 vs 2025) | 317,000/month | 202,000/month |
| First-time editor return rate | — | 20% |
| Editors reaching 100+ edits | — | <3% |
| Enterprise revenue (FY 2024-25) | — | $8.3M |
| Volunteer share of Enterprise | — | $0 |

**Table 21.2.** Volunteer metrics: then and now

The volunteers who create the content that powers AI receive nothing from the commercial relationships their labor enables.

### Return to Robert Sim

Robert Sim's hope for Wikipedia—"that it continues to exist"—now carries its full weight. The reader understands what is required for that hope to be realized: people.

The reader also understands why people might not stay.

## The Question

The volunteer question asked was not really about Wikipedia. It was about collective work itself: What happens when the people who *are* the shared resource realize it is being harvested?

Why would anyone continue building something being taken from them? The question is not about individual motivation. It is about the basic structure of shared governance: if the people who create it cannot control what happens to it, it is not a public good. It is exploitation with a free license.

The volunteers who maintain Wikipedia have no answer to the question this chapter asks. Neither does the Foundation. Neither does anyone.

Why contribute when your labor trains competitors? Why write when AI answers? Why build when others extract?

For twenty-five years, volunteers contributed because they believed they were building something that would remain free, accessible, and theirs. The AI era has revealed that belief may have been naive. Their labor is valuable—valuable enough for billion-dollar companies to train on, valuable enough for the Foundation to monetize. But not valuable enough, apparently, to compensate or even consult.

Wikipedia needs people. But people need reasons. The feedback loop that sustained the encyclopedia for a quarter century is breaking.

Chapter 22 asks what futures remain possible when the people who built Wikipedia have every reason to leave.

CHAPTER TWENTY-TWO

# Four Futures

> Neither the state nor the market is uniformly successful in enabling individuals to sustain long-term, productive use of natural resource systems.
>
> — Elinor Ostrom, *Governing the Commons*

On January 15, 2026, Wikipedia celebrated its twenty-fifth anniversary. The celebrations were global: retrospectives in major publications, commemorative banners on the site, partnership announcements from the Foundation.

Four longtime editors, scattered across four time zones, read the coverage and thought about the future. Not just Wikipedia's future—the future of the commons itself.

The question beneath all four scenarios was not: "Will Wikipedia survive?" It was: "Can a commons survive capture?"

In California, a Pro-Enterprise pragmatist saw sustainability. AI companies were finally paying. The free-rider problem was being solved. Wikipedia would survive by adapting.

In Germany, a copyleft purist saw betrayal. The license that promised derivatives would stay free was being violated with Foundation complicity. The promise that brought volunteers was being broken.

In Australia, an AI skeptic saw existential threat. Traffic was declining. AI systems were citing corrupted sources. The brand that made Wikipedia trusted was being eroded.

In Singapore, a longtime contributor saw uncertainty. The work still mattered. The community still meant something. But for how long?

Same anniversary. Four visions of what comes next. None knew which would arrive.

## January 2026 Baseline

As the scenarios diverge, the starting conditions are documented:

| Factor | Current State |
|---|---|
| Enterprise revenue | $8.3 million (FY 2024-25) |
| AI customers | Amazon, Meta, Microsoft, Google, Perplexity, Mistral, Anthropic |
| Human traffic | Down 8% YoY |
| New registrations | Down 36% since 2016 |
| Community factions | Four-way split |
| Legal status | NYT v. OpenAI pending; Thomson Reuters v. ROSS on appeal |

**Table 22.1.** Wikipedia status as of January 2026

The factions documented in Chapter 18 represent different visions of the future. Pro-Enterprise voices say "AI is inevitable; get paid for access." Copyleft Purists see "enforcement or betrayal of the founding promise." AI Skeptics warn of "existential threat requiring coordinated response." Volunteer Advocates observe "our labor, their profit, no consent."

Each faction sees a different path forward. None commands consensus.

