

Uncolonized Futures: Big Tech's AI Infrastructure and the Rhetoric of Climate Deferral

Asra Zahid^{1*}

¹ Department of Development Studies, School of Social Sciences and Humanities (S3H), National University of Sciences and Technology (NUST), Islamabad, Pakistan. **Corresponding Author:** Email: asrazahidk@gmail.com. ORCID: 0009-0004-4920-4642

ABSTRACT

This paper explores how six major technology companies: Amazon, Microsoft, Apple, NVIDIA, Google, and Meta address the environmental implications of their artificial intelligence (AI) infrastructure through their discourse in different genres of official documents. The research draws on corporate texts, including sustainability reports, executive keynotes, shareholder letters, and regulatory filings, along with independent evidence from non-corporate sources. The research investigates three questions. First, it examines how rhetorical strategies vary by communication genre and target audience. Second, it assesses how temporal deferral functions as a means of legitimation across the dataset. Finally, it evaluates whether discursive confidence enabled by temporal deferral correlates with weaker material emissions performance across the sample. Utilising inductive discourse analysis grounded in sociotechnical imaginaries (Jasanoff & Kim, 2009, 2015); the sociology of expectations (Brown & Michael, 2003; Borup et al., 2006); and Crawford's (2021) materialistic critique of computation, the findings indicate a consistent dual narrative structure: an outward-looking promissory register (performed primarily through keynote addresses and shareholder letters) that represents AI as an active ecological actor, and an inward-looking defensive register (limited to regulatory filings) that re-classifies those same obligations as non binding legal risks. Further comparative analysis of this two-register structure against the 2020-2024 emissions trajectories for each corporation indicates a striking asymmetry. Building on prior use of the term "technowashing" (Perakslis, 2020; Ribeiro & Soromenho-Marques, 2022), this paper refines this concept within climate-AI discourse by analyzing its reliance on uncolonized future times where Big Tech uses speculative promises about a future humans have not reached yet to shape the present debate (Brown & Michael, 2003). It argues that this temporal deferral does not merely exaggerate claims; rather, it allows AI-driven promissory discourse to expand alongside rising rates of material extraction, rather than emissions being reduced by tech solutions.

ARTICLE HISTORY

Received 21 June 2026
Revised 30 July 2026
Accepted 10 August 2026
Published 14 August 2026

KEYWORDS

Technowashing, Temporal Deferral, Sociotechnical Imaginaries, AI Infrastructure, Corporate Climate Discourse

Introduction

A shareholder letter and a regulatory filing describe the same company in different ways. One communicates through a registry of destiny, and the other through a registry of liability. This paper investigates the dual nature of registry across genres by asking how large technology corporations communicate their responsibility for the negative environmental impact of Artificial

CONTACT Asra Zahid ✉ asrazahidk@gmail.com 🇵🇰 Pakistan

© 2026 The Author(s). Published by Dialogues in Humanities and Social Sciences.

This is an Open Access article distributed under the terms of the Creative Commons Attributions-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Intelligence (AI) infrastructure. From 2024 to 2026, six corporations (Amazon, Microsoft, Apple, NVIDIA, Google, Meta) increased their investment in AI infrastructure while simultaneously reaffirming climate commitments made before those investments. While this is not unique among large corporations (i.e., all organizations have conflicting commitments), what is analytically interesting is how each corporation describes its conflicting commitment to climate mitigation differently based on the audience. At a keynote conference, the CEO can say that AI will discover a new industrial coolant (Microsoft, 2025c). However, when communicating with shareholders, the same CEO writes that the company is "thinking in decades and executing in quarters" (Nadella, 2025). Furthermore, regulators read the same statements from the same fiscal year as non-binding, aspirational goals that could expose the company to legal action if it failed to meet them (Microsoft, 2025b). That was not three corporations; it was one corporation speaking in three distinct registers.

The paper treats register-shifting as the primary subject of study, not just background noise behind a secondary climate story. Three research questions organize the analysis. Firstly, what are the recurring rhetorical strategies that are used throughout the corporate corpus, and how do these strategies vary by genre (e.g., formal reports, executive keynotes, shareholder letters, regulatory filings) and target audience? Second, what role do temporally deferred claims (the placing of a claim's verification in some unspecified time frame) play in allowing these different registers to be articulated as part of a singular and consistent corporate identity? Finally, does the discursive confidence enabled by temporal deferral correlate with weaker material performance on emissions, and what does that divergence reveal about the limits of legitimation through unfalsifiable, future-oriented claims?

