

The Dimensions of Digital Coercion

Coraline Ada Ehmke

We are increasingly dependent on the internet to fulfill our basic personal, social, cultural, and political needs. But with the rapid pace of change in online spaces, the broader impact of technological innovation can be difficult to predict or identify. Many advances in technology ultimately reflect and sustain systems of inequity and oppression, whether directly or indirectly, deliberately or unintentionally. And the lack of accountability for how these technologies are deployed, by both public and private actors, makes us particularly vulnerable to coercive forces online.

Coercion is using the threat of harm as an incentive for obedience or cooperation, compelling people to act contrary to their own needs and best interests. The internet facilitates coercion at scale, posing new challenges to freedom, autonomy, and the exercise of our fundamental human rights. Understanding the mechanisms of digital coercion in online spaces is key to developing adaptable strategies for counteracting these threats.

What is digital coercion?

Digital coercion is simply the manifestation of mechanisms of control in online environments. There are four general dimensions of digital coercion: **attention**, **ergonomic**, **trust**, and **cultural**. Each of these dimensions can be understood by examining its aims, scope, threats, incentives, and impact.

Attention coercion

Attention coercion is the demand for more time or cognitive load than a person would voluntarily or consciously consent to otherwise.

Its **intent** is to force users to spend more time “engaging” or interacting with an app, platform, tool, or technology in a way that is secondary to what they are actually trying to accomplish. The **threats** that attention coercion relies on for compliance may include loss of access; loss of social, cultural, or professional status; and “FOMO” (“fear of missing out”). Its **impact** can manifest as hypervigilance, cognitive exhaustion, addiction, erosion of free time, or even radicalization.

Ergonomic coercion

Ergonomics are the product of design and ethical decisions. Ergonomic coercion involves forcing users to trade their autonomy, safety, privacy, time, or effort in exchange for usability.

Its **intent** is to restrict access to fundamental features that make a platform attractive, accessible, or useful. **Compliance** is incentivized through a platform's ubiquity, withholding or willfully violating basic usability and accessibility norms, or an embrace-extend-extinguish approach to open standards. The **impact** of ergonomic coercion includes forcing users to work in ways that best serve the provider, but not the users themselves. This can result in inequitable access for those who are unwilling or unable to satisfy unreasonable requirements imposed by a platform or provider.

Trust coercion

Trust coercion demands an assumption of benign intentions on the part of a provider or community, without any corresponding accountability for its trustworthiness.

Its **aim** is to encourage participation or adoption by instilling a sense of safety that may not actually be present. **Compliance** relies on pressure to conform to Western trust models, performative commitments to diversity/equity/inclusion, insisting on a presumption of "good intent", and obscuring or minimizing evidence of harm when it occurs. The **impact** of trust coercion is that vulnerable people may be exposed to threats to their safety or well-being, with little or no recourse when harm is caused.

Cultural coercion

Cultural coercion involves the imposition or prioritization of (predominantly white and Western) norms and values in exchange for the "privilege" of participation.

Its **aim** is to artificially homogenize a platform's user base by prioritizing the comfort and convenience of participants from the dominant culture. **Compliance** is ensured by a platform's ubiquity; forcing acceptance of inequitable power structures as a precondition for participation; the absence of representative governance; and lack of culturally appropriate alternatives. Its **Impact** can include the erasure of identity, social status, gender, race, religion, or other distinguishing characteristics of diverse participants, in turn propagating systems of exclusion that sustain digital colonization and other forms of slow violence against already vulnerable participants.

Digital coercion on the commercial internet

In her book "Technoculture: From Alphabet to Cybersex", Professor Leilia Green argues that technological advancements are most often shaped by the choices and priorities of those with the most power.

The business models of digital gatekeepers like Google, Amazon, and Meta are reliant on the normalization of data harvesting, surveillance capitalism, and ad-tech. And their distorted priorities are often at odds with the promise of internet technologies as a force for promoting justice and equity worldwide.

Resisting and mitigating corporate control of the internet requires understanding the coercive tactics that these companies employ to ensure and extend their technological, ideological, economic, and cultural control.