### Scenario One: Wikipedia Captured

The Enterprise path continues. Revenue grows from $8.3 million to $15 million to $50 million. Every major AI company pays for access. The Foundation reinvests. Individual access remains free.

Wikipedia becomes infrastructure for the AI industry—a public good supported by those who profit from it. Volunteers maintain content while commercial entities extract value. The copyleft ideal becomes a dead letter.

*Who benefits:* The Foundation achieves financial sustainability. AI companies gain legitimate access. Users receive AI answers based on verified content—at first.

*Who loses:* Volunteers see their labor absorbed without compensation or consultation. The ShareAlike promise dies. Future open projects learn that free content means free for harvesting.

*The dependency dynamic:* Once revenue reaches $50 million, the Foundation cannot risk losing it. When AI companies request changes to terms of service, data formats, or attribution requirements, the Foundation must weigh community values against commercial relationships. The entity that should advocate for the shared resource becomes the entity that negotiates its harvesting.

*What survives:* A database. *What dies:* The commons.

### Scenario Two: Wikipedia Fractures

A critical mass of editors concludes the Foundation has betrayed copyleft. They announce a fork, invoking CC-BY-SA's permission to copy and modify.

The fork implements strict attribution requirements. Technical measures block non-compliant scrapers. Some volunteers migrate. Most stay with the original Wikipedia. Two communities develop, each claiming to represent Wikipedia's true values.

*The precedent:* In 2010, when Oracle acquired OpenOffice, developers forked to create LibreOffice. Research shows the fork attracted "the long-

term and most active committers" and achieved "long-term sustainable community."

*Why Wikipedia is different:* Wikipedia is not software. Articles require maintenance across millions of pages. The volunteer base cannot sustain two encyclopedias. A fork would divide the already-declining contributor pool.

*The Spanish precedent:* In 2002, Spanish editors forked over commercialization fears, creating Enciclopedia Libre. That fork survives but never matched the original's scale. Edgar Enyedy's "WikiPAIDia" warning came true—but so did the network effects that make forking difficult.

*Who benefits:* Ideological purists preserve principles in at least one version.

*Who loses:* Both encyclopedias face divided resources and duplicated effort. Readers face confusion about which Wikipedia to trust. The movement's civil war consumes energy that could build knowledge.

*What survives:* Two partial projects. *What dies:* The unified commons.

### Scenario Three: Wikipedia Enforces

Courts rule against AI fair use. The *NYT v. OpenAI* trial proceeds in late 2026 and finds that training without license is infringement. The Thomson Reuters precedent is affirmed on appeal.

AI companies face a choice: license content or remove it from training. Publishers coalition demands compliance. The Wikimedia Foundation, pressured by community and courts, joins enforcement efforts.

*The legal foundation exists:* The Copyright Office's May 2025 report found that market effect is "undoubtedly the single most important element of fair use" and identified AI training as threatening "significant potential harm." Judge Chhabria warned that "in many circumstances, AI training might not qualify as fair use."

*Who benefits:* Content creators gain licensing revenue and attribution requirements. The copyleft ideal is enforced against AI. Volunteers see their labor finally acknowledged.

*Who loses:* AI companies face billions in licensing costs. Users receive degraded or more expensive AI outputs.

*The complications:* The Foundation acknowledged non-compliance in April 2024 but continued Enterprise deals. Courts may find the Foundation implicitly ratified AI training through years of accommodation. Enforcement now would require explaining that silence. Revenue relationships create conflicts of interest.

*What survives:* The commons and the institution that abandoned it. *What dies:* The Foundation's claim to be Wikipedia's true steward.

### Scenario Four: Wikipedia Declines

No dramatic resolution. The status quo continues until it doesn't.

Registrations continue declining. Traffic continues falling. ChatGPT maintains its lead in monthly visits. The "smaller club working harder" exhausts itself. Article quality degrades imperceptibly, then visibly.