This paper does not independently measure emissions or water use; it draws on existing third-party audits to evaluate the discursive claims against them. The paper begins by situating this argument within two previously published uses of the term "technowashing," then details the study's methodological approach, outlines four focused findings, and concludes by highlighting what a genre-sensitive analysis contributes to debates over accountability.

Literature Review and Theoretical Framework

Sociotechnical Imaginaries

Science and Technology Studies (STS) uses sociotechnical imaginaries to describe how a technology can become socially acceptable. Sociotechnical imaginaries are collectively imagined ways to organise society through designing and implementing scientifically/technologically based, nation-specific projects (Jasanoff & Kim, 2009), which Jasanoff (2015) later defined as "collectively held, institutionalized, and publicly performed visions for desirable futures." Sociotechnical imaginations are long-term cultural frameworks defining what societies consider achievable and worth striving for (Jasanoff & Kim, 2009; Jasanoff, 2015). The importance of sociotechnical imagination to corporate discourse is that they explain how a corporation's profit interests become indistinguishable from a societal vision of progress; embedding an organization's AI systems within a common story of national or global development makes those systems impervious to ordinary scrutiny (Burri, 2015; Jasanoff & Kim, 2009). The authors note that in their own case study of the United States and South Korea developing vastly different "nuclear imaginaries," one based on caution with nuclear weapons and the other on a need for nuclear energy development, both countries developed these two completely different "imaginaries," and did so surrounding the very same technology. As such, the authors conclude that the imaginative process, not the technology itself, provides legitimation.

The Sociology of Expectations and Uncolonized Future Time

Where imaginaries develop and become part of long-lasting culture, the sociology of expectation analyzes the short-term commitments stakeholders make to keep an imagination

relevant in the current period. These short-term commitments are called "technological expectations." They represent "real-time representations of potential technological conditions and abilities" (Borup et al., 2006) and serve several key roles: they encourage investment, build institutional credibility, and stimulate interest in a new technology before it performs. Brown and Michael (2003) identify the fundamental temporal mechanics behind this function as the colonizing of "uncolonized future time." Corporate and scientific actors define an uncolonized future time as having value, and defer until later the time when a promise will be subject to verification. This is the paper's primary analytical tool. A commitment routed into "uncolonized future time" cannot be evaluated today; this is exactly what separates a promise from an ordinary statement of fact. In their example of xenotransplantation research, Brown and Michael (2003) clearly demonstrate the basic mechanism. Despite numerous clinical setbacks, each failure was incorporated into the imagined future outcome of xenotransplantation as "noise" or temporary "technical delays" along an otherwise stable trajectory toward an expected future. Similarly, this paper treats corporate climate-AI rhetoric as using the same mechanism.

Greenwashing as organizational decoupling

While sociotechnical imaginaries and expectations describe how a promise develops and is maintained, organizational theory explains why an organization would make a promise. Delmas and Burbano (2011) define greenwashing as "the intersection of two firm behaviors: poor environmental performance of a company and positive communication about that company's environmental performance." Motivations for these actions include marketplace pressure from environmentally conscious stakeholders, weak or nonexistent regulatory enforcement, and predictable cognitive biases like optimistic bias and hyperbolic discounting among decision makers. This conceptual model is useful; however, it has one limitation in this case: a greenwashing claim can be verified against a present and fixed condition. An audit of false labeling of an efficiency certification can be conducted today (Delmas & Burbano, 2011). For many of the promises examined in this study, it will be impossible to verify them today because their truth conditions have been intentionally placed in the future. This research attempts to bridge that analytical gap: claims made in the future cannot be evaluated until after they become the past. This draws on Suchman's (1995) definition of legitimacy management as a strategic and symbolic practice, not just a compliance error.

Artificial Intelligence as a Material System

The paper draws on materialism via a critique by Kate Crawford (2021) to counterbalance its purely discourse-based analysis. According to Crawford, AI is "neither artificial nor intelligent" and is instead a materially situated system constructed from resources such as minerals, energy, labor, and logistical infrastructure (Crawford, 2021; Crawford & Joler, 2018). A central component of her critique is the concept of a "double operation of abstraction/extraction": the same corporate rhetoric that removes computation from its physical basis and presents it as clean or disembodied. This enables companies to justify physically removing the materials these systems require (Crawford, 2021). Crawford and Joler (2018) trace a single voice command made using Amazon Echo through all of the material components involved in producing that technology, including lithium mined in Nevada, cobalt mined in the DRC, ocean-spanning cables connecting those mines to data centers, and workers abroad who were tasked with labeling the audio files. The paper treats each of the discursive patterns described in the sections that follow as examples of that double operation, i.e., a particular style of speaking, along with the particular material consequences that that style can mask.