Attention coercion

Most popular commercial sites depend on their users' continual engagement, usually in service of a corporation's voracious appetite for data harvesting.

A provider's demand for constant attention is commonly enforced through deceptive design practices (sometimes called "dark patterns"), including popup windows and screen takeovers; relentless requests to complete surveys or provide feedback; interruptive push notifications enabled by default; and attention-grabbing "unread" badges that are visible even when the app is not in use.

Attention coercion is especially prevalent in centralized social media spaces. These platforms employ coercive strategies to drive engagement at all costs, often with devastating consequences. Hyper-engagement tactics may include artificially escalating real or manufactured conflict and controversy; promoting mis- or dis-information to foment interpersonal, social, or political conflict; and even knowingly radicalizing its users.

Ergonomic coercion

Ergonomics are commonly weaponized by platforms on the commercial internet.

Some providers employ "ergonomic paywalls", intercepting fundamental UI/UX affordances like copy-and-paste or screen captures, disabling hyperlinks, or removing basic accessibility features. This ensures that access to content is read-only and exclusively available through their platform.

Coercive ergonomic tactics are used to “capture” users, disincentivizing them from disengaging from a given platform or preventing them from using alternative 3rd party tools.

Forced continuity is another form of ergonomic coercion. Many online services offer “free trials”, at the end of which a user’s credit card is silently charged without any warning. The situation is made even worse by complex and opaque cancellation procedures that are deliberately designed to be as difficult as possible to use, including requiring telephone calls during business hours. It’s also common for platforms to make opting out of marketing emails as difficult as possible, hiding “unsubscribe” links in low-contrast small text, and directing a user to a web page intentionally designed to persuade or even guilt them into changing their mind.

The consequences of these user-hostile tactics include limited access to reliable news sources; preventing discovery and access to academic and scientific research; and interfering with critical, social, and cultural expression.

Trust coercion

Commercial platforms project an illusion of trustworthiness to attract and capture audiences. But in reality, trust and safety are rarely a core concern, especially when it interferes with a provider’s self-interests.

An overarching characteristic of trust coercion is the imposition of the dominant culture’s trust model: typically transactional, rather than affective. For example, many providers refer to their customers and users as a “community”, while refusing to prioritize basic personal and social safety features, let alone fulfilling fundamental community management responsibilities.

One tactic often employed by social media platforms is “friend spam”. Users may be prompted for access to their contacts or social media accounts, so that the provider can exploit the user’s personal network for its own commercial gain. An inverted form of friend spam invokes commutative trust, enticing a user to participate by showing them the profiles of people they may know who are already using the platform.

Without accountability and transparency, users can only see what providers *say*, not what they actually *do*. This makes an informed assessment of a platform’s trustworthiness nearly impossible. Combined with an artificially high cost of opting out of a particular platform, users are routinely coerced into risking their own safety and well-being in exchange for access.

Cultural coercion

In 1976, sociologist Herbert Schiller predicted that in the near future, the cultural lives of people around the world would be shaped and dictated by a small number of private media interests. Today, with the near-total domination of the internet by a handful of US tech corporations, this grim prophecy has become a reality.

Increasingly, the price of participation in online spaces is the tacit acceptance of the dominant culture's worldview and value system. Marginalized communities are most heavily impacted by this coercive imposition, subjecting them to irresistible mechanisms of exploitation, cultural appropriation, and digital colonization.

In her 2018 book "Algorithms of Oppression: How Search Engines Reinforce Racism", internet studies scholar Dr. Safiya Umoja Noble brought attention to how Google reinforced racist attitudes toward Black people— in particular, the sexualization of Black girls— in its search results. Although Google has since intervened in this particular instance of algorithmic violence, it continues to wield unassailable authority over digital taxonomies. In other words, Google has become the final arbiter of what information is true, valid, and relevant, through its distorted lens of opaque, oppressive, and unrepresentative cultural dominance.

