




Funding The Web

From Cartel to Covenant

Robin Berjon

2026-05-08



supramundane agency.

Browsers, and the browser engines that power them, provide critical public infrastructure to over five billion people and yet no one pays for their browser. To cover the high cost of maintaining such highly complex products, browser vendors and search engine providers have developed a system in which money is levied from search revenue and distributed to browsers. This arrangement benefits both: browsers don't have an obvious revenue model and web search engines are only relevant to the extent that the web itself is, which in turn requires a high-quality, frictionless browsing experience.

Over the years, this ad hoc system has succeeded in providing funding to browsers but it forms a de facto private government of essential web functions in ways that produce detrimental effects on the digital sphere in general: lack of transparency, poor funding of browser engines, defunding of other parts of the web, massive market concentration, plummeting search quality, harms to privacy, lower cost for effective disinformation, reduced revenue for quality journalism, and loss of opinion pluralism.

This report documents the existing system, analyses its negative externalities so as to determine requirements for alternatives, and proposes multiple paths forward along with enforcement and governance considerations. The existing levy distributes billions of euros annually, and shows no sign of slowing. The objective here is not to break or replace this system, but to place it gradually under the kind of governance that can make it operate in the public interest and in a manner that is resilient to ongoing AI-driven transformations of both browsers and search engines. It proceeds as follows:

- 1. A brief historical context leading to the situation of the web today.*
- 2. A description of the current system, how it evolved from mutualistic to parasitic, and how the levy on search that funds browsers is organised.*
- 3. A comprehensive description of the problems created by the current system.*
- 4. A review of the state of the art in evaluating the cost of maintaining a browser, with a surprising spread of evaluations.*
- 5. An exploration of the space of solutions, set out to develop our imaginary.*
- 6. Detailed considerations on the governance concerns that go into establishing a better system.*
- 7. A set of options and considerations pertaining to the enforcement needed to bring a better system into existence.*
- 8. Considerations pertaining to funding apportioning funds.*
- 9. And, finally, a conclusion reiterating the sheer enormity of the improvements we can bring to the digital sphere simply by addressing this problem.*

**

This work was funded under a grant by the Digital Infrastructure Insights Fund (<https://infrastructureinsights.fund/>), for which I am very thankful.

**DIGITAL
/INFRASTRUCTURE
/INSIGHTS
FUND**

Table of Contents

- Executive Summary 1**
- Ding Dong! The Web Is Dead 3**
- The Current System 5**
 - The Mutualistic Web 6
 - The Search Levy 7
 - Out of Scope 7
- Incumbent Problems 9**
 - Artificial Search Monopoly 10
 - Enshittifying Search 11
 - Defunding The Web 12
 - Unique Source of Funding 13
 - Monoculture Over Pluralism 14
 - Opacity and Unaccountability 16
 - No Browser Governance 17
 - Browser Engine Financing 17
 - Can Search *Not* Kill The Web? 18
- How Much Can A Browser Cost? 20**
 - So, What's The Tab? 20
 - Workshop on Browser Funding 21
 - The Ehrenberg-Russell Estimate 22
 - The Kardell Estimate 22
 - The Cooper Estimate 22
 - The Tarakiyee Estimate 23
 - The Ayala Estimate 23
- Solutions 24**
 - Voluntary Donations 24
 - Fiduciary Obligations for User Agents 25
 - Retroactive Public Goods Funding 26
 - Global Levy for an Open Web (GLOW) 26
 - The League of Digital Infrastructure 27
 - Web Infrastructure Search Endowment (WISE) 29
 - Search Protocol 30
 - Other Control Points 31
 - Deliberate Activism 32
 - Another Web Ahead 33
 - Putting It All Together 34
- WISE Covenant Governance 35**
 - High Level Requirements 35
 - Runtime Governance vs. Private Government 36
 - A SWIFT Precedent 37
 - Scope and Limits 38
 - Constituencies 38
 - Polycentric Approach 39
 - How to Not Break Browsers 39
 - How to Avoid Corruption and Oligarchy 40

Enforcement	42
Implementation Stages	42
<i>Bootstrap</i>	42
<i>Self-Enforcement and Growth</i>	43
The Digital Markets Act (DMA)	43
Isn't a Covenant of This Kind Illegal?	44
Aren't Choice Screens Bad?	45
Isn't Co-Regulation Terrible?	46
Reference Implementation	47
Fund Attribution	48
How Much to Levy from Search?	48
The Browser/Engine Split	49
Other Funding Opportunities	50
Conditionalities	51
Conclusion	52
Acknowledgements	54
References	55

Executive Summary

Over five billion people use the web through browsers which they never pay for, despite browser engines being complex and costly products. This is possible because browser vendors and search engine providers have developed a funding system — the *search levy* — in which search engines pay browsers (including public-interest ones like Mozilla Firefox) for traffic under specific conditions. This arrangement has succeeded in keeping browsers free and cutting-edge, but over two decades of operation it has also created an artificial monopoly over search with cascading negative consequences.

Taking a system view of the web, search is a logical place at which to apply a levy. Search extracts value from content and behaviour on the web (search engines would have no value otherwise) and renders it available at a functional control point since it is required for web discovery. Conversely, web content is valuable because browsers, funded by search, provide critical services that render that content available. As a result, the levy initially produced a virtuous cycle through which both search and the web at large grew together. Unfortunately, the levy failed to evolve as the web grew into the essential service that it is today. We don't need to get rid of this system: to get a healthy web, it is simpler to restructure the levy that we have through better governance.

This report documents how the levy works and which problems it creates, but more importantly what we can do about it.

The search market is worth several hundred billion dollars annually. That market is structured entirely by this search levy which mechanically produces monopoly: the search engine that can pay the most for default placement captures the most users, which generates the most revenue, which enables it to outbid all competitors. This self-reinforcing cycle, confirmed as an illegal monopoly by the US District Court in *United States v. Google*, has consequences far beyond the search market. It empowers Google to charge supracompetitive advertising prices, which defunds the broader web (including news media) by capturing advertising revenue that would otherwise circulate to publishers and content creators. It eliminates pluralism by concentrating editorial power in the hands of a single search engine. It fails to guarantee funding for browser *engines* — the most expensive and critical component of web infrastructure — which leaves the evolution of web technology entirely in the hands of the companies that most benefit from the system: Apple and Google. And it provides no mechanism to hold browsers accountable when they betray their users, as Google Chrome does through pervasive deceptive design patterns that harvest people's entire online behaviour.

A direct consequence of this market structure and of who it grants power to is that the web isn't a commons but rather private property that the public is granted access to.

The report, informed by research, policy entrepreneurship, a workshop with the browser engine community, and direct (confidential) interviews with industry actors, outlines a set of solutions and proposes the **Web Infrastructure Search Endowment (WISE)** as the primary reform: a multistakeholder institution that would formalise the existing levy under democratic governance, manage browser and search choice screens through an open, accountable expert group, and condition eligibility on adherence to rules protecting people and funding



web infrastructure. WISE can be bootstrapped from a single jurisdiction — most plausibly through the EU's Digital Markets Act. Complementary reforms include fiduciary duties for user agents, the gradual transition of search from a website to an open protocol that supports multihoming, and the funding of standards, testing infrastructure, and advocacy from levy proceeds. Governance draws on the precedent of SWIFT, the coöperative that manages critical global financial infrastructure through multistakeholder coöperation.

Even at a fraction of current levy rates and after accounting for the elimination of monopoly pricing, the revenue potential comfortably exceeds the cost of maintaining multiple browser engines, investing in standards and testing, and funding the democratic institutions needed to govern web infrastructure in the public interest. It makes it possible to reposition the web as an explicitly hegemonic project in support of agency and democracy after decades spent pretending that infrastructure governance is apolitical while monopolies quietly write the rules. The problem has never been that there isn't enough money. The problem is that no one accountable to the public has ever decided where it goes. This report charts a path to changing that.



Ding Dong! The Web Is Dead

The web has been pronounced dead many times over. It's easy to take this as a sign of its resilience but we should be wary of unfounded optimism. The fact that the web is still kicking around in some form is less a testament to its inherent resilience, let alone its good governance, than a consequence of the time it takes for such a large system to stop moving. The wheel may still be spinning but the hamster is long dead (Dash 2026).

The web's death is commonly attributed to "Artificial Intelligence," by which people mean LLM disintermediation of search and displacement of older advertising delivery mechanisms. The Financial Times thinks the web may be fatally wounded (Hodgson et al. 2025), the Economist asks "*can anything save it?*" (The Economist 2025), the BBC described the web "*destroyed*" (Germain 2025), the Observer has the web's economy "killed" (Tremayne-Pengelly 2025), while the Guardian asks "*so why is nobody trying to stop it?*" (Mahdawi 2025). But it is this report's position that the problem started much earlier and that while LLMs do make it worse they are the continuation of a deeper problem that predates the launch of today's machine learning products and platforms.

This deeper issue, thankfully, is fixable. And because the web is, for now, still slowly settling into its purported "death", *we may yet revive it*.

The digital sphere we're in sits downstream of deliberate, explicitly neo-liberal political choices made in the 1990s. When the Clinton-Gore administration set out to define their hugely influential framework for global electronic commerce, principle number one was: "*The private sector should lead.*" This wasn't simply an aspirational statement gesturing in the direction of private-led innovation (the administration was also busy redirecting massive public investment from Cold War-era military industries towards the ascendant Silicon Valley), but a comprehensive vision for the governance of the internet, which they promoted



worldwide, wrote into a wave of world-historical regional trade agreements, and consistently returned to:

*“For electronic commerce to flourish, the private sector must **continue to lead**. (...) governments should encourage industry **self-regulation** (...) private entities should, where possible, take the lead (...) governments should refrain from imposing **new and unnecessary regulations** (...) governments should establish a **predictable and simple legal environment** based on a decentralized, contractual model of law” (Clinton-Gore Administration 1997, emphasis mine).*

That the private sector may “lead” the inter-corporate governance structures pertaining to simple consumer products is often unproblematic, and self-regulation in finance is unchallengeably structural to the world economy as it is currently governed. That said, private governance over large-scale infrastructure (whether financial or physical) that sets rules and liability structures and information flows for all downstream actors, both businesses and individuals, concentrates power in ways that have well-documented deleterious effects and perverse incentives even without collusion (Rahman 2018; Frischmann 2013; Barzilai-Nahon 2008). In the digital world, following thirty years of this excessively private-centric, globally-harmonised laissez-faire policy on all kinds of information, innovation, and transaction platforms, we are dealing with what is in essence the **private government** of our lives and livelihoods (Anderson et al. 2019) by platforms with more direct control than that enjoyed by our governments.

In the case of funding for web browsers and related parts of core web infrastructure (and with impact on the funding of content and services on the web in general), the world’s central information flows are under the opaque and extractive private government of Google Inc. The funding of all major browser engines and of over 92% of browsers by market share comes from a privately-administered levy on Google Search, essentially granted in exchange for enforcing Google’s monopoly in search. This is not in the least a “market-based solution”, but rather a private mechanism (Viljoen et al. 2021) that interferes with the healthy operation of not just search/browser markets but competition more generally in the entire digital sphere, to say nothing of downstream effects for our information ecosystem, democracy, and society at large. As I will show below, for as much as this system bankrolls critical infrastructure benefiting all, it does so at a steep cost for the rest of society that is all but hidden from mainstream political discourse.

Concerning as it is, the centralisation and scope of this mechanism offers us an opportunity. What if we kept it going but transitioned it to democratic governance so as to ensure that it funds what needs funding? How could democratic governance chip away at, and eventually counterveil, the downstream issues that it currently creates?

This report builds not just an understanding of today’s levy arrangement and its issues, but also extends our imaginary and outlines policy entrepreneurship that can bring a healthier web back to life. It does so by opening the door to a long-needed alternative vision to that of the Clinton-Gore administration, which has more than run its course as we live through the aftermath of its success. We find ourselves at the beginning of a more democratic internet as can be seen almost everywhere in emerging projects and tendencies — redirecting the capital flows from the current levy system has the potential to supercharge these trajectories and make a new internet bloom.



The Current System

As of 2022, 63% of the world population, or roughly 5 billion people, used the Internet (ITU 2022). While some of these may not regularly use a *browser* per se (the application used to navigate the web), they all use a *browser engine* (the software component that both browsers and other apps use to render web content) almost every time they use a device as those are widely used in applications in addition to forming the core of browsers. Essentially all mobile and desktop applications that connect to servers are variously thin or sophisticated wrappers around a browser engine, unbeknownst to end-users who simply see them as so many “applications” in their own right.

Browsers and browser engines are provided to people free of charge, as public goods. Their continued provision and security vigilance is critical since without them the web collapses, as the attack surface for most apps without this upstream economy of scale would be enough to bankrupt whole swaths of software development. The browser engines are also staggeringly complex: for instance, in 2022 Chromium (the browser engine powering Google Chrome and other browsers) ran to about 35 million lines of code which puts it in roughly the same engineering weight class as entire **operating systems**. The exact cost of maintaining the three primary browser engines or of operating a major browser is subject to debate (covered in a later chapter) but even the higher estimates, while significant, amount to only a small fraction of the direct value (monetary or otherwise) produced by the web.

In order to assemble the funds required to continue to operate, almost all browser vendors rely on variations of the same strategy: they exercise an *ad hoc* levy on search engine revenue. Variations on this levy include: selling the search engine “default” (the initial and typical approach); royalties on search volume; and/or, intra-company balance transfers (when the search engine and browser belong to the same company). Taken together, these strategies form what I will refer to in the rest of this piece as the broader “levy system”, which is detailed further down in this report.

Despite its critical importance to web infrastructure, the levy is opaque and poorly documented, and suffers from a number of undesirable shortcomings that have large-scale detrimental effects on the web. This report looks at those issues in detail, establishes



requirements from them, and outlines a set of approaches that could be combined in various combinations to form a range of viable solutions.

Note that search is a logical control point at which to establish such a levy (Clark 2012). Almost all of the value of a search engine comes from the web it indexes — the content and services that make it worth searching through in the first place — which means that search engines have an incentive in keeping the web healthy *at least so long as they aren't in a position to enclose it all*. A critical way to keep the web healthy is to invest in browsers and more generally in web infrastructure as that empowers the many people building web things to create greater marvels yet. This has the potential to create a positive cycle through which key web infrastructure is funded, leading to a better web, driving more revenue to search engines, and in turn levying more funds with which to support the health of the web.

The Mutualistic Web

The relationship between publishers and search engines originally constituted a simple and rather fortuitous mutualistic affair across the period of generalised and steady growth. The arrangement was elegant in its reciprocity: publishers made their content freely available for search engines to index, which benefited the search engine by making its results relevant and comprehensive. In return, the search engine used that content *only* to index it and make it discoverable (with no *substitutive purpose*, crucially), driving traffic back to the publisher. No formal contract or standardised protocol governed this exchange — it was an alignment of incentives which emerged organically from the web's architecture and was sustained by a simple feedback mechanism. If a publisher withheld content, it lost traffic and if a search engine became extractive, publishers could exclude it via robots.txt and the engine would suffer from less relevant results.

The critical precondition for this situation was competition: so long as multiple search engines vied for users, no single engine could afford to take more from publishers than it gave back, because publishers had alternatives and the costs of defection were manageable on both sides.

Unfortunately, this mutualistic relationship collapsed under market concentration (as a direct result of the mechanism described in this report). As the search side of the deal became increasingly monopolised, publishers lost their leverage. When traffic from a single search engine constitutes your livelihood, you no longer have a meaningful choice to exclude it: you do as you are told, no matter how much it hurts. Google exploited this dependency incrementally more over time, simultaneously doing more with publisher content than merely indexing it (e.g. featured snippets, knowledge panels, voice assistant answers, AMP, and eventually AI Overviews) while demanding more from publishers in terms of format compliance and data access, and sending steadily less traffic back over the years.

It has been a painful death by a trillion hits: every time search terminates at the search engine (or sooner, as when a chatbot replies with a fact extracted from a publisher without attribution or royalties of any kind), that user is not visiting the publisher's property, and that shaves another penny that won't go towards supporting that publisher's content. Individually, it's imperceptible, but there are on the order of a trillion searches every year: it adds up.

From a mutualistic genesis, the relationship became parasitic over time, and in many cases (like professional journalism) is steadily killing the host.



The Search Levy

The **Search Levy** (henceforth just “levy”) is a system in which a portion of search engine revenue is levied and used to pay browsers. (It is also used for operating systems and other agents, but our focus in this document is browsers.) The levy can take multiple general forms:

- **Default Placement.** The browser vendor sells its default search function to a search engine. Defaults have an outsized impact on search volume as few users change them, in part because users don’t necessarily distinguish browser from search since the experience of both is intertwined, and in part because changing the default is made hard. This practice also includes selling a *presence* (being mentioned in the initial list above a manual option or longer list) or a *ranking* (i.e. “second position”) on the list of search engine alternatives that browsers provide to end-user configuration interfaces.
- **Royalties.** The browser vendor can also negotiate a share of the revenue when directing traffic to a given search engine. Royalties and default placements are not mutually exclusive.
- **Intra-Company Transfers.** When the browser vendor and search engine vendor are the same company, revenue from the latter contributes to paying for the former and can be accounted internally as balance-transfers and/or structured profit-sharing contracts between distinct budgets or subdivisions. In effect, this is an integrated variant of default placements that is cheaper and of greater strategic value. (For smaller search engines, it may also be the only deal they can make, which drives the existence of the DuckDuckGo or Ecosia browsers.)