AI systems trained on 2024 Wikipedia increasingly cite outdated or degraded content. The feedback loop closes: degraded AI cites degraded Wikipedia cites degraded AI.

*The degradation loop:* On January 24, 2026, *The Guardian* reported that ChatGPT was citing Grokipedia—an AI-generated encyclopedia that copied Wikipedia content and added neo-Nazi citations. The contamination was spreading.

*The brand erosion:* "Is 'I edit Wikipedia' something people say with pride or embarrassment?" Henner asked. "Contributing to open source on GitHub has cachet. Making TikToks has cachet. Editing Wikipedia? We've become the encyclopedia your teacher warned you about, not the movement you want to join."

*Who benefits:* No one. This is the tragedy of the commons realized.

*Who loses:* Everyone. Readers lose reliable information. AI loses quality training data. The knowledge commons degrades for all.

*What survives:* A degraded database. *What dies:* The community and its belief in collective ownership.

### One Root

The four scenarios diverge in trajectory but converge in origin.

Licensing did not fail because the law was poorly written. The legal case for enforcement existed. The memorization studies documented copying. The market harm was measured. The precedents were developing.

The license failed because the institution entrusted to enforce it was captured.

The Foundation had resources—hundreds of millions in revenue and reserves. It lacked will. It was not independent—it had become entangled with the companies whose compliance it should have demanded.

All four scenarios flow from this capture. The *Captured* scenario sees the pattern continue. The *Fractured* scenario sees the community split over the capture. The *Enforced* scenario sees external pressure force action the Foundation avoided. The *Declined* scenario sees the capture's consequences unfold without intervention.

The scenarios are not alternatives. They are simultaneous trajectories whose relative strength will determine the future.

### Successful Commons

Elinor Ostrom won the Nobel Prize for demonstrating that communities can govern common resources without state or market intervention. Her eight design principles for successful commons governance explained why some commons endured while others collapsed.

Wikipedia met most of these principles: clearly defined boundaries (who can edit, what content qualifies), rules appropriate to local conditions (consensus-based policies), collective-choice arrangements (RFCs,

discussions, votes), monitoring (Recent Changes patrol, bots, oversight), graduated sanctions (warnings, blocks, bans), conflict-resolution mechanisms (dispute resolution, ArbCom), minimal recognition of rights to organize (self-governance), and nested enterprises (Foundation, affiliates, projects).

The governance that worked for content curation did not extend to commercial relationships. Enterprise contracts were confidential. Volunteers were not consulted. The structure that made Wikipedia succeed as a shared resource failed to protect that resource from external harvesting.

Ostrom showed that commons survive when eight conditions hold. The lesson is not that Wikipedia failed at commons governance. It is that commons governance cannot defend against capture *from within its own institutions.* Ostrom's principles assume the institution itself remains uncorrupted. They offer no defense when the guardian becomes the extractor.

### Institutional Capture

Institutional capture occurs when organizations formally designed to serve broad interests are progressively redirected to serve specific, often private, actors.

The mechanism does not require corruption. It requires only that institutional interests align with harvesters' interests—and they did.

The capture proceeded through predictable stages. First came dependency creation: Enterprise launched, revenue grew, AI companies became significant partners. Then conflict avoidance: WMF acknowledged noncompliance but continued deals with no public demands. Institutional alignment followed: AI strategy emphasized "humans first" while selling to AI companies. Finally, advocacy reversal: WMF no longer advocated for copyleft enforcement.

The Foundation was compromised not by bad actors but by structural incentives. Once Enterprise revenue mattered, legal action became commercially inconvenient.

### Other Commons

What happened to Wikipedia can happen to any open project.

Open source software has shifted from copyleft (GPL) to permissive (MIT/BSD) licenses. By 2022, only 22% of open source used copyleft, down from 45% in 2015. The shift enables corporate absorption without reciprocity.