On social media platforms, transgender women responding to anti-trans rhetoric and abuse [routinely get their accounts suspended](#) for so-called terms of service violations-- for example, a tweet that said "kick TERFs out of your spaces" framed as a "call for violence against people or groups". Major platforms not only refuse to moderate hateful positions on topics like LGBTQIA rights, racial justice, and anti-fascism— they profit by actively extending its reach and policing those who rightly call out hate speech.

The functionally oppressive priorities of the dominant culture are increasingly hidden behind an illusion of technological impartiality. And when this illusion comes under scrutiny, providers often put the blame on "unintentional bias" on the part of their otherwise "objective" algorithms, rather than acknowledging their own conscious complicity. Instead of accepting responsibility for the harm they cause, tech companies are inclined to make hollow promises to reform, draped in the co-opted language of the very communities they are harming the most.

Digital coercion on the open internet

The Open movement arose in response to the looming threat of complete domination of the internet by large commercial interests. In the intervening decades, the movement has

permanently changed how people around the world gain access to information and technology. But it's a mistake to assume that openness alone is a cure-all that leads to equitable outcomes. Good intent is simply not enough; even the most altruistic aims can lead to significant harm.

"Open" can manifest the same dimensions of coercion as the commercial internet, but expressed in different forms and with different motivations. Sometimes this happens through the co-option of open culture and technologies by corporations, but it also occurs independently as bias, unexamined privilege, power disparities, and other forms of systemic inequity are mapped into the open domain.

The impact of open technologies must be assessed not only within the scope of adopters or consensual users, but also with regard to collateral or non-consensual participants. It's important to ask to whom, in the traditional FLOSS model, is freedom actually extended? Whose freedom is ignored or deprioritized?

Attention coercion

While attention coercion on the commercial internet is employed to artificially drive engagement, open alternatives can also demand a disproportionate commitment of ongoing time and energy.

Even when a user manages to install and configure such an application, it may require an ongoing commitment to manually installing updates, security patches, and keeping "dependencies" up-to-date.

Compounding the time and attention required for basic installation, use, and maintenance, open source technologies often suffer from poor, outdated, or no documentation. Even when documentation is available, confused users may be met with an "RTFM" ("read the fucking manual") attitude when they ask for help. They may be directed to often unreliable third-party resources like Stack Overflow, spending even more time and effort sorting through often incomplete, low-quality, or out-of-date advice.

Requiring users to effectively become their own system administrators, combined with steep learning curves and a lack of support resources, is a form of digital gatekeeping that functionally excludes a vast majority of potential adopters and users.

Ergonomic coercion

The time-and-attention demands common to many open source technologies are often compounded by poor user experience design and other ergonomic choices. Some of these technologies, even those intended for general use, require such a high degree of technical aptitude and ergonomic tradeoffs that they are effectively out of reach for many users.

Open source applications may require users to use a terminal or shell, which brings its own set of ergonomic challenges, for basic installation, configuration, and upgrades. Even applications with graphical interfaces are often designed to be platform-agnostic (rather than cross-platform), and therefore fail to meet the ergonomic and accessibility expectations of users of a particular operating system. For example, many applications don't use native UI components, don't integrate with platform-specific UX affordances like global keyboard shortcuts, don't tie into native spell-checking functionality, or don't support OS-level text expansion configurations.

Another common side effect of platform agnosticism is that accessibility frameworks that are native to a given operating system must be independently reproduced by developers— meaning that, frequently, such features are simply not implemented at all. As a result, users who rely on assistive technologies to meet even basic accessibility requirements are once again pushed to the margins.

Without user-friendly, accessible open source alternatives, many people are effectively left with no choice but to rely on commercial technologies.

Trust coercion

It's important to acknowledge that “open” does not automatically mean trustworthy (or even safe). Open platforms that are decidedly altruistic in their aims are still functionally optimized for the safety and comfort of non-marginalized participants, just like their commercial counterparts.

With the rise of right-wing nationalism and authoritarianism worldwide, Wikipedia faces an ever-growing challenge in stemming the harm caused by mis- and dis-information. According to a 2016 study from Stanford University titled "[Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes](#)", the most successful malicious articles on Wikipedia were actually produced by long-time editors and contributors. In cases like this, an open platform's transactional, merit-based trust model actually opens new attack vectors for bad actors.