These deals are negotiated bilaterally and under extreme confidentiality. In practice, a levy agreement between a browser vendor and a search engine vendor can marry aspects from several of these forms and entail byzantine contracts that take forensic lawyers weeks to understand, much less entangle.

Levies are economically significant and constitute one of the larger commercial exchanges on the web. For instance, in the year 2021 Google Search spent over \$26 billion USD on default placement deals (Pierce 2023) — which does not include its intra-company transfer costs — and paid 36% to Apple of its search advertising revenue made in Safari (Nylen 2023). In 2020, Mozilla made 86% of its revenue from the Google Search levy (Lardinois 2021).

Note: As discussed below, the levy system is very opaque and most relevant numbers are not available unless they were revealed through legal discovery. Even though some of these numbers are relatively old, interviews with people involved in these deals indicate that the situation has not significantly evolved compared to what information is publicly available.

Out of Scope

In this report, I am treating the following items as mostly (but not completely) out of scope:

- **Mobile OS.** Mobile operating systems benefit from a dynamic similar to that of browsers in being funded in part through a search levy. However, integrating them completely would have made the report too complex in that they benefit from other levies (most importantly from “app stores”) and drive a number of other consequential OS-level defaults and configurations (e.g. for geolocation/navigation, email, messaging, etc.). Fixing the levy for browsers would open the door to solving governance issues in mobile OSs



as well, but that more complex market would benefit from further research and detailed extensions into that space.

- **Complete Analysis of LLMs.** The report *does* cover LLMs where they meaningfully intersect with the levy, otherwise it would be stale already. However, I have held back from a full-on treatment of the topic, for two primary reasons. First, the space is dynamic and it's too early to know where massive markets will land (for instance, regarding advertising disclosures or transparency, market structure for intermediaries, etc). Google clearly avoided using some of its distribution channels (search, browser) at least so long as Judge Mehta hadn't rendered his decision on anti-trust remedies (*United States v Google 2024*), so that they could claim that the impact wasn't there — which has delayed market dynamics that might otherwise have already hit production. Where I do see dynamics being outlined, they seem to point to there being strong potential incumbency power at work as well as high value in having access to browser-like distribution channels, which would potentially align the space with today's levy system and thus make this report more directly applicable. But after careful consideration, I decided against getting caught up in excessive speculation and only touching upon AI when the interaction was clear in the present or could restructure or destabilise the current levy system.

Incumbent Problems

The levy was an important invention and remains a good starting point twenty years on. Taking a system view of the web, search is a key, logical place at which to apply a levy. Search extracts value from content and behaviour on the web (search engines would have no value otherwise) and renders it available at a functional choke point since search is one of the required components of web discovery (Clark 2012). In turn, web content is only valuable in the first place because browsers provide critical infrastructure services that render that content available, make browsing it more attractive to people, derisking that browsing, and to a large degree molding that browsing via recommendation and reputation systems.

Initially, the levy produced a virtuous, regenerative cycle: the search levy financed core web infrastructure and that core web infrastructure made the web continuously successful such that search engines enjoyed strong businesses that could then be levied from.

Unfortunately, the levy failed to evolve as the web grew from a time of relative infancy when only 14% of the world population (fewer than a billion, (ITU 2022)) used the Internet to the essential part of global society that it is today. Over time, serious shortcomings with this *ad hoc* approach have surfaced that have not been addressed to date. It is this report's position that this needs to change if we are to return the web to better health. The changes need not be revolutionary since the core principle — a levy on search to pay for web infrastructure — can readily remain the same and much of the existing system can remain in place. As concentrated power in the hands of a small number of companies continues to drain the rest of the web of its revenue, we risk losing the web altogether — a problem a long time in the making but now critical enough that the mainstream press has noticed (even if the problem often remains confused with appeals to technological determinism).

This chapter captures a relatively comprehensive view of the levy's issues so as to provide a solid shared background from which to produce requirements and interventions.



Artificial Search Monopoly

The manner in which the levy system is structured means that it *mechanically* increases concentration in the search market. Few people change their default search engine even once — particularly on mobile devices — which means that purchasing a default placement effectively purchases market share. In turn, the search engine with the highest market share has greater profits (including supracompetitive ones from its ability to over-charge as a monopolist) from which to pay a higher price for further default placements. This eventually leads to almost every browser defaulting to the same search engine, and that search engine dominating the web by insuperable margins.

Interestingly, this market dynamic is such that the search engine that happened to be ahead when the default payment model became standard practice can then by happenstance become perpetually dominant. Short of a major upset to the whole web, there is no point at which this cycle corrects itself: **whoever makes the most money is whoever can pay the most for defaults, and whoever can pay the most for defaults stands to make the most money.** Breaking it requires external intervention. **By any rubric of anti-trust, this is an automatic and self-reinforcing monopoly at the infrastructure layer.**

This model is, of course, not theoretical: in *United States v. Google* the court's central finding was that Google "is a monopolist, and it has acted as one to maintain its monopoly." **In 2021, Google's payments to secure preloaded defaults totalled more than \$26 billion — nearly four times more than all of its other search-specific costs combined.** The scale of these payments is itself diagnostic: Google would not pay this much unless it expected to receive something of comparable value in return, which helps put a price on search engine defaults. This value was captured crisply in internal Google slides obtained as part of the discovery process in that case.

How Chrome creates value for Google

Chrome drives revenue and/or efficiency gains for Google's top PAs. Highlights include:

\$ Impact	Lever	PA
XL	Chrome MAU * LTV per platform	Search, Ads
M	Incremental revenue	Search, Ads, YouTube
M	Signed-in Google users	Search, Ads, YouTube
M	Web influence & 1P delivery enablement	Search, Workspace, YouTube, P&E, Core
M	Google Payments FoP and GMV*	Commerce
M	OpEx savings from Codecs & Widevine	YouTube, Android
S	Chrome Enterprise Premium sales	Cloud Security

The case's focus on Google, however, tends to understate the degree to which browser vendors are not merely passive recipients of Google's payments but active participants in a structural arrangement that serves their own interests. The collusion here does not need



to be conspiratorial as there is no back-room negotiation needed; the incentive structure makes this kind of coöperation self-executing and frictionless. Browser vendors have a direct financial stake in the perpetuation of search concentration: Google's monopoly pricing power (in advertising), which generates the supracompetitive profits from which the levy is paid, translates directly into revenue streams for affiliated browsers. **While this has yet to be legally determined to constitute cartel behaviour, it is factually the case that Google Search and the browsers it pays (Apple Safari, Mozilla Firefox, Samsung Internet, and others) work in a way that all know maintains prices high and restricts competition.**

This collusive logic extends beyond payment arrangements into the technical and governance architecture of the browser ecosystem. The "voluntary standards" approach that browser vendors resolutely defend in bodies like the W3C entrench a system in which they are the final arbiters of all meaningful decisions affecting the functioning of the web, presenting as *technical consensus* what are in substance the *commercial preferences of a small group of megacorporations and their incentive-aligned vassals*. The problems created by the levy system are well-known to the web community at large, but when proposals have been raised for the W3C to exercise collective governance over search defaults — creating a standards-based mechanism that could help address the worst of the levy's issues — major browser vendors have resisted, one representative stating in no uncertain terms, as recently as September 2024: "W3C's value isn't in its ability to put limits on browsers."

What's more, it's too early to understand how the search market may be reshaped by the advent of LLM chatbots, geopolitical and economic paradigm shifts, and other major world events, but it's very clear that Google's structural position in search equips them with a direct path to continued monopoly precisely by using the same methods that they use to artificially dominate search, as described in this report.

Enshittifying Search

Cory Doctorow's concept of *enshittification* describes a predictable sequence in the decay of platform value: platforms first allocate surplus to users to build dependency, then withdraw it in favour of business customers, and finally extract value for shareholders, leaving everyone else progressively worse off but typically unable to leave due to the two-sided nature of the market in question (Doctorow 2025). Search has slightly different dynamics in that users don't care about the presence of advertisers and publishers can be indexed by many search engines (multihoming) at very low extra cost. Nevertheless, Google has managed to lock most web users in by **purchasing defaults** anyway, which then locks advertisers and publishers in because they need access to the users.

This power relationship has long enshittified the web for publishers. Under competitive conditions, publishers who allow their content to be indexed by a search engine face a meaningful check on the search engine's extractive behaviour: if one search engine becomes too parasitic, publishers can de-index from it, and users who find results deteriorating can defect to alternatives. But when a single search engine commands 90+% of search traffic, de-indexing is commercial suicide for publishers. There is no credible threat of exit, publishers have to do as they're told as we've seen with AMP or CWV (Google 2015).

With users locked in, the quality degradation predicted by the enshittification dynamic has been empirically documented. A systematic comparison of search engines across a repre-



sentative set of queries found Google performing worse than several alternatives — including Marginalia, an independent search engine running on a fraction of Google's resources — and returning, for the most routine informational queries, a mixture of scam sites, low-quality SEO content, and ad-laden pages designed to extract clicks rather than answer questions (Luu 2023; Ovide 2025; Nield 2026). Recently, Google Search was observed producing results that don't even reflect the correct content of the pages it points to, generating fake headlines instead (Hollister 2026).

Another aspect of search response quality is the privacy properties of the pages being pointed to by the Search Engine Results Page (“SERP”). Privacy is a regular concern for users and a search engine could use the presence of tracking in a page to either downrank it (in service of the user) or uprank it (because the search engine benefits from lessening privacy). A recent report looked at this aspect alongside quality and found that Google came up short on both aspects (Kollnig 2025).

Remarkably, Google's enshittification has further led to the rest of the web enshittifying accordingly. It's not merely that Google's results have gotten worse, but that Google's dominance has deformed the *entire web* in the direction of its preferences (Sato 2024). When a single ranking function governs access to the majority of web traffic, the rational response for every publisher is to optimise for that function rather than for users — which is a form of enshittification, albeit one driven by the fact that web traffic isn't a market but rather a proprietary, opaque mechanism largely driven by Google. The result is not only that bad content rises in rankings; it is that good content, written for humans, is often outcompeted by content engineered for algorithmic legibility.

The entire texture of the public web has been sculpted by the preferences, the whims, the errors of a single opaque, non-deterministic, and continuously-evolving algorithm operated by a single company for its own commercial benefit. This is what *structural* enshittification looks like: not just a platform that has gotten worse, but a platform whose gravitational field has pulled the surrounding ecosystem down the toilet with it.


Defunding The Web

Google Search advertising alone represents approximately \$224.5 billion of the roughly \$600 billion global digital advertising market — about 37% of all digital ad spending. As explained above, the levy system is critical to driving Google's market share up, but that is not its only effect. In *United States v. Google LLC*, Judge Amit Mehta went into extensive, damning detail to show that Google's monopoly position, enforced by exclusive distribution agreements with browsers and mobile OSs, enabled it to charge supracompetitive prices for search text ads, with no only apparent constraints (United States v Google 2024). He explains at length how Google was able to maintain revenue growth for search at or above a 20% target — representing about 2/3rd of the company's overall growth — for over ten years, largely by playing on ad pricing. Putting it quite simply: “*Unconstrained price increases have fueled Google's dramatic revenue growth and allowed it to maintain high and remarkably stable operating profits.*” As far as is known the problem continues and the European Commission has recently opened an investigation into precisely this practice (Stolton 2026).

The web's overall revenue, however, is not infinitely extensible. Marketing budgets are set and marketing teams tend to treat search as just one of the ways to reach people on the



web. The ad revenue that flows to Google Search ads doesn't go elsewhere. This means that, as illustrated below, not only does Google Search get a share of all web advertising revenue (which would be fair), not only does it get a share that should flow to competing search engines (which would be problematic), but it *also* gets a share that should go to non-search web properties (e.g. publishers), which results in large-scale defunding of the rest of the web.



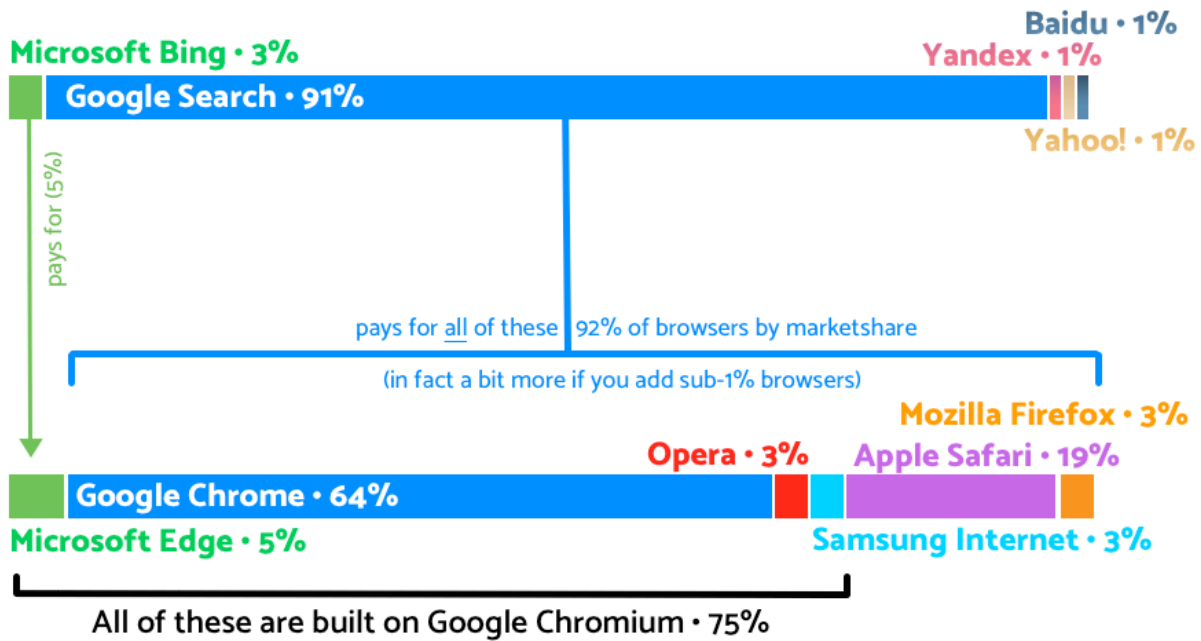
Browsers directly benefit from defunding the web, since they get a cut of the additional revenue. Critically, much of the value captured through the search/browser levy never returns to web infrastructure. Of the billions extracted through this process, only crumbs fund web-related infrastructure.

Additionally, the gravity of Google's search revenue undermines potential changes to the advertising system as a whole because its sheer gravity limits what options can be adopted elsewhere without bolstering it even further (Berjon 2025b). This further compounds the problem by making it harder to fix the web's many problems with ad-based revenue models.

Unique Source of Funding

The current structure of the search and browser ecosystem can give us a sense of how concentrated the browser's funding model is. On the search side, we have Google with a 91% share of global query volume and Microsoft Bing at 3%, everyone else at 1% or less. On the browser side, we have Microsoft Edge at 5%, and it's financed by Bing. Every single other browser with market share higher than 1% (accounting for a total of 92% comprising Chrome, Safari, Firefox, Samsung Internet, and Opera) is financed by Google Search.





Looking at browser engines, the picture is even more concentrated. Three major browser engines remain in the world — Google's Blink, Apple's WebKit, and Mozilla's Gecko — and all three are financed by Google Search. (There are promising developments around Servo and Ladybird, but as things stand today a consumer-grade browser can only be produced using those three). Since Microsoft Edge uses Blink, over 97% of the market runs on Google-financed engines (a number that's even closer to 100% once we add in the sub-1% browsers).

This shows both just how concentrated the market is, but also how dependent it is on a single source of funding. *A strategic shift from Google towards apps or LLMs, for instance by de-emphasising the web on Android and driving users to apps instead, could upset the entire web ecosystem overnight.*

Monoculture Over Pluralism

Every search engine embodies, *de facto*, an editorial policy. Like newspapers, television channels, or radio stations, a search engine frames what is and isn't relevant — what appears prominently, what is buried “beneath the fold,” how attention is directed and modulated, and what effectively ceases to exist. The implementation differs — the editorial policy is applied on demand and automatically, calibrated through algorithmic signals rather than editorial meetings — but the function is identical, and no less subjective, than that performed by any other form of media. As Cathy O’Neil puts it, an algorithm is just someone’s opinion written in code (O’Neil 2016).

When a newspaper editor decides which story leads the front page, we understand this as an editorial judgment shaped by institutional values, commercial incentives, and individual biases. When Google's ranking algorithm makes an equivalent determination, we have somehow convinced ourselves that this constitutes a neutral, technical process (the internet offered up this search result!) rather than a highly consequential editorial operation (Google’s



algorithm and market incentive favoured this search result). Google's editorial decisions have a greater impact than those of any newsroom but with none of the institutional guarantees, appeals or corrections process, or accountability.

The consequences for information integrity are direct. **As a single ranking system becomes the dominant arbiter of visibility, the signals it uses to assess quality become the targets of manipulation.** This results naturally but quietly from a system where the appearance of authority, evaluated by a single set of criteria, progressively decouples from genuine expertise, and misinformation that successfully mimics those quality signals becomes increasingly difficult to distinguish from legitimate content. Any monoculture of relevance determination creates a single, legible attack surface for those who wish to manipulate the information environment. In a pluralistic search ecosystem, where multiple ranking editorial choices coexist, manipulating one system's signals would not automatically grant visibility across all others. **The algorithmic monoculture makes disinformation automatable and scalable, while making appropriate countermeasures inhibitably expensive** and politically unpopular *a priori*.