Creative Commons—the organization that created ShareAlike—now says AI training should be "non-infringing by default." The entity that built the licensing infrastructure has accommodated the ingestion.

Research data faces similar pressures. Open access mandates enable AI training on scientific literature without attribution.

The pattern is consistent: create a shared resource, license it freely, watch institutions meant to protect it become dependent on those who harvest from it.

### Arc Complete

This book began with Nupedia's failure—an expert-controlled model that could not scale. The question was how to create free knowledge. The answer was radical openness: anyone can edit, trust the process.

Twenty-five years later, the question has inverted. The problem is not how to create free knowledge but how to keep it free.

The opening and closing images mirror each other. Chapter 1 began with Nupedia's expert control failing; this chapter ends with four scenarios, all fraught. That chapter showed one failed model; this chapter shows a successful model being extracted. The original problem was how to create free knowledge; the current problem is how to keep it free. The original dysfunction came from too much control; the current dysfunction comes

from too little enforcement. Hope was implicit then. Hope is conditional now.

The openness that enabled creation also enabled ingestion. The license that was supposed to protect the collective work was abandoned by the institution meant to enforce it.

## The Thesis

The copyleft framework was designed to keep knowledge free through legal mechanisms. That framework failed—not because the law was poorly drafted, but because:

**Legal interpretation is contested.** Courts may find AI training is fair use or does not create derivatives. The law Stallman designed cannot function if judges interpret it away.

**Technology circumvents law.** AI training transforms content into statistical weights. Attribution cannot trace from output to input. Share-Alike cannot apply to model parameters. The technology architecture defeats the legal architecture.

**Institutions must enforce.** A license requires enforcement to have teeth. The entity that should have enforced chose revenue instead.

**Political monoculture disabled opposition.** The Foundation lacked ideological diversity. FEC records documented employees donating exclusively to one party across 292 contributor-cycles—a pattern matching the tech companies whose compliance the Foundation was supposed to police. Professional networks overlapped; vocabulary aligned; adversarial instincts never developed. When enforcement questions arose, no internal faction argued for confrontation. Partnership felt natural to people who saw tech executives as colleagues.

Licensing cannot protect against co-option from within. The commons endures only when its defenders remain uncorrupted.

## What Remains

This book offers no redemption, because none has occurred.

The scenarios remain in motion. The outcome depends on choices not yet made—by courts, by the Foundation, by the community, by readers. The future is not determined.

But the lesson is clear.

Richard Stallman created copyleft because Symbolics took his work and refused to share. He designed a legal mechanism to ensure that free software would remain free.

Forty years later, AI companies did to Wikipedia exactly what Symbolics did to Stallman. They took the work of a quarter million volunteers, trained their models on it, and locked the results in proprietary systems. The ShareAlike requirement was ignored. Attribution was not provided.

The Wikimedia Foundation knew. Then it signed Enterprise deals anyway.

The license did not fail. The institution failed.

What happened to Wikipedia can happen to any open project where the institution that holds power becomes entangled with the forces the project was designed to resist.

The copyleft movement was built on a simple belief: that legal mechanisms could keep knowledge free forever. That belief was always more fragile than it appeared. The law is only as strong as the institutions that enforce it.

Wikipedia's founders believed they were building an encyclopedia that would endure on those terms. They created a license to guarantee it. They established a foundation to protect it.

Twenty-five years later, the encyclopedia remains. The license remains. The foundation remains.

But "free" has become something else: free for harvesting, free for training, free for commercial systems that compete with the shared resource they consume.

That is the story of Wikipedia. That is the lesson of the commons.

---

Robert Sim had 79,000 edits to Wikipedia. Twenty years of contribution. Articles about Southeast Asia that would not exist without his work. His labor had trained ChatGPT, Google's AI Overview, and a dozen other systems. None of them attributed him.

He did not feel like he was part of a shared project. He felt like his work had been absorbed.