Even beneficent would-be contributors face safety and equity obstacles. For example, in the [MediaWiki archive documenting the Black Lives Matter protests](#) in the US in 2020, there is a startling lack of photographs of Black protestors. This is due in part to the very real threat of retribution through state or interpersonal violence; the portrayal of identifiable Black participants in these historic events can be literally life-threatening. Despite the egalitarian promise of open content platforms, lack of safe access means that our view of history continues to be distorted by the lens of the dominant culture.

Open source communities are also historically safest for non-marginalized contributors. Since 2014, with the creation of Contributor Covenant, codes of conduct have steadily gained acceptance (despite vocal and often violent opposition). But their normalization means that while the presence of a code of conduct used to be a positive signal of a community's inclusive intentions, today it's the *lack* of one that sends a stronger signal. Even well-meaning open source communities may be unprepared for actual enforcement, and scramble to react when a violation occurs and harm is done.

Without safe and trustworthy spaces, there cannot be equitable contribution. As a result, even open technologies are shaped by the priorities of those who wield the greatest privilege.

Cultural coercion

The philosophy of traditional open source is based on three core tenets: that meritocracy is an equitable measure of value; that technology is fundamentally neutral; and that unrestricted access to source code is an unqualified good.

Meritocracy has been widely criticized by tech justice advocates as a system that sustains inequity. Its fundamental assumption is that there is a level playing field, where everyone starts from the same place, with the same access and the same privilege of participation. Proponents of meritocracy also believe that the identities, lived experiences, and unique perspectives of participants are irrelevant: the only measure of value is what you produce, and how often. "Merit" itself remains undefined, a privilege that is granted in proportion to how well an individual mirrors the in-group's image of itself.

Ignoring decades of research by social scientists, many traditionalist free and open-source institutions continue to promote the myth of the neutrality of technology. In fact, the canonical definition of open source explicitly prohibits the exercise of any kind of moral or ethical authority, enshrining this dangerous attitude in "Freedom Zero" and insisting that creators and

contributors accept the unrestricted use of their creations, even for purposes that are explicitly malevolent. The annotated [Open Source Definition](#) even claims that “giving everyone freedom means giving evil people freedom, too.” The truth is that technology is only ever neutral towards its creators, and only to the degree to which it preserves the dominant culture’s social order.

The final tenet of traditionalist open source is that open access is a one-size-fits-all solution to technological inequity. But in practice, open access isn’t really open to everyone. Open access favors those with the most free time and the lowest degree of familial or community responsibilities. It prioritizes contributors from well-established, English-speaking technology hubs.

According to a 2022 paper titled “[The Geography of Open Source Software](#)”, 7.4% of global open source contributions come from developers in the San Francisco Bay Area. Even though only 17% of the world’s population speaks English, and 64% of those as a second language, over a third of all programming languages were developed in the US, UK, Canada, or Australia, and there are only a handful of multilingual programming languages.

The growing popularity of open source around the world facilitates the export of these norms, ideals, and assumptions. Participants from different countries and cultures are coerced into silently accepting these tenets as the price of participation. A history of open source contributions is often used as a proxy for professional competence, and is increasingly a gatekeeping factor in career development, so participants from outside the dominant culture have little choice but to accept, if not embrace, this ethos.

All of these factors help prop up systemic supremacy and accelerate digital colonization, enshrining a status quo that coerces participants into accepting the ethical framework of the dominant technological culture. Promoting, defending, or tolerating these systems, even under the auspices of openness, is functionally indistinguishable from intentional exclusion.

Standards as an accountability technology

Digital coercion, in all its forms and on both sides of the open/closed divide, derives its power from lack of accountability. Without accountability, there can be no trust. Without trust, there can be no consent. If we have no real agency to make informed choices about our safety and well-being, we have no defense against coercion.

Standards can help create the conditions for accountability, opening possibilities for more humane, ergonomic, trustworthy, and culturally appropriate alternatives to coercive platforms and technologies.