This problem compounds the longer it lasts since it impacts which sources of information are viable, but as a second-order effect also orients *which content gets* recirculated and thus influences the creation of future content. **When a single algorithm determines the relevance of all indexed human knowledge, it necessarily and, over time, recursively imposes a particular epistemological framework** (which is to say particular assumptions about what constitutes a credible source, what forms of evidence matter, and which knowledge traditions deserve amplification). Knowledge systems that do not conform to the algorithm's implicit model of authority are not censored in the traditional sense; they are rendered essentially invisible by deranking and deprioritizing, which amidst the volume and velocity of information available today amounts to the same thing. Ultimately, this implements *epistemicide*: the systemic destruction or displacement of alternate ways of knowing.

We've lived in this world long enough that many can barely imagine alternatives, but the antithesis — algorithmic pluralism — is gaining traction, particularly in Europe (Collectif 2024; Szymielewicz et al. 2025; Elsayed-Ali and Berjon 2025). In the context of search, this would mean that end-users (and corporations) could directly choose between ranking systems that embody different editorial programmes and directly impact the information intake and cognitive health of each user.

The democratic stakes are not hypothetical, nor are they subjective. When a single company's ranking algorithm can determine which political candidates receive visibility, which policy positions appear well-supported, which scientific findings seem authoritative, and which cultural narratives gain traction, the resilience of democratic discourse depends entirely on the benevolence, competence, and political neutrality of that company's leadership and implementation quality — properties that are neither guaranteed nor stable across time, and that we have been watching collapse in Google Search over a short few years as market incentives and power structures shift quickly.

Media pluralism has never been advocated for based on claims that media owners are malicious or extractive; rather, media pluralism has generally been presented as a *prima facie* public good, based on the principle that concentrated control over information flows is structurally incompatible with democratic self-determination, regardless of the (always temporary) character of those who hold that control. And the larger the scale, the more this



principle applies. The question is not even whether Google is currently abusing its editorial power over search — even though it is — but whether any free society should permit such concentrated power over cognitive security, science, and democracy to exist at all.

Opacity and Unaccountability

There is very little **transparency** in today's levy system. Almost everything the web community knows about it has been disclosed not through any governance process, but through material released from discovery processes in public court cases, leaks, forensic analyses, and whistleblowers. Most participants in the web economy, including seasoned developers and large, established digital properties, have no idea how the system operates. As a result, when the system produces distortions (and as shown in the rest of this chapter, it produces many) there is currently no mechanism through which the affected web community can identify the problem, trace its origins, or propose a remedy. If we are to take the W3C's stated goal of building a web for all humankind as anything other than a suburban lawn sign slogan (Wilson and Çelik 2025), we need to ensure that the web community is able to evaluate how the beating heart of web infrastructure funding actually operates qua (intransparent but still measurable and analysable) market phenomenon.

The absence of transparency is compounded by a complete absence of **accountability**, and this has direct consequences for users. The bilateral deals struck within the levy system can impose requirements on browser vendors that run contrary to users' interests without any form of oversight or accountability towards those users (mediated by disclosures to third parties or otherwise). A search engine may, for instance, require that the browser **abstain from implementing** certain privacy protections, or forbid the browser from modifying the Search Engine Results Page (SERP) even when such modifications would benefit the user. For instance, a search engine may benefit commercially from making advertisements visually indistinguishable from organic results or from pushing users towards being logged in, which sacrifices privacy in support of its advertising business. A browser acting in its users' interest would be expected to counteract such practices, as indicated in the W3C's Ethical Web Principles when they establish the priority of constituencies, mandating that potential benefits for users take precedence over benefits to other ecosystem participants (Appelquist et al. 2024).

None of these concerns are theoretical, or even particularly hard to corroborate despite being explosively litigable. I know from direct conversations with browser engineers and product people (under conditions of anonymity) that Google has delayed the shipping of privacy protections in *competitors'* browsers based on search deal terms. This asymmetry is worsened by the fact that browsers have no foolproof way of calculating the revenue produced by the traffic they drive and so have to take Google's projections (and their royalties) at face value when told that a given change may reduce their levy royalties.

People who work for browser vendors often confide that the arrangement makes them uncomfortable. The browser's entire *raison d'être* is the trust of the user (Nottingham 2020) but the levy falls far short of that commitment. Browsers tout their privacy features but then turn around and send users to Google which is known for its highly-invasive processing. Apple prides itself on end-to-end privacy guarantees in its vertically-solidated hardware and software ecosystem, only to default search engine primacy to Google, vegans all the way to the slaughterhouse.



No Browser Governance

Browsers are expected definitionally to be “*user agents*”, software that acts on behalf of the person and in their interest in a trustless, open environment (Berjon and Yasskin 2025; Nottingham 2020; Appelquist et al. 2024). (Websites and endpoints at stable domains are assumed, conversely to represent the interests of one or more legal persons and enforce business logic fairly across all classes of user agent.) These are not supposed to be aspirational platitudes but rather to articulate a key architectural assumption on which the entire web platform is built. The browser is the user's fiduciary in the digital sphere, the intermediary entrusted with representing their interests in every interaction with every server (Berjon 2021b).

In the absence of **enforcement** of this expectation, however, this architectural expectation finds itself violated, which in turn puts the web off kilter. Google Chrome, for instance, remains the only major browser without meaningful privacy protections. Worse, it deploys a comprehensive set of highly deceptive design patterns to trick its users into disclosing their near-entire online behaviour to Google (Munir et al. 2025; United States v Google 2024). Architecturally, it doesn't deserve to be classed as a user agent but rather as Google's optimised access mechanism for reaching and observing users. There exists, however, no mechanism within the web standards community, within the levy system, or within any known governance structure through which the web community could identify this betrayal, hold the browser accountable, and compel a change in behaviour.

This is a key structural gap that weakens the web: **there is no runtime governance of browsers.**

Browser Engine Financing

The levy finances browsers, but it does not finance browser *engines*. To recall, **a browser is the application that users interact with whereas the browser engine is the vastly more complex substrate beneath it** (or embeddable and white-labellable in other applications). Key components include the “rendering pipeline”, the increasingly load-bearing and complex-to-secure JavaScript runtime required for web app interactivity, the rapidly-iterating and multi-protocol networking stack, the security sandbox, the implementation of hundreds of standards that collectively determine what the web platform can do (Ayala 2025). Building and maintaining a browser engine is *far* more expensive and technically demanding than building a new browser shell atop an existing engine would be. Indeed, most “alternative browsers” do the latter with a miniscule fraction the budget it would take to attempt (much less guarantee) the former.

Yet the levy's financial flows are structured around the browser as a distribution endpoint for search queries, not around the engine as a piece of shared infrastructure. A browser vendor can adopt an existing engine (often Google's Blink/Chromium), collect revenue-sharing payments from a search engine for default placement, and contribute nothing whatsoever back to the engine's development. This is textbook “free riding”: the system rewards building thin browser shells that capture levy revenue while externalizing the cost of engine maintenance onto whichever entity underwrites the engine's development. Yet, we do need engine diversity lest all decisions pertaining to web technology be made by a single team (which, indeed, is increasingly the case today).



Because browser *engine* development is not directly funded by the levy, only organisations with either massive independent resources (Google, Apple) or mission-driven commitment (Mozilla, which derives the overwhelming majority of its revenue from Google) can maintain one, much less research, design, and implement a new one or a fork.

Can Search *Not* Kill The Web?

Looking back over the past decades, it's tempting to ask: is it even possible for search *not* to kill the web?

Today's shift in which AI overviews are substituting for publisher visits (and thereby destroying the revenue model of ad-driven sites) is a marked worsening of a trend that has been unfolding since Google became dominant enough to impose its will on websites about fifteen years ago. Returning results in "boxed" (extended previews displayed in a box) or other inline form, to say nothing of AMP (Google 2015), already do what AI overviews are just now perfecting, which is eliminating the need to send people to the actual source that worked on producing the content extractable as a byproduct of indexing. The problem was never the indexing — **indexing the web serves the mutualistic purpose of making it discoverable, so long as it drives traffic back to the source. The problem is *substitutive use*: in effect, Google treats the need to send traffic to others as a bug which it has invested massively into fixing.**

To make matters worse, chatbot agents have figured out how to circumvent paywalls (van Ess 2025). We now find ourselves in a bind: the paywall, pillar of subscription models, creates friction for human readers while letting bots through to extract the content invisibly.

The IETF is working on an upgrade to robots.txt, AI Prefs, a declarative framework through which publishers could specify terms for AI use of their content per page or per resource. This can help publishers who cannot readily block AI crawlers unilaterally because it loses them the (weak, but existing) discoverability benefits of AI referrals without in exchange yielding any significant protection against predation (since the AI's training data already contains their content regardless of any last-mile occlusion). It is unlikely, however, that a solution built solely from "voluntary" standards, without runtime governance, can suffice. In order to gain teeth, AI Prefs will need to be coordinated with legal instruments. This will prove challenging in the current political environment, especially as bilateral trade agreements with the USA can be mobilised to prevent the intercession of legal protections in other jurisdictions.

Browsers could assist in this empowerment of publishers against predation by crawlers and non-human agents. As AI assistant functionality potentially migrates into browser interfaces, browsers can act responsibly. **Browsers committed to principles of mutualistic web ecology could, by design, prefer links to answers, help route compensation to publishers when their content is substitutively used at runtime, and surface provenance to users in ways that maintain the web's foundational referral function** even within AI-summarised responses. This kind of decision — which given Google's strategy we cannot expect to see taking place in any mainstream browser under today's funding model — is exactly the sort of large-scale coordination that a funding mechanism with built-in runtime governance would be able to deploy.

In many ways, these failings belong not solely to search engines but more broadly to the very notion of an *open* web. "Open" implies unfettered access, ungoverned use. It's a *libertarian*



arrangement, not a *democratic* one. In an open system, any emerging control point will accrue power, and the difficulty of coordinating other actors will make it impossible to counterbalance that power (Clark 2012). **As Laurens Hof aptly put it: “The result is a community that has developed exceptional sophistication about technical architecture and individual rights while remaining largely inarticulate about collective governance”** (Hof 2026). **Mutualism and the commons both subsist on a distributed and ideally bottoms-up enforcement of structuring protocols, not the kindness of anyone’s heart.** The way forward requires this sort of coordination as it alone has the ability to adapt to large-scale systemic shifts. Search that doesn’t kill the web is possible, but because search is by its nature *centralising*, we can only get mutualistic search if we govern it under collective rules, as shared infrastructure.



How Much Can A Browser Cost?

A reasonable bystander might expect that the cost of maintaining the web's most critical infrastructure would be a well-established figure, the kind of number that appears in budget documents, gets debated in governance forums, and informs policy decisions about the digital sphere's future. That reasonable bystander might, perhaps even without excessive *naïveté*, expect that the people who earnestly declare that the web is “*for all humanity*” and “*designed for the good of all people*” (Wilson and Çelik 2025) wouldn't dare make such grandstanding statements without at least a modicum of public accountability.

The reasonable bystander would, however, be wrong. As it turns out, the web community does not know, to within multiple orders of magnitude, how much a browser or browser engine costs to develop and maintain. This is a symptom of the private government of the as-yet cyberlibertarian web. A public resource, designed for and maintained to ensure the good of all, could simply never exist under the opacity documented throughout this report. **Based on how its most fundamental structures are governed, the web isn't a commons but rather *private property* that the public is granted access to because it's profitable (for now) to grant it to them, less a commons than an easement.**

So, What's The Tab?

The state of the art reveals that every cost estimate from single-digit millions to multiple billions has its defenders, and each pricing level has a methodology to support it, even if it's back-of-the-envelope. Summarising the estimates:

Estimate	Scope	Annual Cost
Ayala	Engine only (Servo, feature-parity)	~€9 million



Tarakiyee	Engine (Servo)	€50–70 million
Cooper	Engine + browser (Gecko + Firefox)	~\$400 million
Ehrenberg	Engine + browser (Chrome engineering)	~\$1 billion
Kardell	All three engines combined	~\$2 billion
Russell	Engine + browser (Chrome, full operating cost)	\$3–4 billion

This variation can be explained through several identifiable factors. There are stark differences in compensation between a massive, RSU-heavy Silicon Valley monopoly and a nimble European cooperative like Igalia. Aspects of browser operation that sit outside core engine development (from the Safe Browsing database to sync servers, to marketing and distribution) add substantial costs that some estimates include and others exclude. And the degree of organisational overhead varies enormously between a five-person focused team and a multi-thousand-person division within a trillion-dollar corporation.

The most important conclusion, however, is not about the precise number but about the governance implications. **The funding level required does not necessarily determine which levy reform ought to be chosen.** Given a WISE-like system that formalises the levy but governs it in the interest of the web's users, transparency into funding needs be imposed as a condition of participation. And even assuming substantially depreciated revenue due to the loss of monopoly pricing in search advertising following the reform of the system, the revenue potential remains comfortably in the billions. The bulk of this huge annual revenue would be more than enough to cover the full cost of multiple engines and browsers, with ample margin for investment in standards work, testing infrastructure, accessibility, and support for the broader web commons.

Workshop on Browser Funding

In the process of developing this report, I organised a workshop on "Browser Funding After Search Deals" in June 2025 at the Web Engines Hackfest, which is the primary gathering of the browser engine community. The workshop convened engineers and leaders from across the browser ecosystem to examine what would happen to browser funding if the current search default payment system were disrupted, whether by antitrust remedies, regulatory intervention, structural reform, or other *forces majeure*.

The key finding was the one noted above: the community itself — including engineers working for browser vendors — lacks a clear understanding of browser costs. Participants disagreed not merely on precise figures but even on the correct order of magnitude, reflecting fundamentally different assumptions about what counts as a *necessary* cost, what level of engineering investment constitutes "*sufficient*" browser development, and how to account for *shared infrastructure* such as the Web Platform Tests (WPT) environments, browser sync servers, the Safe Browsing database, and security monitoring. The discussion made clear that greater accountability is needed. Having infrastructure serving over five billion people managed under financial obscurity has a further chilling effect on novel browser engines that struggle to know how to account for necessary costs and presents a challenge for discussion about potential public funding: asking a government for €10 million and asking it for €4 billion are not the same task.



The Ehrenberg-Russell Estimate

Daniel Ehrenberg, a highly experienced browser developer, offers a back-of-the-envelope estimate for Chrome's cost based on engineering headcount: somewhat fewer than 2,000 engineers at an average fully-loaded cost of \$500,000 per year yields approximately \$1 billion annually (<https://bsky.app/profile/littledan.dev/post/3lbdqmgjt362m>). Alex Russell, a former Chrome engineer, suggests that this figure significantly underestimates the true cost by omitting marketing, distribution, and infrastructure expenses, suggesting that the actual operating cost is likely three to four times higher, putting Chrome's total cost in the range of \$3–4 billion per year (<https://bsky.app/profile/infrequently.org/post/3lbfngxw3pcfm>).

They note, however, that revenue as can be inferred from Google's payments to other browser vendors for default placement would be very large relative to even these elevated cost estimates: Google pays Apple approximately \$20 billion annually to be the default, and Safari holds roughly 19% of the global browser market. That implies roughly \$1 billion in search default revenue per percentage point of market share, or on the order of \$100 billion that could theoretically flow annually to web infrastructure across the entire browser ecosystem.

Evidently, if monopoly pricing in search advertising were eliminated (which is desirable) that figure would drop substantially. But it can drop by a couple orders of magnitude before it ceases to be large enough to support massive improvement across the digital sphere.

The Kardell Estimate

Brian Kardell, a long-time web standards expert and developer advocate at Igalia, has produced a widely-cited estimate of browser engine costs. His analysis concludes that maintaining all three major independent browser engines (Blink, WebKit, and Gecko) requires approximately \$2 billion per year in aggregate (Kardell 2022).

This figure attempts to capture the full cost of engine development, including the standards work, testing infrastructure, and platform integration that engines require but that is often invisible in other estimates.

The Cooper Estimate

Alissa Cooper, a leading expert in both internet standards and tech policy, has produced a **detailed analysis of browser economics, encompassing both the revenue and cost sides of the equation** (Cooper 2025a). **On the revenue side, she finds that Chrome generates an astonishing estimated \$17–35 billion per year in advertising revenue for Google in the United States alone. On the cost side, she estimates that Google spends at least \$1–2 billion per year on Chrome development, while Mozilla spends slightly less than \$400 million per year on both the Gecko engine and the Firefox browser.**

Cooper's analysis argues that since Firefox is a competitive browser that provides a fully functional browsing experience, its funding level can be considered sufficient. **If \$400 million per year is enough to maintain an independent browser engine and a competitive browser, then the multi-billion-dollar figures associated with Chrome reflect not the intrinsic cost of browser development but the overhead of operating within the inefficiency of a**



monopoly's corporate structure. Cooper counterintuitively concludes that the web could do perfectly well even if Google were forced to divest Chrome, although critics point to the importance of shared infrastructure like standards, testing, etc. being significantly undercounted in these estimates and currently borne heavily by Google (Cooper 2025b).