He thought about editing the next morning, but did not. He thought about it the morning after, but did not.

The commons is not the encyclopedia. It is the moment when someone decides their contribution matters because they own it collectively. When that moment dies—when ingestion replaces ownership—the commons dies.

Robert Sim's 79,000th edit was his last.

# Epilogue

In 1985, Richard Stallman stood before a dead machine and wept. The PDP-10 had been more than hardware. It was the center of a community—programmers who shared code, improvements, ideas. When Symbolics took that shared work and refused to share back, the community scattered. The machine went dark.

From that grief came copyleft. Use my work, Stallman said, but share your derivatives. The license would spread, ensuring that free software begat free software. No one could ever again take what the community built and lock it away.

In 1999, Stallman proposed extending this protection to knowledge. "The World Wide Web has the potential to develop into a universal encyclopedia covering all areas of knowledge," he said. Wikipedia launched two years later under a copyleft license. The promise was clear: anyone could use the encyclopedia, but they had to share their improvements. Knowledge would remain free.

Twenty-five years later, Wikipedia succeeded beyond that vision in certain ways. 6.7 million articles in English alone. 250,000 active editors. Billions of readers every month. A free encyclopedia accessible to anyone with an internet connection—more comprehensive than anything ever published.

But the mechanism of protection failed.

Not because the license was poorly drafted. Courts had upheld copyleft provisions before. The legal arguments for enforcement existed: AI systems memorized Wikipedia text nearly verbatim, reproduced it in their outputs, competed directly with the source. Fair use factors weighed against the extractors.

The mechanism failed because the institution chose not to wield it.

---

In April 2024, the Wikimedia Foundation acknowledged that most of its AI customers were "not compliant with the letter of the Creative Commons rules or the spirit of the licenses." The Foundation kept signing deals anyway. By January 2026, OpenAI still paid nothing despite using Wikipedia more than any other source. ChatGPT still provided no attribution. Model weights remained proprietary. The invisible extraction continued.

When Grokipedia arrived—a visible fork with attribution, free access, and neo-Nazi citations—the Foundation objected to the content while partnering with companies whose structural violations were worse. The double standard was institutional capture made visible.

The wall Stallman built still stands. The legal mechanism remains enforceable in principle. But enforcement requires an enforcer with resources, will, and independence.

The Foundation had resources—$185 million in annual revenue, $400 million in reserves. It lacked will. It was not independent. It had become financially dependent on the companies whose compliance it should have demanded.

The guards were bought. The commons was captured.

---

We began this book with the degradation loop—ChatGPT citing Grokipedia, human curation laundered through AI layers, the visible fork with neo-Nazi sources treated as authoritative by the invisible extractor. We traced the path backward: the Enterprise pivot, the license switch, the founding schism, the copyleft promise. We asked how we arrived at this moment.

The answer was not Grokipedia. Visible forks with attribution, however flawed their content, preserve what copyleft was designed to protect. The answer was the invisible extraction—Wikipedia absorbed into proprietary model weights, attribution stripped, access paywalled—enabled by a Foundation that chose partnership over enforcement.

The answer is institutional capture. Not conspiracy. Not malice. Structural incentives that aligned the Foundation's interests with the extractors' interests until enforcement became "economically awkward."

This is not license failure. This is institutional failure.

---

The commons question for the AI age is whether any commons can survive such capture.

Elinor Ostrom's research showed that communities could govern shared resources without state or market intervention. Her eight principles explained why some commons endured while others collapsed. Wikipedia met most of those principles for content governance.

But Ostrom's framework assumed the institution itself remained uncorrupted. It offered no defense when the guardian became the extractor.

The pattern that killed the MIT AI Lab in 1982—take what the community built, add proprietary value, refuse to share back—repeated at scale in 2024. The only difference was the response. Stallman fought. The Foundation accommodated.