There are many forms of standards that apply to digital spaces. Regulatory standards establish basic rights and responsibilities, for example those defined in the [General Data Protection Regulations \(GDPR\)](#). Technical standards, like the [Unicode Standard](#), ensure basic interoperability across implementations. Normative standards apply to behavioral and social expectations and are often outlined in codes of conduct like [Contributor Covenant](#). And foundational ethical standards, such as the [Ethical Source Principles](#), express what is and is not morally acceptable in a given context.

Fundamentally, standards reflect our ideals. The rules, constraints, and guidelines defined by a standard express what we want, and what we don't; what we will tolerate, and what we won't; what we hope to happen, and what we hope to prevent.

These ideals can manifest across multiple layers of standards. For example, Tim Berners-Lee has declared that [access to the web by everyone, regardless of disability, is essential](#); this frames the ideal of accessibility as an ethical standard. Accessibility manifests as a normative standard on platforms like Mastodon, where there are cultural expectations for providing alt-text descriptions of images that are posted. The accessibility ideal manifests as a technological standard in the [W3C Accessibility Standards](#), and as a regulatory standard in [Article 9 of the UN Convention on the Rights of Persons with Disabilities](#).

Understanding standards as expressions of ideals highlights the importance of equitable collaboration across social, cultural, and disciplinary boundaries. Cultural hegemony in the design of techno-social systems is an anti-pattern that inevitably leads to harm at scale.

Conclusion

The coercive forces we are subject to in digital spaces are extensions of broader systems of inequity. And while these systems can feel too big or too ubiquitous to overcome, we cannot simply resign ourselves to the status quo.

Interrupting systems of technosocial oppression requires constant vigilance, examination, collaboration, and creativity. It requires understanding that these systems are dynamic, and continually adapt in response to changing conditions, often outpacing our efforts to mitigate

their harms. But regardless of their particular manifestations, they are all heavily dependent on different (or interwoven) forms of digital coercion.

The framework presented here is intended to inform strategies for reshaping the realities of online participation. Understanding different manifestations of digital coercion can help us develop more effective strategies for resisting, subverting, and replacing the systems that they sustain.

Coercion can operate across proprietary/open boundaries, so our strategies must also transcend these boundaries. The complex ecology of the modern internet requires us to take a dialogical, rather than confrontational, approach to this divide, to counter the many manifestations of digital coercion.

Agency, accessibility, privacy, autonomy, and other digital rights are not individual concerns that can be addressed by individual choices. Real accountability means prioritizing pro-social outcomes over both the profit motives of "closed" and the philosophical purity of "open" technologies.

References & further reading

- [Afrofeminist Data Futures](#)
- [Algorithms of Oppression](#)
- “Communication and Cultural Domination” by Herbert Schiller
- [Critical about Critical and Speculative Design](#)
- [Decentralized social media platform Mastodon deals with an influx of Gab users](#)
- [Declaration of Digital Rights](#)
- [The Dehumanizing Myth of Meritocracy](#)
- [Design for Safety, A Book Apart](#)
- [Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes](#)
- [Does Open Source Need Its Own Priority of Constituencies?](#)
- [The Ethical Source Principles](#)
- [The Geography of Open Source Software: Evidence from GitHub - ScienceDirect](#)
- [Human-centric perspective on digital consenting](#)
- [International Organization for Standardization \(ISO\) - Deliverables](#)
- [Learning to Code in One's Own Language](#)
- [Race Equity and Inclusion Action Guide](#)
- [Slow Violence](#)
- [Technoculture](#)
- [Trans people keep getting suspended from Twitter—and they want answers](#)
- [Types of deceptive design](#)
- [The unseen 'slow violence' that affects millions](#)
- [Violence, Peace, and Peace Research](#)
- [Wikipedia:Disinformation](#)

About the Organization for Ethical Source

The [Organization for Ethical Source \(OES\)](#) is a diverse, multidisciplinary, and global community that is revolutionizing how tech culture works. We are investing in tools like Contributor Covenant as part of our commitment to creating a better future for open source communities around the world. If you'd like to help us shape that future, consider [becoming an OES member](#).