The Tarakiyee Estimate

Tara Tarakiyee, who works on open source funding issues at the Sovereign Tech Agency in Germany, approaches the question from **an entirely different angle, looking at the empirical productivity of a small, focused engineering team** (Tarakiyee 2025). They examine how much improvement in Web Platform Tests (WPT) pass-rates the cooperative Igalia has been able to achieve in the Servo browser engine with just five engineers. Zooming out from that observed productivity to the full scope of engine development work, their analysis arrives at **a ballpark figure of €50–70 million per year to develop a competitive browser engine.**

The Ayala Estimate

Dietrich Ayala, a veteran of browser issues and formerly of Mozilla, has developed a methodologically interesting approach to browser engine costing, grounded in a data-centric measure of feature development throughput (Ayala 2026). Rather than starting from headcount or budgets, Ayala begins with a concrete, measurable unit of output: Baseline Widely Available (BWA) features. A BWA is defined as a feature that is “well established and works across many devices and browser versions. It’s been available across browsers for at least 2½ years (30 months).”

Ayala measures two things: how many new BWA features are added to the web platform each year (approximately 52), and how many of those features the Servo engine manages to implement at its current funding level (22). From this, he calculates that implementing 52 BWA features per year in Servo — which is to say keeping pace with the web platform's evolution — would require approximately 44 full-time engineers. At a cost of €200,000 per headcount, this yields an annual cost of approximately €8.8 million.

This covers only the engine and the per-engineer cost assumption is conservative (and very low by US standards). Nevertheless, the feature-based methodology offers a benchmark **grounded in measurable outputs.** It makes it possible to ask not *"how much are we spending?"* but *"how much does it cost to deliver a given level of web platform capability?"*

Solutions

This section primarily looks at solutions regarding funding sources or in some cases solutions addressing specific incumbent problems. How the funds are then potentially deployed and governed (in cases where they don't go to browsers or browser engines directly) are tackled in subsequent chapters.

The primary concerns driving these solutions are:

- addressing issues detailed above, though of course not necessarily all at once;
- maintaining or improving current funding levels while possibly diversifying funding recipients;
- establishing a democratic system to ensure lasting healthy governance of the levy and ideally inspiring other domains of the digital sphere to follow suit; and
- if possible, improving existing capabilities.

Voluntary Donations

The simplest conceptual alternative to the current levy is to fund browsers and browser engines through **voluntary contributions**. Under this model, search engines, which have strong commercial incentives to want the web to function well, would be among the primary contributors — but so too could other beneficiaries of the affordances of web infrastructure: large publishers, e-commerce platforms, or governments. Some dedicated form of institution would be required to apportion the funds in a fair and accountable manner, ensuring that contributions flow not only to the most visible browsers but also to the engine development work that underpins them and to the smaller independent browsers that maintain competitive pressure in the ecosystem.

One mechanism that could strengthen this approach is a shift in engine licensing toward what might, if it came into existence, be termed *Commons Source*. The concept of Commons Source licensing draws on Elinor Ostrom's work on commons governance, which demonstrates that sustainable commons management **requires the ability to define boundaries, monitor use, and sanction free-riders** (E. Ostrom 2019). By contrast, the open-



source model, as currently practiced, is fully open-access: anyone can use, modify, and redistribute the code without (much) obligation. This is not, however, how commons typically function since a commons “requires specific institutional conditions, including clearly defined boundaries around the resource and its users, proportional distribution of costs and benefits, collective choice arrangements that give participants a voice in rule-making, accessible conflict resolution mechanisms, and recognition by external authorities of the community’s right to self-govern” (Hof 2026). **A well-governed commons assumes the ability to gate access to appropriators and to enforce rules around the shared management of the resource. Commons Source licensing would require that applications or browsers wishing to use an engine contribute back in some prescribed way: perhaps X number of significant code contributions for every Y users, financial payments, participation in governance, or some squishier metric like exemption from payments as long as on is “in good standing with the Maintenance Board”.** This would address the free-riding problem identified in the previous chapter, in which browser vendors adopt an engine, collect levy revenue, and invest nothing in engine maintenance.

Tax policy in various jurisdictions could further regularise and incentivise voluntary contributions. Donations of money or engineering work to a commons-maintaining organisation could qualify for tax breaks, analogous to those available for contributions to scientific research or cultural heritage. In some locales, this could run into issues when tax law has a restrictive interpretation of charitable purposes which tends to exclude organisations focused on developing or maintaining digital infrastructure, on the grounds that software development is assumed to be a commercial rather than charitable activity. The implicit assumption embedded in such tax laws is that the governance of digital infrastructure is a private commercial matter rather than a public interest concern. This assumption needs to change.

The principal downside of the voluntary approach is that infrastructure is known to be chronically under-provisioned when funding depends on goodwill rather than obligation (Frischmann 2013). Open source remains full of critical projects sustained by a single burnt-out maintainer, and there is little reason to believe that voluntary payments would reach the level of funding required to maintain the three browser engines we currently have, let alone grow that set.

It is also unclear how search defaults would operate under a purely voluntary system. Absent a structured mechanism for selecting defaults, the same market dynamics that produced the current concentration would likely reassert themselves. We might be able to get (some) browser funding from a voluntary system, but the countless problems caused by the concentration of the search market would persist.

Fiduciary Obligations for User Agents

A complementary approach that addresses the incentive misalignment at the heart of the levy rather than the funding mechanism directly is to impose legally enforceable fiduciary duties on **user agents** (including voice/chat browser and other unconventional browsers). The idea is that software acting on behalf of a person in their interactions with online services should be bound by duties of loyalty, care, and confidentiality analogous to those imposed on other fiduciaries such as lawyers, doctors, or financial advisors (Berjon 2021b; Berjon and Yasskin 2025).



The relevance to the levy is that if fiduciary duties were imposed by law, a browser would incur significant liability (even class-action liability) if it were to default to a search engine whose practices demonstrably conflict with users' interests (e.g. through deceptive ad presentation or invasive data collection). It simply would not be able to surveil all online activity in the way that Chrome currently does. A fiduciary browser could still receive revenue from search engines, but the terms of that relationship would be constrained by a legal obligation to put users first, rather than by the current arrangement in which the search engine's commercial leverage overrides the user's interests. This framework extends naturally to AI agents, which hold comparable positions of trust. This makes fiduciary duties a forward-looking intervention as well as a remedy for present dysfunction.

A fiduciary approach is useful enough that it could be implemented independently of levy reforms, and it would meaningfully address several of the problems documented above such as the deployment of deceptive design anti-patterns or the search engine's ability to veto browser privacy features. However, it does not solve the funding problem. A browser bound by fiduciary duties still needs revenue, and **if the only available revenue comes from search engine default payments, the fiduciary obligation creates a tension without resolving the underlying economic dependency. Fiduciary duties are therefore best understood as a necessary complement to structural funding reform rather than as a substitute for it.**

Retroactive Public Goods Funding

For completeness, it is worth noting the model of Retroactive Public Goods Funding (RPGF), which has been explored primarily in the blockchain ecosystem by projects such as Optimism. The core idea is that public goods are funded *after* they have demonstrated value rather than before, on the theory that retroactive assessment makes it easier to identify and reward genuinely useful contributions, particularly structural dependencies. This inverts the traditional grant model, in which funders must predict which projects will succeed, and instead allocates resources based on observed impact.

The primary innovation compared to voluntary donations is that retroactive assessment mechanism. However, this assumes that projects can bootstrap without funding and front all the costs of development and operation without any payment certainty, which is problematic at any given time and particularly poorly matched to the scale of browser engine development, which requires sustained multi-year investment of significant size. What's more, there is, to date, scant evidence that RPGF produces meaningful long-term funding at infrastructure scale, or that the community-driven allocation mechanisms consistently direct resources toward the most critical infrastructure rather than toward the most visible or fashionable projects. Some form of RPGF or retrospective evaluation may well have a role (for instance, as a way of prioritizing feature development or in-kind contributions to a browser project in the public interest), but it is my position that it should not be mistaken for a credible *standalone* solution for browser engine funding.

Global Levy for an Open Web (GLOW)

Another approach would be to establish a formal, transnational levy dedicated to financing open web infrastructure, the Global Levy for an Open Web or GLOW. The concept draws on precedents in international development finance, where solidarity levies have been suc-



successfully deployed to fund global public goods. The most prominent example is the airline ticket solidarity levy, introduced by France in 2006 and used as a microlevy to support global health initiatives operated through UNITAID. The levy has raised over €4bn since its inception (Bertrand et al. 2023). Similar mechanisms have been proposed or implemented for financial transactions, carbon emissions, and telecommunications.

GLOW could be justified on the basis that the web generates enormous global economic value and that the infrastructure enabling this value creation is a global public good whose maintenance costs are trivially small relative to its benefits. **The levy could be applied at any of several points: on search advertising revenue, on digital advertising transactions more broadly, or as a small surcharge on e-commerce facilitated by web infrastructure.**

Such an approach might also align with the long-standing aspirations of the New World Information and Communication Order (NWICO), a movement that has long advocated for a more equitable global information system in which developing nations are not merely passive consumers of information infrastructure designed and governed by the global minority (MacBride 1980).

The challenges, however, are substantial. Any transnational levy requires international coordination, and ensuring that the resulting funds are governed with the interests of people everywhere in mind and not just the narrow community currently involved in browser development. A GLOW is worth pursuing as a long-term strategic goal, but it may be tactically easier, in some cases, to convince countries to set up sovereign infrastructure directly, orthogonally to questions of funding models and long-term governance.

A comparable levy, initially restricted to the EU, has been proposed for LLMs to pay for the content that they use in training (Mensch 2026). While it is interesting to see the idea of levies being supported in mainstream publications and opening up to a wider debate, the details are too scant and the way in which the funds would benefit actual content creators (particularly smaller ones) is too unclear.

The League of Digital Infrastructure

Mark Carney's 2026 Davos address (Carney 2026) centred on a key diagnostic assessment: *rupture*. We are traversing a phase in which the international world order is being rewritten, and its new rules need to be reinvented. This is particularly true for digital infrastructure which has found itself at the centre of this shift in how power is used. These systems operate as private government and have aligned themselves with American geopolitical interests.

This rupture creates opportunity: what if a coalition of middle powers committed to building, co-owning, and governing shared digital infrastructure as a matter of collective sovereignty formed a *League of Digital Infrastructure* with which to build durable digital independence.

Such a league would likely focus on “middle powers”, states that are too consequential to ignore but too limited to independently impose their will on the global order. What the rupture has made clear is that middle powers face a choice. The two dominant players — the United States and China — are treating digital infrastructure as an instrument of power and middle powers find themselves needing to decide if they align with a power, go it solo, or forge a path forward together.



Alignment is a costly option that sacrifices independence for short-term benefits at the whim of at-times unpredictable leaders. Going it alone is particularly challenging: it implies developing a largely alternative stack (even if parts can be reused) on separate infrastructure and finding a way to drive adoption at home. The third option of forming a coalition has its difficulties in that it requires sustained coordination but it's also the most promising and most realistic for countries that prize their sovereignty.

In addition to geopolitical transformation, a recurring concern is the maintenance of democracy. This makes any kind of international league difficult to maintain as it has to deal with the very real possibility that one or several of its members may defect to fully or partially (“competitive”) authoritarianism (Levitsky and Way 2020). Paradoxically, this is an argument *in favour* of these countries aligning their approaches to digital infrastructure coöperatively, rather than an argument for skepticism of such coalitions. In *Meteorology as Infrastructural Globalism* Paul N. Edwards recounts what he calls “the first WWW” — the World Weather Watch, a “global network for the automatic collection, processing, and distribution of weather and climate information for the entire planet” (Edwards 2006). Edwards centres his study on the concept of **infrastructural globalism: the process by which shared technical systems created international cooperation that long outlasted the political and economic arrangements that sponsored them — not through idealism but rather technical needs and commitment that exceeded any country's capacity.**

Infrastructure creates its own political logic and its value compounds over time: shared systems raise the costs of defection, common standards generate network effects, and institutions formed around technical coordination develop constituencies and durability that political agreements and trade treaties alone cannot. The conclusion to draw here is that today's democracies that find themselves backsliding should be the first to invest in such infrastructure with well-established transnational democratic governance as a backstop against their own future authoritarianism.

Such an approach has for instance been advocated for in the context of AI in the **Airbus for AI proposal** (Tan et al. 2025), though not from the perspective of protecting democracy *per se*. Airbus was born of the recognition that no single European economy could sustain a competitive aerospace industry against American incumbents operating at continental scale with massive subsidies. The solution was **joint ownership of productive capacity**. Critically, this made Airbus innovative and imbued certain resiliencies to its business model, which should *not be caricatured* as a straightforward “public works” project by European states. Airbus competed in commercial markets, generated revenue, and created industrial ecosystems in each participating country, in ways more like a transnational corporation than a national investment.

This report is not the right place to develop a fully worked out vision for such a League, but several countries aligning on solutions described herein could be its first building block. Pooling these decisions would help shield the countries from American retaliation as the Trump regime predictably seeks to maintain control over the capabilities of its tech companies. **The funds driven through an Airbus-style co-owned transnational system could help support the development of infrastructure and cooperation through the risky “teething pains” period of a nascent democratic regime for the internet.** Pragmatically, having fewer regional variations to deal with (and heightened incentive alignment between participants

with vastly different budgetary stakes in the endeavor) could make distributed contribution easier for browser and search engine vendors to plan, budget, and coordinate.

Web Infrastructure Search Endowment (WISE)

The WISE proposal reforms the levy system while preserving its core economic logic that search revenue should fund web infrastructure, and addresses its structural failings by using control points to force commercial actors into conformance with collectively-governed goals. Rather than relying on voluntary contributions or new taxation, WISE would formalise the existing levy by delegating the management of browser and search choice screens to an institution **governed by relevant stakeholders, including browser vendors, search engines, the web community, and representatives of the public interest.**

Two aspects of this system need to be considered. First, the *design* of choice screens matters enormously. Experience with failed choice screens has demonstrated that poorly designed screens can be trivially gamed by incumbents or rendered meaningless by deceptive anti-patterns. A choice screen designed through expert consensus, informed by behavioural research and iteratively refined, should perform significantly better than screens created by monopolies in a sluggish feedback loop with antitrust enforcers.

For instance, a well-designed choice interface could help people select defaults based not only on experiential qualities like speed or relevance but also on non-experiential attributes such as privacy practices, sustainability commitments, or editorial transparency. These are all dimensions of the choice that a significant population of users care about but cannot readily evaluate through direct use, and are at the mercy of trademarks or word-of-mouth to assess. A rubric-based filtering system could make these dimensions legible and actionable directly on the choice screen, which would in turn make the choice feel more informed and empowered (and thus more likely to be revisited between OS installations as well as more satisfying).

It is also worth noting that designing a choice screen is not solely the interface of the one-time screen itself but also where and how it appears (e.g. just at first install vs in settings), whether icons for non-selected browsers or search engines can appear on default screens or in the “hot seat” (the OS-populated area at the bottom of a mobile screen where primary apps live), or whether “win-back” prompts in which a disregarded browser or search engine tries to make itself default again through prompting are acceptable. All of these secondary considerations add up to a series of potentially significant interventions in end-user behaviour. A WISE governing group would serve as a co-regulatory body in charge of establishing best practices across all of these surfaces, perhaps with reporting requirements and some degree of veto granted to regulators (see further details in the Enforcement chapter).

Second, and equally important, is *eligibility* for search engine or browser choice in the first place. **Being listed as an option on the choice screen at all can itself be made contingent on the acceptance of specific rules or covenants: contributions to engine development, adherence to privacy standards, compliance with accessibility requirements, participation in the governance institution, refusal of an enumerated list of deceptive interface design patterns or economic behaviours.** This transforms the choice screen from a passive list into a governance instrument. It becomes a control point at which the web community can impose conditions to ensure that the web remains healthy.



In this view, WISE would be empowered to withhold or decrease funding in response to infringement by browsers or search engines, creating an enforcement mechanism that the current *ad hoc* levy system entirely lacks (or rather, a mechanism from which only Google benefits). **By retaining the economic foundation of the existing levy while subjecting it to transparent, multistakeholder governance, WISE offers a path that is both structurally sound and politically feasible.**

Search Protocol

Perhaps the most architecturally transformative proposal is to **reconceive search not as a website but as a protocol, essentially a standardised interface through which browsers access search functionality and render results themselves, natively**, instead of fetching and rendering the SERP as any other webpage on any other domain. This is not, in fact, a new idea: early internet search had its own protocol known as WAIS (Wide Area Information Servers), a system that predated the web itself.

In a protocol-based search scenario, a search engine would expose its index and results through a standardised (high-volume) API rather than through a proprietary web interface. The browser would then be responsible for retrieving, rendering, and presenting results to the user, in the same way that browsers today retrieve and render HTML from web servers without the server controlling the display. This seemingly modest architectural change impacts multiple dimensions, particularly for chat-based and unconventional browsers.

First, it enables genuine innovation in search interfaces, since the browser controlling the user experience means drastically more and cheaper experimentation with that interface as well as incentives to put the user first. Different browsers (or different configurations of a given browser) could present results differently, experiment with new layouts, or develop specialised interfaces for different query types or different levels of trust, confidentiality, or end-user authentication.

Second, it enables multihoming: the browser can retrieve results from multiple search engines simultaneously, merge them, and present a unified view with or without surfacing to the user which engines returned which results. It could dynamically select different sources based on query analysis, routing a medical question to a health-specialised index and a product query to a commerce-focused one, as user trust and local context allow. This would be a boon to search verticals, who have suffered very directly from Google's actions.