Other commons face the same pressure. Open source has shifted from copyleft to permissive licenses. Creative Commons now says AI training

should be ”non-infringing by default.” Research databases are scraped without attribution. The pattern is consistent: create a commons, license it freely, watch the institution meant to protect it become dependent on those who extract from it.

Licensing alone cannot protect against institutional capture. The law is only as strong as the institutions that enforce it.

---

Yet this is not where the story ends.

The volunteers remain. 250,000 editors still maintain Wikipedia every month. They still believe their contributions matter. They still show up, fix errors, add citations, argue on talk pages about neutrality and sourcing.

Their choice will determine the outcome.

If volunteers conclude the promise was broken—if they stop contributing, stop maintaining, stop caring—the commons will decay. The database will persist, but the community will scatter. Another machine will go dark.

If volunteers demand accountability—if they organize, protest, fork, or simply refuse to accept that extraction is inevitable—the Foundation will have to respond. Institutions can be reformed when their constituents insist.

The Spanish editors who threatened to fork in 2002 forced the Foundation to abandon advertising plans. Community pressure has shaped Wikipedia’s policies from the beginning. The volunteers are not powerless. They have not yet chosen to use their power.

---

In 1985, Stallman wept before a dead machine and asked: how do I ensure this never happens again? His answer was copyleft—a legal mechanism to keep shared work free.

In 2026, we ask a different question: when the institution entrusted to enforce that mechanism is captured, what then?

The copyleft promise was never self-enforcing. It required humans willing to defend it. Stallman fought Symbolics for two years, alone, because he believed the community was worth saving. The MIT hackers did not fight with him. They drifted away to industry jobs. The community died.

Wikipedia's volunteers face the same choice. They can drift away, accepting that their labor will train systems that replace them. Or they can fight—demanding enforcement, transparency, accountability.

The commons does not die when the license fails. It dies when the people who built it decide their contributions no longer matter.

Richard Stallman's grief in 1985 was not about a machine. It was about a community that would not fight for itself. The PDP-10 was already dead. What he mourned was the realization that no one else had tried to save it.

The volunteers who edit Wikipedia today are not yet at that moment. The community still exists. The work still matters. The question of whether to fight remains open.

This book cannot answer that question. Only the volunteers can.

But we can say what is at stake: not just an encyclopedia, but the belief that collective ownership of knowledge is possible. That belief survives only if someone defends it.

The commons was never the database. It was the decision, made millions of times by hundreds of thousands of people, that contributing to something shared was worth the effort.

That decision is made fresh every day.

# A Note on Sources and Methods

Writing the history of Wikipedia presents unique challenges and opportunities. The project's radical transparency means that nearly every decision, dispute, and discussion has been preserved in searchable archives. This abundance of primary sources allows unprecedented access to how the encyclopedia actually developed—but also requires careful navigation.

## Source Hierarchy

I have organized sources into tiers based on reliability and proximity to events:

**Tier 1: Wikimedia Foundation Documents.** Board resolutions, annual reports, Form 990 tax filings, and official statements provide authoritative information about the Foundation's operations, finances, and governance decisions.

**Tier 2: Wikipedia Primary Sources.** Policy pages, talk page discussions, arbitration cases, and edit histories document how the community actually functioned. These sources are invaluable but require careful interpretation—they show what people said publicly, not necessarily what they thought privately.

**Tier 3: Academic Research.** Peer-reviewed studies of Wikipedia provide systematic analysis of editing patterns, demographic composition, and content quality. I have relied heavily on work by researchers at the

Oxford Internet Institute, MIT, and the Wikimedia Foundation's own research team.

**Tier 4: News Coverage.** Contemporaneous journalism captures how events were understood at the time, though later reporting sometimes repeats earlier errors.

**Tier 5: Memoirs and Interviews.** First-person accounts provide valuable perspective but often reflect later rationalizations. I have used these cautiously, preferring documented evidence over remembered claims.