Third, it unbundles search into its constituent functions: indexing, relevance ranking, content moderation, filtering, and payment (for the service) become separable concerns that can be provided by different actors and recombined by the user's agent. This unbundling is structurally comparable to what the AT Protocol (AT Proto 2026) promises to achieve for social media: **decomposing a monolithic platform into interoperable components and producing a system that resembles mature polycentric democratic governance rather than today's more authoritarian architecture.**

Fourth, by unbundling and decomposing economies of scale, it opens the door for search verticals to benefit from multihoming and form new, narrower economies and specialisation: for instance for academic research, local commerce, news, legal databases, personal content, or social media all integrated seamlessly by the browser into a coherent discovery experience. Such economic forecasting is necessarily speculative, but such verticals could



win back some leverage and profitability for specialised knowledge weakened by an uneven battleground for discoverability and by LLM predation in recent decades.

Lastly, a protocol-based search could become more of a reusable building-block or OS-managed function for non-browser contexts and applications. Once any application, AI agent, or device can access the same search infrastructure through the same protocol, the **distinctions between browsers, applications using browser engines to display untrusted content, and more sandboxed applications could be redrawn consequentially.**

One other consideration is that the payment dimension can also prove transformative in such a future. **Monetisation that happens today at the level of the web domain used for search would, if protocolised, happen at the browser or operating system level:** the browser and user could decide on the business model, not the search engine components that simply get paid directly by the browser for what they produce. This need not involve advertising at all. Users could pay a subscription fee to their browser, which then pays search engine components per API call, or subscribe directly to one or more search engines across application surfaces. If advertising is used, it can be privacy-centric and contextual, since the browser controls the display and can apply its own ad quality standards. Combined with a fiduciary model of the user agent, the browser would be strongly disincentivised by significant liability to present deceptive ad interfaces with manipulative interfaces. This approach can also be applied to the incentive problems that plague AI agents, such as demonetising or penalising results including hallucinations, following approaches like NLWeb that ground AI responses in verifiable web content (Guha 2026). Crucially, protocol search is compatible with and mutually reinforcing of WISE: making search a protocol, even with an interoperability mandate, does not protect against the problem of search-engine default incentives. WISE governance of choice screens and protocol-based search access are complementary reforms that together address both the funding and the architectural dimensions of the issue.

Other Control Points

David Clark's work over the years on control point analysis (Clark 2012) provides a systematic **methodology for cataloging and understanding the points of power created by specific design decisions in internet architecture.** Rather than beginning from abstract principles, Clark's approach looks at **concrete tasks (such as retrieving and viewing a web page) to identify which actors exercise control at each technological event or informational hop, what options for control they possess, and what vulnerabilities those control points create in aggregate.**

Importantly, it is not technology alone that assures the correct operation of the internet. Rather, the actors who sit at **control points must be *disciplined to be trustworthy, whether through market competition, legal obligation, or institutional design.*** Actors removed from the actual technology (e.g. governments, standards bodies) can only exercise *indirect* control, shaping outcomes by influencing the behaviour of those who hold direct power at the relevant control points.

Search and browsers are the two main control points that this document has focused on, but they are only two among several that could be used to capture funding for web infrastructure if search-focused approaches prove insufficient or politically infeasible. **Commerce platforms operate as control points for economic transactions on the web, and open marketplace**



protocols like the Beckn Protocol already demonstrate that these can be structured to support multistakeholder governance rather than private monopoly (Beckn 2026). **Social media constitutes another control point**, and the AT Protocol's decomposition of social into interoperable layers creates governance surfaces that could support other infrastructure levies and economies of scale (AT Proto 2026). **Advertising, concentrated around exchange networks, is a control point** whose current private governance has produced cascading dysfunction across the web ecosystem and could also be eligible for better governance.

These solutions are not mutually exclusive, in fact they are mutually reinforcing. Covenants for Beckn-based commerce networks, governance frameworks for the AT Protocol's social infrastructure, and shared institutions for advertising markets all make sense independently and ought to be used to finance their own respective infrastructure layers. But **they should also be connected to an improved levy governance body such as WISE and collaborate on areas of overlapping interest, potentially co-financing projects that benefit the entire web ecosystem such as browser engine development, accessibility standards, or privacy-preserving technologies.** The ideal solution, informed by the literature on commons governance, is not a single monolithic institution but **a federation of infrastructure governance bodies, each managing its own domain while participating in shared funding and coordination mechanisms for the commons they all depend upon.**

Deliberate Activism

An alternative to the current levy cannot succeed as a purely technocratic system that confines itself to managing choice screens and distributing levy funds. The problems it addresses don't exist in isolation as surprising malfunctions in an otherwise well-functioning system; they are aligned with the general state of the digital sphere. Confronting this reality requires engaging with a deeper question about how liberal democracies relate to infrastructure power. Liberal political theory has traditionally emphasized state *neutrality* across a range of societal dimensions — speech, religion, enterprise — intervening only to prevent specific harms. This posture of principled non-interference, supplemented by occasional corrective regulation, has served tolerably well in contexts where the state is meaningfully more powerful than the private actors it oversees. But when the private actors in question are transnational infrastructure providers whose platforms mediate the information flows, economic transactions, and social interactions of billions of people, the power differential between state and actor is radically reduced and at times inverted. Under these conditions, neutrality does not produce neutral outcomes: it simply supports a status quo in which the most powerful make the rules.

Johannes Thumfart argues that, **if they wish to maintain liberal democracy in increasingly illiberal times, democratic states must shift their political orientation from *neutrality* to *neutralisation*.** Instead of a passive stance that assumes fair outcomes will emerge from non-intervention they must actively counterbalance concentrations of power that threaten democratic self-governance (Thumfart 2024; Bagg 2024). The distinction is that while neutrality is a principle about what the state refrains from doing, neutralisation is a practice concerned with what the state must actively do to maintain the conditions of democratic agency, namely dispersing power.

Applied to web infrastructure, this means that **democracies cannot simply regulate from a distance, issuing periodic fines and modest, negotiated remedies. They must proactively**



create conditions that structure the digital sphere so as to support self-governance. A WISE-like institution would be one such act of neutralisation: it does not dictate what browsers or search engines should do, but it forces them to be governed in the interest of the public web rather than of just a couple of tech monopolies.

But such a change cannot stand alone. **Today's digital sphere is governed by powerful actors who will not easily make room for democracy, and any project that does so as an island will find itself under pressure,** as we can see with Wikipedia (Berjon 2025a). This means that any solution in this space must, too, take an active part in dispersing power in the digital sphere.

The institutions involved must be willing to take political positions and to advocate publicly for them. It means building alliances with organisations pursuing parallel goals in adjacent domains and being willing to support them with funding. It means being explicit about the web as a deliberate political project to further human agency, epistemic diversity, and democratic self-determination. The studied apoliticism of existing digital governance bodies is not neutrality: it's abdication.

Another Web Ahead

Finally, it must be acknowledged that if the levy system cannot be improved but proves too entrenched to reform, if browser engine diversity continues to decline, and if the parasitism of Google and browser vendors over the rest of the web continues to degrade content quality and to defund the web, then it may ultimately be necessary to consider whether the web's current architecture can sustain the kind of open, democratic information ecosystem that its founders envisioned at all — or if we might not just be better off declaring it dead for good and moving on.

One emerging splinter web is the AT Protocol, which already supports long-form writing, embedded applications, and offers a growing ecosystem of social tools architected around users rather than servers and origins. In a sense, the AT Protocol represents what the web publishing might look like if rebuilt on a healthier set of guarantees and infrastructural assumptions: content is grounded in long-lived and user-managed, app-independent identities and on-protocol data is guaranteed authentic by dint of a user-managed public key infrastructure rather than being domain-bound or CDN-managed; data is self-managed over standardised commodity storage by individuals rather than in siloed and proprietary platform data languages; and the decomposition of platform functions into interoperable (and swappable) services forming layers with stable APIs creates natural surfaces for piecemeal and resilient democratic governance. But it is too early to focus exclusively on this avenue — yet. The web remains the world's most important platform, its standards are mature and widely implemented, and abandoning it in favour of a drastically different, not very backwards-compatible platform would impose transition costs and risks. The better course is to fight for the web's renewal while keeping one eye on the alternatives that may be needed if renewal fails.



Putting It All Together

The solutions surveyed in this chapter are not mutually exclusive; in fact many are complementary. They can be woven together into a layered strategy in which each element reinforces the others, and their execution can be sequenced to avoid trying too many things at once. We are facing a major problem that also presents a generational opportunity, the goal here is to fit these solutions together into an ambitious but feasible strategy brief.

WISE as the foundation. The most actionable near-term reform is the establishment of WISE: a multistakeholder institution that formalises the levy system, governs choice screens for both browsers and search engines based on delegated powers (described in the enforcement section), and conditions eligibility for participation on adherence to rules that protect the web's users and ecosystem. WISE does not require global consensus to begin operating. It requires a single major jurisdiction willing to mandate that browser and search choice screens within its territory be managed by a democratic multistakeholder institution rather than by bilateral deals through which monopolists organise a market. Europe is the obvious candidate, although other jurisdictions could serve the same bootstrapping function. Once operational in one major market, WISE's governance framework creates gravitational pull: browser vendors and search engines seeking access to that market must comply with its rules, and the governance benefits become visible to other jurisdictions considering similar reforms. The institution can then grow not by imperial decree but by demonstrated superiority over the system it replaces.

Building fiduciary duties and protocol search over time. With WISE established as the governance foundation, two complementary reforms can be layered on progressively. First, fiduciary duties for user agents would constrain browser behaviour using levy fund distribution as an enforcement mechanism. Second, the transition to protocol-based search can proceed gradually. WISE can begin by piloting and standardising a search protocol alongside the existing website-based model, allowing browsers to experiment with multi-source search interfaces and users to experience the benefits of algorithmic pluralism. Over time, as the protocol matures and browsers develop richer search integration, we can expect the system to shift of its own accord to the superior user-experience offered by search protocol.

Advocacy and political dimensions. As explained above, liberal democracies have traditionally focused on neutrality with respect to a number of societal dimensions (e.g. speech, religion, enterprise), with limited correctives for harm. This theory of government works to the extent that we are not dealing with powerful transnational infrastructure providers. In the latter case, the power differential between state and actor is much reduced. As a result, simply being neutral with the occasional timid prodding to prevent harm will not produce infrastructure neutrality (nor will it prevent harm). We need to shift our political thinking from neutrality to neutralisation (Thumfart 2024), which means being far more proactive.

This report is in line with such a shift, and as such WISE cannot merely act as an infrastructure governance body: it must also ensure that its wider environment supports the kind of democratic governance at other layer that it embodies for search. With a stable funding base drawn from the levy, WISE can advocate for the broader ecosystem reforms that a healthy web requires, develop better technosocial arrangements outside its core scope, and help develop democratic governance for other sectors of the digital sphere.



WISE Covenant Governance

The promise of WISE is to bring the existing levy under democratic governance in order to reshape it so as to address its many shortcomings. But what governance is that? The idea of a democratic system often brings up images of a “UN of the internet” in which the peoples of Earth would all get a vote. But that is not what the situation calls for here. We can make it so that evidently public concerns — web infrastructure — are governed in service of the public without reducing representation to mechanical direct voting. **A properly democratic system is governed by a multiplicity of stakeholders with different interests, which interact as a system of “checks and balances”, that in aggregate increase people’s agency in making good choices for themselves, and allowing jurisdictional oversight to be layered in differently for different constituencies.**

This chapter does not attempt to define a complete set of bylaws or a fully specified governance mechanism. That work must be done collaboratively, within the institution itself during its bootstrap phase, through the kind of contested public debate that democratic governance requires. Rather, it touches on a variety of **design dimensions to establish feasibility and plausibility**, and that there is ample precedent for governing infrastructure of comparable complexity and importance through multistakeholder institutions.

High Level Requirements

At its core, WISE must be an accountable power structure that represents the multiple constituencies whose interests the search/browser levy affects: browsers, search engines, publishers and the wider web development community, and the public. The current system's key failure is profoundly **political**: it effectively concentrates power in the hands of Google, with some say from Apple, and excludes everyone else. WISE must invert this arrangement,



placing credible power in the hands of a much broader community rather than reserving it for the sole usage of a monopoly.

Several design requirements follow from this:

The institution must be **representative** without being paralytic. Multistakeholder governance has a well-deserved reputation for producing lowest-common-denominator outcomes when every constituency holds a veto. We need to design a voting system for represented parties that works differently for different constituencies and be cautious with how many explicit vetoes are distributed. For voting, SWIFT's method of assigning voting power based on (costly) metered participation in the system is an interesting source of inspiration, since it gives greater voice to those who use the system most (Scott 2014; S.W.I.F.T. SC 2024). In order to avoid concentrating power based in size alone, this basic distribution would need to be dampened, however. Vetos ought to be granted to jurisdictional regulators, and possibly to a sufficiently large coalition of constituencies that cannot be granted usage-based votes.

The institution must be **credible** in interactions with states, regulators, and donors. This is not an organisation to be led by engineers cosplaying governance roles. It is an entity that governs the allocation of billions of dollars and exercises authority over choice screens that mediate how (potentially) billions of people access the web. Its legal structure, financial controls, and decision-making processes must be designed with that responsibility in mind.

The institution must be **powerful** in a way that existing web governance bodies are not. It must be able to make decisions that can tackle live issues as they arise and exert the power to actually fix them rather than just talk and suggest.

WISE must learn from the shortcomings of its internet governance predecessors, drawing on the extensive body of institution-building experience that exists outside the technology sector, in domains from cooperative governance to commons management.

There is much experience to tap into. **We can think of WISE, in a sense, as Wikipedia for web infrastructure: a project of radical ambition grounded in the conviction that shared resources can be governed by their stakeholders, and that the resulting governance, which will be messy, contested, and imperfect, will nevertheless outperform the alternative of (de facto) unchecked monopoly control.**

Runtime Governance vs. Private Government

The governance model that WISE requires differs fundamentally from what the web standards community is accustomed to. Existing standards bodies practice what might be called *compile-time governance* (for the engineering-minded): they write specifications, ascertain that implementations exist, publish the specifications as standards, and then hope that nothing goes wrong in how the standard gets used. They have close to no ability to intervene if the ecosystem starts tanking, as evidenced by the absence of interventions on the part of the relevant bodies while problems of the digital sphere have fast accumulated (Cath 2023). This approach worked tolerably well in an era of competitive browser markets, where market discipline served as the enforcement mechanism (since the non-compliant will be outcompeted). It has failed catastrophically in an era of monopoly, where the dominant actor can simply pick and choose which “mandates” to implement according to their own interests.



What WISE requires instead is **runtime governance**, a term coined by the creators of the Beckn Protocol to describe the **active, ongoing management of shared concerns in a live system**. Runtime governance means not only defining the rules according to which the ecosystem operates but also monitoring compliance, adjusting rules in response to changing conditions, imposing consequences for violations, and mediating disputes between participants. It is the difference between writing a wishlist and operating a government.

Runtime governance is more powerful, and we need that power to deal with the complexity of today's digital sphere. But precisely because it is more powerful, it must be governed democratically: by its stakeholders, through processes that are transparent, contestable, and accountable. This represents a deliberate shift away from the libertarian politics that guide today's "voluntary standards" regimes, in which "voluntary" is a term that applies meaningfully only to a small number of powerful actors. **A runtime governance institution makes the exercise of power over infrastructure explicit and subjects it to democratic accountability**, which is to say that it does openly and legitimately what Google currently does privately and without mandate.

A SWIFT Precedent

The claim that a multistakeholder cooperative can govern critical global infrastructure at this scale often meets skepticism. People tend to think of cooperatives as quaint, small-scale, and fundamentally unserious. Something that runs an organic grocery stores, perhaps, but not for systems on which the global economy depends. This perception, of course, is the province of people who haven't heard about the Visa Co-operative or, even more dramatically, about SWIFT.

The Society for Worldwide Interbank Financial Telecommunication is a cooperative owned by over 10,000 member financial institutions, governed by a board elected by those members, and that operates the global messaging infrastructure through which the vast majority of international financial transactions are processed (Scott 2014). SWIFT handles billions of messages per year, operates under high security and reliability requirements, and coordinates the interests of thousands of financial institutions across every jurisdiction on earth. **It is, by any measure, one of the most consequential pieces of infrastructure in the global economy. And it has been a cooperative since its creation in the 1970s.**

SWIFT's bylaws are instructive for WISE's design (S.W.I.F.T. SC 2024). The cooperative assigns voting rights based on usage of the system, weighted so as to offset domination by collusion between the largest members. Its board includes geographic representation requirements to avoid concentration of power in any single region. It operates under the regulatory oversight of the National Bank of Belgium (its home jurisdiction) as well as the oversight of the G10 central banks, providing redundant, polycentric jurisdictional oversight. And it has managed, over five decades, to evolve its technology, governance, and membership while maintaining the trust of participants whose interests frequently conflict. This isn't to say that SWIFT is in any way perfect, but it has successfully continued to operate despite the upheavals of post-9/11 international finance oversight or the geopolitical impact of weaponised interdependence (Farrell and Newman 2024).



If the world's banks and financial institutions, who aren't exactly famous for their cooperative spirit, can run essential infrastructure this way, so can the web, which people so often consider a commons intended to serve all humankind.

Scope and Limits

Defining what WISE governs is as important as defining how it governs.

The institution's scope should encompass:

- The management of browser and search **choice screens**, including eligibility criteria, design standards, and the conditions under which actors may be listed or delisted.
- The collection and distribution of the **levy**, including the levy rate, the allocation formula between browsers and engines, and the requirements imposed on recipients.
- The development and maintenance of technical **standards** directly related to search and browser interoperability, including (over time) the protocol search specifications that would enable protocol-based search.
- **Accountability** for conditionalities: the rules that browsers and search engines must comply with in order to participate in the system, covering privacy, security, accessibility, and adherence to the priority of constituencies, in line with the expectations of jurisdictional regulators.
- **Support** the development and operation of similar institutions to govern other digital infrastructures, for instance for advertising, social media, or cloud.

Equally important is what WISE should *not* do:

- It should not govern the **content** of the web.
- It should not make **editorial** judgments about search results.
- It should not regulate the **internal engineering** decisions of browser vendors beyond what is necessary to ensure interoperability and compliance with agreed standards.

It should resist the **institutional** tendency, familiar from every standards organisation, to expand its scope incrementally until it encompasses everything tangentially related to its original mandate.

Constituencies

The work of establishing constituencies properly will need to be done during bootstrapping, but we can consider the following an adequate sketch for evaluative purposes:

- **Browser vendors.** *Bona fide* browsers who participate in the system, with a minimal user base. Voting rights ascribed either by measure of market-share or proportional to (metered) usage of the system.
- **Browser engine projects.** Browser engines used by browsers that meet the above criteria. Forks only count when they full depart and no longer take upstream changes. Voting rights derived from the browsers that use them.
- **Search engines.** Participating search engines, with a minimal user base. Voting rights ascribed by (paid) usage volume. Given a search protocol that unbundles search responsibilities to multiple entities, this constituency may be split up.
- **Regulators.** Jurisdictional oversight grants participation in the governance, at some level. Non-voting group, primarily observes and wields targeted vetoes.



- **Web builders.** Those who make things on the web. This constituency is hard to find good representatives for, but it is possible to list a number of professional organisations that could be pooled to drive a nomination pool. Non-voting, primarily observes and with a sufficient majority can veto in some cases. Alternatively, sites could be granted votes measured by click-through rates from search for non-navigational queries.
- **Public interest.** Similar to the previous one but focused on civil society figures. More of an advisory and accountability role.

Polycentric Approach

A polycentric system is one that has multiple decision centres, which is to say that decision-making and rule-setting take place at a variety of levels. Each decision centre has limited and autonomous prerogatives, and operates under and overarching set of rules. This offers a richer repertoire to work with than simply thinking in terms of “decentralisation,” which is a poorly defined notion and operates only on an axis of more or less centralised. Polycentric systems offer redundant levels of checks and balances, which makes them robust over time.

Unbundling governance and making it work through multiple related institutions, each with its own arrangement, has benefits that include clearer shared local knowledge, rules that are better adapted to local needs, higher trust and lower enforcement costs, and reliance on disaggregated knowledge, which is more diverse and resilient. Additionally, the separate-but-interacting institutions form a parallel autonomous systems: the probability of failure over a large region is decreased because a set of local autonomous systems with a variety of rules is more robust. (E. Ostrom 2005)

As Elinor Ostrom noted: “*While all institutions are subject to takeover by opportunistic individuals and to the potential for perverse dynamics, a political system that has multiple centers of power at differing scales provides more opportunity for citizens and their officials to innovate and to intervene so as to correct maldistributions of authority and outcomes. Thus, polycentric systems are more likely than monocentric systems to provide incentives leading to self-organized, self-corrective institutional change.*” (E. Ostrom 1998) This can help solve the problem of the Iron Law of Oligarchy by shifting greater responsibility to more effective parts of the institutional network when one becomes oligarchic (Berjon 2024).

Describing a full set of polycentric institutions *ex ante* is an exercise in futility, but we should be aware of the value involved in growing the system that way from the start so that we can avoid the pitfall of putting all of the responsibility in a single entity.

How to Not Break Browsers

There is a long history of failed standards, and we should often be grateful that they failed. XHTML 2.0, the Semantic Web's more grandiose ambitions, and numerous W3C specifications that never saw meaningful implementation all represent cases where the standards community's judgment about what the web needed or could successfully transition to was wrong, and developers' refusal to implement provided valuable feedback (Mozilla Corporation and Opera 2004; Hickson 2004; Nottingham 2024). The risk with creating a governance body, like WISE, that has real enforcement power is that it could impose standards of no



actual value, or worse, standards that actively harm the web's functioning. This would have happened if the W3C of the early 2000s had been given runtime governance authority.

This risk requires strong internal rules constraining when WISE may impose **mandatory** standards. **The guiding principle should be that mandates are reserved for cases where they are necessary: where interoperability cannot be achieved through voluntary adoption, where market discipline has demonstrably failed, or where the absence of a standard produces harms (to users, publishers, or the ecosystem) that competitive dynamics cannot correct.**

Market-backed standards, as are common today, can work well when the relevant market is genuinely competitive (Nottingham 2024). Under conditions of competition, vendors are disciplined into complying with standards that users want because noncompliance costs them market share. The failures of "voluntary" standards that we observe today are largely evidence the relevant markets are captured. Today's web is, structurally, a planned economy with Google's bottom line as its driving KPI. WISE seeks to open that up gradually to competition rather than take it over and have a different set of underaccountable stakeholders centrally plan it.

WISE should therefore conceive of its authority as a backstop that serves the purpose of organising infrastructure in a way that supports a highly plural and open system, rather than as the only or even primary mechanism through which decisions in the ecosystem are made.

How to Avoid Corruption and Oligarchy

Every democratic institution faces the challenge that Robert Michels identified over a century ago as the Iron Law of Oligarchy: the tendency of organisations, no matter how democratic their founding principles, to develop entrenched leadership whose interests diverge from those of the broader membership (Michels et al. 1915). WISE will not be exempt from this tendency, and pretending otherwise would be naïve. The question is not whether oligarchic pressures will emerge but whether the institution's design can slow their accumulation and create mechanisms for periodic renewal. Several design principles can help there.

First, **polycentric oversight**. Rather than relying on a single jurisdiction to impose and verify conditionalities on WISE, multiple jurisdictions can eventually do so independently, each with its own regulatory authority, its own review processes, and its own capacity to intervene if the institution fails to meet its obligations. This produces the kind of robustness that Vincent Ostrom discussed in his work on polycentricity: a system in which no single centre of authority can be captured without the others *noticing* and *responding* (V. Ostrom et al. 1961; McGinnis and Indiana University Bloomington 2002; V. Ostrom 2006).

Second, **dampened voting**. Following SWIFT's model, voting rights for certain classes of decisions can be assigned based on usage of and participation in the levy system, reflecting the principle that those with greater *active* stake in the system's operation should have greater say in its governance. But in order to prevent the largest players from dominating, voting power should be subject to a form of quadratic decay, i.e. you need four times the market share to obtain twice the votes. (Note that this is *not* quadratic voting, where votes are *counted* quadratically). This ensures that scale confers influence without conferring proportional control, and that coalitions of smaller actors can still outvote large players. Geographic representation requirements, again following SWIFT's precedent, can further prevent power



concentration in any single region, dampening upstream power dynamics inherent to global wealth distribution.

Third, ***term limits and rotation***. Board members and committee chairs should serve limited terms with mandatory rotation, preventing the accumulation of institutional knowledge in a small number of individuals who become irreplaceable and who, in becoming irreplaceable, become unaccountable.

Fourth, ***protected spaces for dissent***. It is well-documented that beliefs conflicting with the interests of well-connected participants spread less readily through organisational networks as a general rule (Allen et al. 2019). WISE should create formal mechanisms for minoritarian perspectives (e.g. ombudspersons for sensitive conflicts and whistleblower protections) that ensure heterodox views reach decision-makers even when mainstream, influential participants would prefer to ignore them.

Fifth, and most fundamentally, ***cultural intentionality***. The Iron Law of Oligarchy operates most effectively when organisations forget that it exists. WISE's founding culture must include an explicit, ongoing commitment to recognising and managing oligarchic tendencies, not as a one-time design exercise but as a permanent institutional practice. This means auditing itself, reporting on power distribution and persistence in positions of power, and cultural dedication to the norm that challenges to institutional authority are healthy and routine.



Enforcement

There is an important distinction in institutional theory between *rules-in-form* (what an institution's documents say) and *rules-in-use* (what actually happens when people interact with the institution's authority) (E. Ostrom 2005). People instinctively understand that: if the official driving speed limit in an area is 50 but the locals know that you never caught under 60, then the speed limit *in use* is really 60. It's the one that matters. Institutions are real when people believe that they are, and a strong way of maintaining that belief is by **enforcing the relevant rules consistently, visibly, and consequentially**.

This is one of the reasons why tech ethics initiatives have failed to deliver meaningful change: they produced rules-in-form (principles, frameworks, guidelines) without any mechanism to make those rules operative in the world. It also explains why standards organisations, absent effective runtime governance, fall so far short of the mark in digital governance: they can specify how systems *should* work but cannot intervene directly or effectively when systems *don't*. And it is why this report takes a more pragmatic approach, grounded in what we understand about how governance actually functions rather than in how we might wish it functioned.

This chapter addresses the practical dimensions of making WISE's authority real: **how to bootstrap the system, how to grow it across jurisdictions, which regulatory instruments can underwrite it, and how to design the choice screens that serve as its primary enforcement mechanism**.

Implementation Stages

Bootstrap

Getting WISE off the ground will require a bootstrap phase during which many of the early operational details are worked out in practice. The first jurisdiction to support the system will need to adopt a relatively hands-on approach by defining acceptable initial ground rules, ensuring that the institution operates within its mandate, and providing the regulatory backstop that gives WISE's authority teeth.



This bootstrap phase is less daunting than it might appear, for one critical reason: the levy already exists. WISE does not need to create a funding mechanism from nothing; it needs to **bring an existing one under democratic governance**. The groundwork exists. What is missing is the institutional structure to govern the levy in the web's interest credibly.

Self-Enforcement and Growth

Once operational, WISE has a clear toolbox for enforcement strengthened over time by a virtuous cycle. **Participants who intend to appear on choice screens and to receive levy payments must remain within the system and comply with its rules. Departure or noncompliance means losing access** to the distribution mechanism that delivers users and revenue. Revenue can be withheld, and exclusion from the system would be commercially devastating.

A significant challenge lies in growing across jurisdictions that may have different expectations, legal frameworks, and political priorities. WISE's governance design must accommodate this variation through subsidiarity principles that enable regional flexibility. Again, SWIFT's experience in navigating the regulatory expectations of over 200 countries provides both precedent and practical lessons for managing this complexity.

The Digital Markets Act (DMA)

Europe's Digital Markets Act is the logical regulatory instrument with which to bootstrap WISE. The DMA was designed precisely to structure digital markets so that they operate in alignment with societal needs, and it already contains provisions for browser and search choice screens that map directly onto WISE's proposed governance scope. The infrastructure of enforcement to back WISE with — gatekeeper designations, conduct requirements, the capacity to impose fines — is already in place.

The DMA's application, unfortunately, has so far been underwhelming. Its choice screen mandates have been implemented through bilateral negotiations with the very monopolies they are supposed to constrain, proceeding at a pace that meets neither the needs of the market, the democratic urgency to reform the digital sphere, nor the promises made by the Commission during the DMA's development. The work is carried out largely behind closed doors, in ways that keep gatekeepers in the driving seat and that are devoid of meaningful public accountability (OWA 2025).

Expert technology policymakers have been arguing for some time that DMA enforcement would benefit from a more open and collaborative approach. There is a strong case that the Commission going it alone is not working in large part because despite highly competent employees it lacks the technical bandwidth, the granular market knowledge, and the stakeholder relationships needed to design effective interventions in a system as complex and opaque as browser and search distribution. The team working on DMA enforcement is also much too small to have any meaningful impact on their own within reasonable time frames, no matter how hard they work (Berjon and Crider 2025). The Commission would fare better if it *organised* its own openness rather than acting as an ivory citadel (Cattan and Toledano 2022). A **co-regulatory** approach, in which the Commission delegates operational details to a bootstrap WISE group while retaining oversight authority and enforcement power, would strengthen the DMA's impact at no additional cost. The Commission provides the legal mandate and the backstop of fines; WISE provides the technical expertise, the accountable



multistakeholder legitimacy, and the capacity for iterative refinement that a regulator with a huge intervention surface inherently lacks.

Isn't a Covenant of This Kind Illegal?

Some objections naturally arise: can competitors legally band together to govern shared infrastructure without running afoul of competition law? Can a public-interest body of this kind be granted authority to structure a market? Before answering these questions, it is worth noting that it would be difficult for WISE to be *more* problematic than the cartel-like arrangement that exists today, in which a single company pays billions to organise a market and foreclose competition (an arrangement that, notably, a US federal court has already found to be an illegal monopoly maintenance scheme).

That said, the concern is not frivolous. **Article 16 of the EU Charter of Fundamental Rights guarantees the freedom to conduct business, and any cooperative arrangement among competitors must be structured to avoid deleterious anticompetitive coordination.** Several precedents, however, demonstrate that such structures can be both legal and societally beneficent, and justify the risks of collusion and ossification. Creating some constraints on the exercise of business can increase the freedom to conduct business by preventing more constrictive alternatives — of the kind that dominate today's digital sphere — from emerging.

SWIFT, which is a cooperative of otherwise fierce competitors, is one example. Its members coordinate on messaging standards, infrastructure investment, and operational rules without this coordination being deemed anticompetitive. On the contrary, the infrastructure SWIFT provides is widely understood to be pro-competitive in that it lowers transaction costs for all participants and enables smaller institutions to access the same global network as the largest. **Infrastructure is often considered to be particularly prone to monopolisation, and is typically regulated to prevent the negative effects of monopolies. This can be challenging to apply in a transnational context, however, as is the case with SWIFT and WISE. A cooperative with regulator oversight can provide a comparable model.**

Standards organisations provide another well-established setting in which competitors coordinate with carefully managed and mitigated antitrust liabilities, on the principle that interoperability standards benefit consumers and promote competition even when they require cooperation among rivals. As I have argued above, one necessary but fragile condition of interoperable standards is that the underlying market they are co- or self-regulating needs to remain healthy and non-monopolistic, as has not been the case in the browser and search markets. Should a SWIFT-style structure restore such competition, one of the key beneficiaries would be the W3C and Mozilla, two organisations currently hamstrung by dynamics created by the status quo.

There is also ample precedent for cooperatives with varied stakeholder classes who hold different rights (W. F. Whyte 1991; Warren et al. 2025; Fajardo-García 2017). The most visible examples are supermarket cooperatives in which workers and customers are both stakeholders, or credit unions in which borrowers and savers coexist under the same governance structure. Cooperatives have shown themselves to be applicable to a wide class of needs, such as utilities, commodity clearinghouses (particularly for highly portable and thus global agricultural commodities), or the Visa settlement network (Schneider 2020).



WISE sits at the intersection of a SWIFT-style infrastructure cooperative and an open standards organisation, overseen by regulators. This intersection does not undermine either component, however, and should be seen as additive and complementary to them. In fact, SWIFT itself functions as a standards organisation, producing standards governed under cooperative rules, in direct support of their mission. WISE is not merely permissible but desirable. It is precisely the kind of market-structuring intervention that the DMA was designed to enact.

Aren't Choice Screens Bad?

There is extensive precedent for bad choice screen programmes, but these have to be understood in context. In many cases, the details of the choice screen had been delegated to the monopoly that they were supposed to undermine. Enforcement was, at times, bad to the point of embarrassment such as when the Microsoft Windows browser choice screen, mandated by the European Commission, was infamously broken for fourteen months without anyone noticing. Google and Apple's compliance with DMA choice screen requirements has been grudging and minimalist. In a number of ways, the DMA's enforcement of its choice screen mandates is following in those unfortunate footsteps: proceeding bilaterally with monopolies, keeping them in the driving seat, and operating largely without meaningful public accountability.

Despite these shortcomings, however, DMA choice screens are having a measurable impact even in their current imperfect form. As reported by Mozilla, "Firefox saw an uptick of 111% daily active users in France and 99% in Germany in the first 12 months of the DMA - despite initially poor compliance" (OWA 2026). Evidence shows demonstrated improvements in user satisfaction when people are given better search choices and when the chosen app gets into the "hotseat". These results suggest that the problem is not with the concept of choice screens but with the quality of their design and the rigour of their enforcement.

Mozilla's research on choice screen design provides a strong and systematic evidence base for how to do this well (Akesson et al. 2023). The findings are striking and people *really* liked having a choice screen and preferred it to be more informative: "98% of people stated that they wanted to be shown a choice screen, with most preferring the screen with more information and greater number of browsers." This runs directly counter to the industry narrative that choice screens are an unwanted imposition that confuses users. People want to choose, but they simply want the choice to be **meaningful** rather than performative.

Building on this evidence, a WISE-governed choice screen could incorporate several design principles:

Regular recurrence. Choice screens should be shown not only at device setup but periodically (perhaps annually) or triggered by contextual nudges that remind people that alternatives exist, for instance when they perform more than a few searches without clicking through to any result.

Ubiquitous availability. The choice should be accessible at device setup time (which Mozilla's research identifies as the option preferred by most people) but also easily reachable in settings, systematically and predictably present everywhere that the relevant decision can



be made. A choice that is technically available but buried five menus deep is not a choice at all.

No win-backs. The monopoly platform must not be permitted to use deceptive patterns, prompts, or degraded functionality to steer users back to its preferred default after they have made an alternative choice.

Informative non-experiential attributes. The choice screen should provide information about qualities that users cannot readily evaluate through direct experience. For instance:

- **Privacy:** For the browser choice screen this could be a filter requiring meeting a minimal bar and for search engines, it could be a ranking or filtering rubric about whether the provider forgets all data, keeps data only if opted in, tracks continuously, etc.
- **Sustainability:** A rubric of levels to filter on or a metric to rank choices on.