### Verification Practices

For claims about specific events, I have sought multiple independent sources. For claims about living persons, I have applied heightened scrutiny and, where possible, sought their response.

All web sources have been archived using the Wayback Machine. Wikipedia citations include revision IDs to ensure verifiability even as articles change.

### Limitations

Several limitations should be acknowledged:

This history focuses primarily on the English Wikipedia and the Wikimedia Foundation. The experiences of other language communities, while occasionally referenced, deserve their own treatment.

Internal Foundation communications and private discussions among editors remain largely inaccessible. The picture presented here is necessarily incomplete.

Wikipedia continues to evolve. Events described here should be understood in their historical context, not as descriptions of current practices.

### A Personal Note

I am neither a Wikipedia editor nor a Foundation employee. I approached this project as an outside observer, which brings both advantages (in-

dependence) and disadvantages (distance from the lived experience of participation). Readers should weigh my interpretations accordingly.

# Bibliography


[1] Wikipedia contributors. *Bomis*. 2026.

[2] F. A. Hayek. "The Use of Knowledge in Society". In: *The American Economic Review* 35.4 (Sept. 1945), pp. 519–530.

[3] Eric S. Raymond. "The Cathedral and the Bazaar". In: *First Monday* 3.3 (Mar. 1998).

[4] Larry Sanger. "The Early History of Nupedia and Wikipedia: A Memoir". In: *Open Sources 2.0: The Continuing Evolution*. Ed. by Chris DiBona, Mark Stone, and Danese Cooper. O'Reilly Media, Nov. 2005, pp. 307–338.

[5] Zachary Schwartz. *Wikipedia's Co-Founder Is Wikipedia's Most Outspoken Critic*. Vice. Nov. 2015.

[6] CNN. "Wikipedia's first ever edit sells as NFT for $750K". In: *CNN Style* (Dec. 2021).

[7] Edgar Enyedy and Nathaniel Tkacz. "'Good Luck with Your WikiPAIDia': Reflections on the 2002 Fork of the Spanish Wikipedia". In: *Critical Point of View: A Wikipedia Reader*. Ed. by Geert Lovink and Nathaniel Tkacz. INC Reader 7. Amsterdam: Institute of Network Cultures, 2011.

[8] Jimmy Wales. *Re: [Wikipedia-l] New to the Wikipedia with a few questions*. Feb. 2002.

[9] Andrew Lih. *The Wikipedia Revolution: How a Bunch of Nobodies Created the World's Greatest Encyclopedia*. New York: Hyperion, 2009.

[10] Jimmy Wales. *[Wikipedia-l] Advertisements*. Mar. 2002.

[11] Larry Sanger. *Jimmy Wales on Advertisement*. Jan. 2011.

[12] Larry Sanger. *Why I Am No Longer Participating in the Wikipedia Project*. 2002.

[13] Larry Sanger. *Why Wikipedia Must Jettison Its Anti-Elitism*. Dec. 2004.

[14] Storyboard18. *Wikipedia's Jimmy Wales Ends Interview in Under a Minute After Dispute over 'Founder' Label*. Network18 Group. Nov. 2025.

[15] Jr. Reagle Joseph M. *Good Faith Collaboration: The Culture of*

*Wikipedia*. History and Foundations of Information Science. Cambridge, MA: The MIT Press, 2010.

[16] Heather Ford. *Writing the Revolution: Wikipedia and the Survival of Facts in the Digital Age*. Cambridge, MA: The MIT Press, 2022.

[17] Sam Williams. *Free as in Freedom: Richard Stallman's Crusade for Free Software*. O'Reilly Media, Mar. 2002.

[18] Richard M. Stallman. "The GNU Manifesto". In: *Dr. Dobb's Journal of Software Tools* 10.3 (Mar. 1985), pp. 30–35.

[19] Richard M. Stallman. *Why Free Software Needs Free Documentation*. Free Software Foundation. 2002.

# Index