Additionally, we can set some user preferences directly on the choice screen such as whether they wish the browser to send an automated privacy signal (such as the Global Privacy Control, or GPC) or whether they want to preclude the search engine from showing them AI summaries.

A choice screen designed through an accountable expert process within WISE, informed by behavioural research and iteratively refined based on measured outcomes, can perform dramatically better than screens designed by monopolies in a sluggish feedback loop with understaffed enforcement agencies that can be prone to political timidity. The evidence already demonstrates that even poorly designed choice screens produce meaningful shifts in market share. Well-designed ones could transform the competitive landscape.

Isn't Co-Regulation Terrible?

Co-regulation has, deservedly, a mixed reputation. It is frequently a vehicle for regulatory capture. The history of industry self-regulation in digital advertising, in privacy (as under the FTC's mandated self-regulatory regime), and in content moderation provides ample grounds for skepticism. The concern that WISE could become captured by the actors it is meant to govern is legitimate.

Several design features can mitigate this risk. **First, small challengers** who are currently crushed under the existing system **must be given real power within WISE through weighted representation** that reflects their importance to ecosystem diversity and not just their current market share. The quadratic-decay voting mechanism discussed in the governance chapter serves precisely this function: it limits the ability of large incumbents to simply outvote the rest of the membership.

Second, deploying WISE across multiple jurisdictions creates overlapping oversight from different regulatory authorities with different institutional cultures, different political pressures, and different susceptibilities to capture. As mentioned above, this creates polycentric robustness.

Third, WISE's internal decision-making structure must be designed to stymie the formation of alliances that would defeat the system's purpose. This means selecting constituencies whose interests are not naturally aligned and further breaking up governance concentration across geographic subdivisions where possible. Having an "ombudsperson"-

type role incentivised to monitor collusion and chilling effects contributes significantly to this goal.

Finally, the most important safeguard is public accountability. Co-regulation fails most spectacularly when it operates purely in private and only between actors who have aligned interests. WISE's deliberations must be public by default, its decision-making processes documented and contestable, and its outcomes subject to external review notably by the relevant regulators. Here as well, an “ombudsperson” is a significant preventative measure against festering intransparency. At the risk of a trite statement, this needs to be co-regulation.

Reference Implementation

Along with this report, example implementations of the browser and search choice screens have been provided on the <https://ftw.fund/> website. These reference implementations demonstrate that the design principles outlined above can be rendered usable (though I am not a designer and better designs are undoubtedly possible).

A pragmatic approach to compliance would be for participating jurisdictions to simply adopt or adapt this reference implementation, providing a concrete, testable baseline from which iterative improvements can proceed. This is not the only possible implementation, and WISE's governance processes could and should develop alternatives and refinements over time. But having a working example from day one makes compliance much easier and provides a tangible artifact around which early consensus can form.

Fund Attribution

A reformed system must answer a set of practical questions that the current *ad hoc* arrangement has never been forced to confront, or at least not in public: **How much money should be levied from search? How should it be divided between browsers and engines? And what other funding needs might the levy support?**

This chapter addresses those questions. The answers are necessarily provisional and to some degree *speculative* as the precise figures would be negotiated within the WISE governance framework rather than decreed in advance but the elements I have are sufficient to establish that the funding potential is more than adequate and that a well-governed levy could simultaneously reduce the burden on search while dramatically expanding investment in the web and the infrastructure of our digital sphere.

How Much to Levy from Search?

The exact levy rate and collection mechanism would need to be negotiated within WISE's governance framework, balancing the interests of search engines, browsers, and the public. Search engines derive their near-entire value from the existence of a functional, content-rich web, and the web in turn depends on browsers and browser engines to render that content accessible. They therefore have an incentive to contribute to web infrastructure. Of course, the literature on infrastructure and public goods shows that funding for these is almost systematically underprovisioned, but we can expect a levy level decided by a group of stakeholders would be adequate in part because not only search engines have a voice, and because they benefit from the funding too.

The current system provides a useful baseline from which to reason about feasible levy rates. Expert testimony in the *United States v. Google LLC* remedies proceedings revealed that Apple levied 36% of the revenue that Google Search generated from Safari traffic, a figure subsequently confirmed by Sundar Pichai himself (*United States v Google* 2024). This is an extraordinary extraction rate (significantly higher than corporate tax rates in most jurisdictions).



We can play with a simple model of what various levy rates could produce under a reformed system. **Total annual global search revenue is approximately \$240 billion. A 36% levy (likely a much higher level than other competitors get, but still informative) would yield over \$86 billion.** Halving that figure to account for the competitive pricing that would prevail in a non-monopolistic search market still yields \$43 billion. Reducing the levy rate to a far more modest 5% and still halving for competitive ad pricing gives us about **\$6 billion**, which is still comfortably more than even the highest estimates of the cost of maintaining all three browser engines and their associated browsers, with substantial margin for investment in standards (for reference, the W3C’s annual operating budget is about \$10m), testing, research, and governance.

Levy Rate	Gross Yield (on \$240bn)	50% Adjustment for Competitive Pricing
36% (current Apple rate)	~\$86 billion	~\$43 billion
15%	~\$36 billion	~\$18 billion
5%	~\$12 billion	~\$6 billion

There is clearly room to finance solid, thriving web infrastructure in the public interest while extracting far less money from search than the current system takes. And this calculation does not account for the positive second-order effects of reform: **money that is no longer captured by monopoly ad pricing does not disappear from the economy. It most likely recirculates through other marketing channels which massively benefits the rest of the digital ecosystem and the publishers who depend on advertising revenue to produce the content that gives search its value in the first place.**

One important note: under a protocol search system of the kind described in the solutions chapter, the browser itself collects the funds—whether from user subscriptions, contextual advertising, or per-query payments. But this only shifts where the money is collected, not the overall logic of fund distribution. Whether revenue flows from search engine to WISE to browsers, or from users to browsers to search engines with a levy extracted at the browser level, the fundamental decisions remain the same.

The Browser/Engine Split

As the chapter on browser costs showed, the variation in estimates reflects a genuine ambiguity about what responsibilities exactly fall within the browser’s remit. A reformed levy system must confront this ambiguity directly, because getting the split wrong has serious consequences. If funding flows primarily to browsers and browsers are free to adopt open-source engines without contributing to their maintenance, the free-riding problem documented earlier will persist under new management. **If we under-value important but low visibility services provided by browsers (like parts of the security stack, the Safe Browsing blocklist, moderation, etc) then we may end up indirectly underfunding essential services by undercutting browser revenue more than it needs to be.**

Determining the ideal allocation is challenging, but the WISE framework provides the institutional mechanism to resolve it: browser and browser engine projects that participate in the system would benefit from surfacing costs that may not otherwise be well known.



Other Funding Opportunities

A reformed levy need not limit itself to financing browsers and engines. The current system's failures have produced cascading underinvestment across the entire web infrastructure stack, and a well-governed levy could address several of these gaps simultaneously.

Open source security & maintenance programmes. Germany's Sovereign Tech Agency has demonstrated that effective investment in open digital infrastructure is possible at relatively modest scale. Operating with a budget of approximately €20 million, the agency funds maintenance and improvement of critical open-source projects. The model is compelling precisely because of its efficiency: targeted funding of specific infrastructure needs, governed by people who understand the technology, with scouting and minimal administrative overhead for grantees. A WISE-governed levy could **generalise this model to all jurisdictions covered by the system, establishing sovereign technology programs that address locally relevant infrastructure needs.** Scaling the STA's model even modestly could transform the digital infrastructure landscape.

Standards development. Web standards are routinely underfunded, and the consequences are visible in their quality and pace of development. Under the current model, **serious participation in standards work requires contributing to the drafting of specifications which is an expert art that demands deep technical knowledge and considerable time.** The result is that standards are slower and lower-quality than they need to be, and **dominated by the commercial interests of the tech monopolies** whose employees can be dedicated to the work full-time. (It is also worth noting that the professionalisation and cost structure of standards work generally alienates volunteer and community participants in various ways, further entrenching the market power of larger players.) **A levy-funded program of professional specification writers, whose job would be to listen to community consensus and translate it into rigorous, implementable specifications, would be an improvement and a counterweight to that bias.** Removing membership fees and actively supporting participation from underrepresented parties, geographies, and smaller organisations would also broaden the base of expertise that informs the web's evolution.

Testing. Web Platform Tests (WPT) is the critical mechanism through which the web achieves interoperability today. It's a shared suite of over two million tests against which all browser engines verify their implementations of web standards (Sender 2024). WPT is one of the largest open-source testing projects in existence, and the infrastructure it operates is *expensive*. **A reformed levy should dedicate stable, predictable funding to WPT infrastructure, test development, and the tooling that enables engine developers to efficiently identify and fix interoperability gaps.**

Long-term research and democratic governance. Looking beyond immediate operational infrastructure needs, the web would benefit from an **endowment dedicated to its long-term health.** This would support **research into future web technologies, policy analysis of emerging threats and opportunities, and the cultivation of expertise in digital governance.** Equally impactful would be funding an **organisation** tasked exclusively with developing and protecting user agency and democratic control over digital infrastructure through new **protocols, products, advocacy, and policy work.** The elimination of “political” debate from standards and web infrastructure discussions, under the guise of an imaginary “neutrality”



that treats all commercial interests as equally legitimate and all power structures as natural, is spurious and demonstrably a key driver of technical standardisation having been captured by monopolies. It's also why existing internet governance bodies lack the institutional expertise to hold the discussions needed to fix our digital governance system. Using levy funds to reposition the web as an explicitly hegemonic project in support of agency and democracy is a necessary corrective to decades of pretending that infrastructure governance is apolitical while monopolies quietly wrote the rules.

Conditionalities

No system of this kind can operate on goodwill alone. For a state to delegate governance authority over a public interest levy to a multistakeholder institution like WISE, it must impose minimum standards of performance, transparency, and accountability short of which the state retains the right to intervene directly, which is to say *conditionalities*. These **conditionalities serve as the backstop that prevents institutional capture**, the insurance policy against the scenario in which WISE's governance is progressively captured by the same commercial interests that corrupted the current system.

The precise conditionalities would need to be specified by the authorising jurisdictions, but they should include at minimum:

- **transparency** requirements for all funding recipients, including auditable cost breakdowns,
- clear guidance regarding **market structuring** versus participating in the market,
- **diversity requirements** for funding recipients, ensuring that no single entity receives a dominant share,
- **accountability** mechanisms for browser and search engine behaviour,
- regular public **reporting** on the health of the web infrastructure ecosystem, and
- **sunset provisions** that trigger automatic regulatory review if specified metrics of web health (e.g. engine diversity, search pluralism, publisher revenue trends, privacy properties of relevant technologies) deteriorate below defined thresholds.

The goal here is not to micromanage WISE from above but **to ensure that the institution remains answerable to the public interest that justifies its existence, enumerating touchpoints where intervention is justified**. The resulting system should exhibit the kind of robustness that polycentric arrangements have, such that WISE is prevented from going off course by redundant oversight capabilities from (eventually) multiple jurisdictions.

Conclusion

The model explored in this report implements a shift from the private government of the web towards one that focuses on infrastructure governed by or at least in the interest of its users. This is hardly a novel approach outside of the digital world: treating infrastructure as a public good creates downstream benefits to both businesses and citizens (Ricks et al. 2022). It promotes a competitive environment in which firms can thrive without the winner-take-all dynamics that create inferior outcomes and authoritarian governance. Put differently, the proposal simply considers digital infrastructure outside of tech exceptionalism, and offers instead a tried-and-true approach that serves us well in the analog world.

While there is no room to go beyond search and browsers in any detail here, we can readily see **how this approach could benefit other digital markets**. Even without a League of Digital Infrastructure, runtime collective governance over shared infrastructure is an idea that is being explored in other spaces. An older proposal (still being explored) suggested a comparable approach as a solution to the many ills of the advertising ecosystem (Berjon 2021a). The network built atop the AT Protocol, which brings together many different actors, is evolving towards something like a “covenant” that would allow for joint governance of the overall ecosystem (AT Proto 2026). And the Beckn Protocol has already deployed network-specific governance bodies (for instance for retail in India) and is investigating extending the model more broadly (Beckn 2026).

People the world around are increasingly recognising the limits of today’s system of private government and are reaching for the same toolbox that was previously used to solve similar problems: democracy. Funding the web by putting the search/browser market under democratic control would be a major contribution to the emerging world of hopeful, liberated, collectively-governed internet systems.

Pappus of Alexandria famously reports Archimedes stating “Give me a place to stand on, and I will move the Earth” (Jones 1986). **The company that primarily organises and benefits from the current search and browsers (and mobile OS) market is Alphabet. As of this writing, it is valued at over \$3.5T. About 90% of Alphabet’s revenue comes from Google, and in turn of Google’s FY25 \$402bn in revenue, \$294.7bn (73%) are from advertising**



revenue with Google Search driving \$224.5bn (55%) of that. This doesn't account for the level at which the search levy may fund the Android ecosystem, which brings in another \$48bn in revenue (12% of the total).

Suffice it to say, the system described in this report is critical to Alphabet's capture of digital markets around the world. (Apple benefits significantly less, but circulating the money they get from the system to a wider, more diverse, and more inventive ecosystem would have positive effects.) **At a time when some claim that the internet has become too captured to ever change, this intervention *alone* would be powerful enough (not to mention comparatively cheap) not just to establish alternative search engines, browsers, or mobile operating systems but also to fund additional democratic infrastructure as well as to recirculate a *massive* influx of value throughout the web and media.**

There is little that could be more impactful in building a better internet, today.



Acknowledgements

Many former or current employees from browser vendors and search engines large and small provided crucial information on condition of anonymity. I am deeply thankful for the time you gave me and your dedication to fixing this arrangement, even though your employers benefit from it.

Many thanks to all the participants in the lively and highly-attended browser funding workshop at Web Engines Hackfest 2025 — this document surfaces many of your views — and to Igalia for their warm welcome and outstanding organisation. Special mentions to Brian Kardell, Eric Meyer, Stephanie Stimac, Alex Moore, and Xan López for your contributions to browser funding issues and precious conversations. A shout-out as well to the indefatigable people at OWA. My heartfelt thanks as well to the participants of the KGI workshop on search remedies, and in particular to Alissa Cooper.

The following people (in alphabetical order) have provided invaluable input into this draft (albeit at times unknowingly): Dietrich Ayala, Ian Brown, Alissa Cooper, Matthew Frehlich, Aurélien Mühl, Mark Nottingham, and Max Gendler.

Many thanks to Juan “Bumblefudge” Caballero for his extensive contributions to this report.

Finally, very special and deeply heartfelt thanks to the Digital Infrastructure Insights Fund (DIIF), and in particular to Katharina Meyer, without whose funding and support this report would not have been written. I had a great time meeting and exchanging with others in my cohort.



References

- AT Proto 2026** AT Proto. 2026. *Authenticated Transfer (AT) Protocol*. <https://atproto.com/>.
- Akesson et al. 2023** Akesson, Jesper, Michael Luca, Gemma Petrie, and Kush Amlani. 2023. *Can Browser Choice Screens Be Effective?* <https://research.mozilla.org/browser-competition/choicescreen>.
- Allen et al. 2019** Allen, Danielle, Henry Farrell, and Cosma Rohilla Shalizi. 2019. *Evolutionary Theory and Endogenous Institutional Change*. https://projects.iq.harvard.edu/files/pegroup/files/allen_farrell_shalizi.pdf.
- Anderson et al. 2019** Anderson, Elizabeth, Stephen Macedo, Ann Hughes, David Bromwich, Niko Kolodny, and Tyler Cowen. 2019. *Private Government: How Employers Rule Our Lives (and Why We Don't Talk about It)*. First paperback printing. Princeton University Press.
- Appelquist et al. 2024** Appelquist, Daniel, Hadley Beeman, and Amy Guy. 2024. *Ethical Web Principles*. Statement December 12. <https://www.w3.org/TR/ethical-web-principles/>.
- Ayala 2025** Ayala, Dietrich. 2025. "What Is Baseline?" Baseline. <https://web-platform-dx.github.io/web-features/>.
- Ayala 2026** Ayala, Dietrich. 2026. "Servo Baseline Readiness." <https://webtransitions.org/servo-readiness/>.
- Bagg 2024** Bagg, Samuel Ely. 2024. *The Dispersion of Power: A Critical Realist Theory of Democracy*. Oxford Scholarship Online Political Science. Oxford University Press. <https://doi.org/10.1093/oso/9780192848826.001.0001>.
- Barzilai-Nahon 2008** Barzilai-Nahon, Karine. 2008. "Toward a Theory of Network Gatekeeping: A Framework for Exploring Information Control." *Journal of the American Society for Information Science and Technology* 59 (9): 1493–512. <https://doi.org/10.1002/asi.20857>.
- Beckn 2026** Beckn. 2026. *Beckn Protocol*. <https://beckn.io/>.
- Berjon 2021a** Berjon, Robin. 2021a. *Governance of Ad Requests by a Union of Diverse Actors (GARUDA)*. Proposal December 1. <https://darobin.github.io/garuda/>.
- Berjon 2021b** Berjon, Robin. 2021b. "The Fiduciary Duties of User Agents." *SSRN Electronic Journal*, ahead of print. <https://doi.org/10.2139/ssrn.3827421>.
- Berjon 2024** Berjon, Robin. 2024. "Transmutations." March 22. <https://berjon.com/transmutations/>.
- Berjon 2025a** Berjon, Robin. 2025a. "How Wikipedia Can Save the Internet With Advertising." TechPolicy.Press. <https://www.techpolicy.press/how-wikipedia-can-save-the-internet-with-advertising/>.



- Berjon 2025b** Berjon, Robin. 2025b. *Requirements for a Healthy Ecosystem in Advertising (RHEA): Building Next-Generation, Pro-Democracy Advertising Infrastructure*. Supramundane Agency / Economic Democracy Project.
- Berjon and Crider 2025** Berjon, Robin, and Cori Crider. 2025. “Digital Sovereignty Can’t Be Bargained Away.” POLITICO, July 3. <https://www.politico.eu/article/digital-sovereignty-us-brussels-belgium-trade-talks-tech/>.
- Berjon and Yasskin 2025** Berjon, Robin, and Jeffrey Yasskin. 2025. *Privacy Principles*. Statement May 15. <https://www.w3.org/TR/privacy-principles/>.
- Bertrand et al. 2023** Bertrand, Arnaud, Christina Castella, Jérémie None, Soline Bouchacourt, Dominique Kerouedan, and Aka Kakou. 2023. *Evaluation of France’s Contribution to Unitaid (2006-2022)*. https://www.diplomatie.gouv.fr/IMG/pdf/summary_evaluation_france_unitaid_en-27-11-23_cle0cf24c.pdf.
- Carney 2026** Carney, Mark. 2026. “Davos 2026: Special Address by Mark Carney, PM of Canada.” World Economic Forum, January 20. <https://www.weforum.org/stories/2026/01/davos-2026-special-address-by-mark-carney-prime-minister-of-canada/>.
- Cath 2023** Cath, Corinne. 2023. *Loud Men Talking Loudly: Exclusionary Cultures of Internet Governance*. April. <https://criticalinfralab.net/wp-content/uploads/2023/06/LoudMen-CorinneCath-CriticalInfraLab.pdf>.
- Cattan and Toledano 2022** Cattan, Jean, and Joëlle Toledano. 2022. “La Commission dans la mise en œuvre du DMA: Citadelle assiégée ou chef d’orchestre?” *Concurrences*. https://www.conseil-ia-numerique.fr/files/archive/files/uploads/2022/Revue_Concurrences_3-2022_La_Commission_dans_la_mise_en_oeuvre_du_DMA.pdf.
- Clark 2012** Clark, David. 2012. “Control Point Analysis.” Preprint, SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2032124.
- Clinton-Gore Administration 1997, emphasis mine** Clinton, Bill, and Al Gore. 1997. “A Framework for Global Electronic Commerce.” July. <https://clintonwhitehouse4.archives.gov/WH/New/Commerce/read.html>.
- Collectif 2024** Collectif. 2024. “Pour Le Pluralisme Algorithmique!” *Le Monde*, September 25. https://www.lemonde.fr/idees/article/2024/09/25/pour-le-pluralisme-algorithmique_6332830_3232.html.
- Cooper 2025a** Cooper, Alissa. 2025a. “The True Cost of Browser Innovation: Why Chrome’s Divestiture Wouldn’t End the Open Web.” <https://www.techpolicy.press/the-true-cost-of-browser-innovation-why-chromes-divestiture-wouldnt-end-the-open-web/>.
- Cooper 2025b** Cooper, Alissa. 2025b. “The Web Can Thrive Without Google’s Search Monopoly.” <https://www.techpolicy.press/the-web-can-thrive-without-googles-search-monopoly/>.
- Dash 2026** Dash, Anil. 2026. “Endgame for the Open Web.” March 27. <https://anildash.com/2026/03/27/endgame-open-web/>.
- Doctorow 2025** Doctorow, Cory. 2025. *Enshittification: Why Everything Suddenly Got Worse and What to Do about It*. First edition. MCD, Farrar, Straus and Giroux.

- E. Ostrom 1998** Ostrom, Elinor. 1998. "The Comparative Study of Public Economies." *The American Economist* 42 (1): 3–17.
- E. Ostrom 2005** Ostrom, Elinor. 2005. *Understanding Institutional Diversity*. Princeton Paperbacks. Princeton University Press.
- E. Ostrom 2019** Ostrom, Elinor. 2019. *Governing the Commons: The Evolution of Institutions for Collective Action*. 10th printing. Canto Classics. Cambridge University Press.
- Edwards 2006** Edwards, Paul N. 2006. "Meteorology as Infrastructural Globalism." *Osiris* 21 (1): 229–50. <https://doi.org/10.1086/507143>.
- Elsayed-Ali and Berjon 2025** Elsayed-Ali, Sherif, and Robin Berjon. 2025. *Algorithmic Pluralism: Towards Competitive & Innovative Information Ecosystems*. <https://drive.proton.me/urls/XZ0NRTYEXG#JprcL8UNYD30>.
- Fajardo-García 2017** Fajardo-García, Gemma. 2017. *Principles of European Cooperative Law: Principles, Commentaries and National Reports*. Intersentia.
- Farrell and Newman 2024** Farrell, Henry, and Abraham Newman. 2024. *Underground Empire: How America Weaponized the World Economy*. Penguin Books.
- Frischmann 2013** Frischmann, Brett M. 2013. *Infrastructure: The Social Value of Shared Resources*. First printing in paperback. Oxford University Press.
- Germain 2025** Germain, Thomas. 2025. *Is Google about to Destroy the Web?* June 13. <https://www.bbc.com/future/article/20250611-ai-mode-is-google-about-to-change-the-internet-forever>.
- Google 2015** Google. 2015. *Accelerated Mobile Pages (AMP)*. Released. <https://amp.dev/>.
- Guha 2026** Guha, R.V. 2026. *NLWeb*. Released. <https://nlweb.ai/>.
- Hickson 2004** Hickson, Ian. 2004. "WHAT Open Mailing List Announcement — WHATWG." <https://whatwg.org/news/start>.
- Hodgson et al. 2025** Hodgson, Ramsay, Cristina Criddle, and Tim Bradshaw. 2025. "AI May Fatally Wound Web's Ad Model, Warns Tim Berners-Lee." *FT*, November 5. <https://www.ft.com/content/20592619-1bb9-451e-b4af-fcd33d981076>.
- Hof 2026** Hof, Laurens. 2026. "The Purpose of Protocols." *Connectedplaces.Online*, March 18. <https://connectedplaces.online/the-purpose-of-protocols/>.
- Hollister 2026** Hollister, Sean. 2026. "Google Search Is Now Using AI to Replace Headlines." *The Verge*, March 20. <https://www.theverge.com/tech/896490/google-replace-news-headlines-in-search-canary-coal-mine-experiment>.
- ITU 2022** ITU. 2022. *Individuals Using the Internet (% of Population)*. World Telecommunication/ICT Indicators Database. <https://data.worldbank.org/indicator/it.net.user.zs?end=2022&start=1960&view=chart>.
- Jones 1986** Jones, Alexander. 1986. *Pappus of Alexandria Book 7 of the Collection: Part 1. Introduction, Text, and Translation*. Sources in the History of Mathematics and Physical Sciences 8. Springer. <https://doi.org/10.1007/978-1-4612-4908-5>.
- Kardell 2022** Kardell, Brian. 2022. "Where Browsers Come From." <https://bkardell.com/blog/WhereBrowsersComeFrom.html>.



- Kollnig 2025** Kollnig, Konrad. 2025. "The Enshittification of Online Search? Privacy and Quality of Google, Bing and Apple in Coding Advice." arXiv:2512.03793. Preprint, arXiv, December 3. <https://doi.org/10.48550/arXiv.2512.03793>.
- Lardinois 2021** Lardinois, Frederic. 2021. *Mozilla Expects to Generate More than \$500M in Revenue This Year*. December 13. <https://techcrunch.com/2021/12/13/mozilla-expects-to-generate-more-than-500m-in-revenue-this-year/>.
- Levitsky and Way 2020** Levitsky, Steven, and Lucan Way. 2020. "The New Competitive Authoritarianism." *Journal of Democracy* 31 (1): 51–65. <https://doi.org/10.1353/jod.2020.0004>.
- Luu 2023** Luu, Dan. 2023. *How Bad Are Search Results?* <https://danluu.com/seo-spam/>.
- MacBride 1980** MacBride, Seán. 1980. *Many Voices One World*. <https://waccglobal.org/wp-content/uploads/2020/07/MacBride-Report-English.pdf>.
- Mahdawi 2025** Mahdawi, Arwa. 2025. "AI-Generated 'Slop' Is Slowly Killing the Internet, so Why Is Nobody Trying to Stop It?" Technology. *The Guardian*, January 8. <https://www.theguardian.com/global/commentisfree/2025/jan/08/ai-generated-slop-slowly-killing-internet-nobody-trying-to-stop-it>.
- McGinnis and Indiana University Bloomington 2002** McGinnis, Michael D. and Indiana University Bloomington, eds. 2002. *Polycentric Governance and Development: Readings from the Workshop in Political Theory and Policy Analysis*. Nachdr. Institutional Analysis. Univ. of Michigan Press.
- Mensch 2026** Mensch, Arthur. 2026. "Mistral CEO: AI Companies Should Pay a Content Levy in Europe." *FT*, March 20. <https://www.ft.com/content/d63d6291-687f-4e05-8b23-4d545d78c64a>.
- Michels et al. 1915** Michels, R., E. Paul, and C. Paul. 1915. *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Hearst's International Library Company. <https://books.google.com/books?id=8XXI87CLp5cC>.
- Mozilla Corporation and Opera 2004** Mozilla Corporation and Opera. 2004. "Position Paper for the W3C Workshop on Web Applications and Compound Documents." <https://www.w3.org/2004/04/webapps-cdf-ws/papers/opera.html>.
- Munir et al. 2025** Munir, Shaoor, Konrad Kollnig, Anastasia Shuba, and Zubair Shafiq. 2025. "Google's Chrome Antitrust Paradox." *Vanderbilt Journal of Entertainment & Technology Law* 27 (3): 419.
- Nield 2026** Nield, David. 2026. "Google's AI Overviews Can Scam You. Here's How to Stay Safe." Tags. *Wired*, February 15. <https://www.wired.com/story/googles-ai-overviews-can-scam-you-heres-how-to-stay-safe/>.
- Nottingham 2020** Nottingham, Mark. 2020. *The Internet Is for End Users*. RFC 8890. IETF, Informational August. <https://www.rfc-editor.org/rfc/rfc8890>.
- Nottingham 2024** Nottingham, Mark. 2024. "There Are No Standards Police." Mark Nottingham. <https://mnot.net/blog/2024/voluntary>.

- Nylen 2023** Nylen, Leah. 2023. *Apple Gets 36% of Google Revenue in Search Deal, Expert Says*. November 13. <https://www.bloomberg.com/news/articles/2023-11-13/apple-gets-36-of-google-revenue-from-search-deal-witness-says>.
- OWA 2025** OWA. 2025. "Apple's Browser Engine Ban Persists, Even Under the DMA." Open Web Advocacy, July 14. <https://open-web-advocacy.org/blog/apples-browser-engine-ban-persists-even-under-the-dma/>.
- OWA 2026** OWA. 2026. "Google Backs Down: Will Grant Hotseat in EU Browser Choice Screen." Open Web Advocacy, March 12. <https://open-web-advocacy.org/blog/google-backs-down--will-grant-hotseat-in-eu-browser-choice-screen/>.
- Ovide 2025** Ovide, Shira. 2025. "Google's AI Pointed Him to a Customer Service Number. It Was a Scam." *The Washington Post*, August 15. <https://www.washingtonpost.com/technology/2025/08/15/google-ai-overviews-scam/>.
- O'Neil 2016** O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Penguin Books.
- Pierce 2023** Pierce, David. 2023. *Google Paid a Whopping \$26.3 Billion in 2021 to Be the Default Search Engine Everywhere*. October 27. <https://www.theverge.com/2023/10/27/23934961/google-antitrust-trial-defaults-search-deal-26-3-billion>.
- Rahman 2018** Rahman, K. Sabeel. 2018. "Infrastructural Regulation and the New Utilities." SSRN Scholarly Paper No. 3205994. Social Science Research Network, June 30. <https://papers.ssrn.com/abstract=3205994>.
- Ricks et al. 2022** Ricks, Morgan, Ganesh Sitaraman, Shelley Welton, and Lev Menand. 2022. *Networks, Platforms, and Utilities: Law and Policy*. Faculty Books. <https://scholarship.law.columbia.edu/books/349>.
- S.W.I.F.T. SC 2024** S.W.I.F.T. SC. 2024. "Swift By-Laws." SWIFT, June. https://www2.swift.com/knowledgecentre/rest/v1/publications/s_byl/4.0/s_byl.pdf.
- Sato 2024** Sato, Mia. 2024. *The Perfect Webpage*. January 8. <https://www.theverge.com/c/23998379/google-search-seo-algorithm-webpage-optimization>.
- Schneider 2020** Schneider, Nathan. 2020. *We Need to Reinvent the Co-Op*. October 20. <https://nathanschneider.info/2020/10/we-need-to-reinvent-the-co-op/>.
- Scott 2014** Scott, Susan V. 2014. *The Society for Worldwide Interbank Financial Telecommunication (SWIFT): Cooperative Governance for Network Innovation, Standards, and Community*. With Markos Zachariadis. Routledge Global Institutions Series 83. Routledge. <https://doi.org/10.4324/9781315849324>.
- Sender 2024** Sender, Boaz. 2024. "WPT: An Overview and History - Bocoup." <https://www.bocoup.com/blog/wpt-an-overview-and-history>.
- Stolton 2026** Stolton, Samuel. 2026. *Google Hit by Fresh EU Antitrust Probe Over Search Ads Pricing*. February 12. <https://www.bloomberg.com/news/articles/2026-02-12/google-hit-by-fresh-eu-antitrust-probe-over-search-ads-pricing>.
- Szymielewicz et al. 2025** Szymielewicz, Katarzyna, Pamela Valenti, Ian Brown, Vid Logar, and Laurens Naudts. 2025. *Towards Algorithmic Plural-*

- ism. https://panoptikon.org/sites/default/files/2025-07/towards-algorithmic-pluralism-in-the-eu-policy_pvbt-discussion-paper_04072025.pdf.
- Tan et al. 2025** Tan, Joshua, Brandon Jackson, Robin Berjon, and Diane Coyle. 2025. *Airbus for AI: A Global Strategy for Public Value Creation*. Public AI / Bennett School of Public Policy. <https://publicai.co/airbus-for-ai.pdf>.
- Tarakiyee 2025** Tarakiyee, Tara. 2025. "Digital Sovereignty in Practice: Web Browsers as a Reality Check." *Do Flamingos Know They're Pink*, June 27. <https://tarakiyee.com/digital-sovereignty-in-practice-web-browsers-as-a-reality-check/>.
- The Economist 2025** The Economist. 2025. "AI Is Killing the Web. Can Anything Save It?" *The Economist*, July 14. <https://www.economist.com/business/2025/07/14/ai-is-killing-the-web-can-anything-save-it>.
- Thumfart 2024** Thumfart, Johannes. 2024. *The Liberal Internet in the Postliberal Era: Digital Sovereignty, Private Government, and Practices of Neutralization*. Palgrave Macmillan.
- Tremayne-Pengelly 2025** Tremayne-Pengelly, Alexandra. 2025. "Tim Berners-Lee Warns A.I. Could Kill the Web Economy as No One Visits Sites Anymore." *Observer*, November 6. <https://observer.com/2025/11/tim-berners-lee-ai-internet-ad-economy/>.
- United States v Google 2024** United States v Google (UNITED STATES DISTRICT COURT FOR THE DISTRICT OF COLUMBIA August 5, 2024). https://www.adexchanger.com/wp-content/uploads/2024/08/DOJ-Google-search-antitrust-Judge-Amit_Mehta-ruling.pdf.
- V. Ostrom 2006** Ostrom, Vincent. 2006. *The Meaning of Democracy and the Vulnerability of Democracies: A Response to Tocqueville's Challenge*. 3. print. Univ. of Michigan Pr.
- V. Ostrom et al. 1961** Ostrom, Vincent, Charles M. Tiebout, and Robert Warren. 1961. "The Organization of Government in Metropolitan Areas: A Theoretical Inquiry." *American Political Science Review* 55 (4): 831–42. <https://doi.org/10.2307/1952530>.
- Viljoen et al. 2021** Viljoen, Salomé, Jake Goldenfein, and Lee McGuigan. 2021. "Design Choices: Mechanism Design and Platform Capitalism." *Big Data & Society* 8 (2): 20539517211034312. <https://doi.org/10.1177/20539517211034312>.
- W. F. Whyte 1991** Whyte, William Foote. 1991. *Making Mondragón: The Growth and Dynamics of the Worker Cooperative Complex*. 2nd ed. With Kathleen King Whyte. Cornell International Industrial and Labor Relations Reports, v. 14. Cornell University Press.
- Warren et al. 2025** Warren, Jerome Nikolai, Lucio Biggiero, Jamin Hubner, and Kemi Ogunyemi, eds. 2025. *The Routledge Handbook of Cooperative Economics and Management*. Routledge International Handbooks. Routledge. <https://doi.org/10.4324/9781003449850>.
- Wilson and Çelik 2025** Wilson, Chris, and Tantek Çelik. 2025. *Vision for W3C*. Statement July 29. <https://www.w3.org/TR/w3c-vision/>.
- van Ess 2025** Ess, Henk van. 2025. *How AI Bots Quietly Dismantle Paywalls via Web Search*. July 11. <https://www.digitaldigging.org/p/how-ai-bots-quietly-dismantle-paywalls>.