Positioning Technowashing against prior use of terms

Perakslis (2020) defines technowashing as a metaphorical means of concealing reality in the digital realm and calls it one of many ways people have concealed realities ("whitewashing,"

"greenwashing," etc.). She focuses her research on how companies create digital dependency and the effects (such as unmapped informal settlements and systems that do not account for gender and social inclusion) when they attempt to hide their designs. Perakslis' definition of technowashing also limits it to hiding evidence of digital dependency and does not include organizations making false environmental claims based on new technologies. Ribeiro and Soromenho-Marques (2022) expand upon Perakslis' idea of technowashing and apply it to macro-political discourses. They refer to technowashing as "the organizational and political laundry" that produces overly optimistic projections about technological advancement to draw public attention away from real-world issues. Their examples of technowashing include proposals to colonize Mars and geo-engineer. These are large-scale, futuristic projects that operate outside of typical corporate actions. Furthermore, their example of technowashing occurs via distraction and simplification of complex ideas.

In contrast to Perakslis (2020), this paper differs in its conceptualization of technowashing because AI infrastructure is not a speculative concept operating in the distant future. Instead, it currently exists as an active system affecting power grids and water tables (Crawford, 2021). Therefore, technowashing operates through abstraction and extraction (as Crawford describes) rather than as a way to distract the audience from a separate reality. This paper also identifies a different mechanism behind technowashing: temporal structures. Corporate AI-climate rhetoric projects into Brown and Michael's (2003) uncolonized future time, allowing corporations to treat current emissions increases as part of the required technical lag rather than as broken promises. This paper introduces a specific temporal mechanism as its primary analytical addition to the previously established term. To trace this temporal mechanism empirically rather than assume it, the study proceeds inductively, building categories from the corpus itself before applying the theoretical framework outlined above.

Methodology

Research Design

This study employs a qualitative discourse analysis grounded in Science and Technology Studies (STS) discourse to examine how genre and target audiences shape corporate legitimacy claims about their use of AI Infrastructure. Coding was inductive: we built categories from the corpus. After coding, we applied a theoretical framework once initial patterns emerged across sources.

Corpus

Three categories within the corpus are analytically distinct from the start of coding. The first category is Corporate Discourse, spanning four genres: Sustainability & Environmental Reports, Executive Keynote Addresses, Shareholder Letters, and Securities & Exchange Commission Regulatory Filings (10-Ks) produced by six corporations: Alphabet (Google), Amazon, Apple, Meta, Microsoft, and NVIDIA. The second category is Independent Evidence generated by organizations with no direct or indirect interest in the claims being made. Organizations in this category include the AI Now Institute, Greenpeace East Asia, Greenpeace International, and the Government Digital Sustainability Alliance (GDSA). A much smaller third category called Supplemental Evidence consisted of United States Congressional testimony regarding the implications of an artificial intelligence workforce on public policy and government procurement; however, the testimony was limited to public policy and did not address climate change issues; therefore, only one reference (Castor, 2026) discussing both artificial intelligence infrastructure and electricity ratepayer's interests will be analyzed and this limitation is clearly identified to maintain transparency.

Unit Analysis and Coding Procedure

The unit of analysis is the discursive claim, a bounded statement asserting a relationship between Big Tech's AI infrastructure and an environmental, financial, or social outcome. Coding occurred in two stages. In stage one, I coded individual corporate sources for three variables: the

promise made by the source, linguistic markers indicating confidence/hedging, and attribution of responsibility (credit claimed vs. blame assigned to external parties).

In stage two, sources from the same firm were compared across genres, i.e., a passage from a Microsoft keynote compared with the corresponding sustainability report claim from Microsoft, as well as the related claim from its 10-K statement. This assessed how different genres affect the register of the same underlying claim when directed toward different audiences. The independent evidence corpus was coded separately for material findings and cross-referenced to the specific corporate claim area each finding addressed.

Ethical Considerations

This study analyzed only publicly published documents prepared by corporations and institutions. No human subjects were involved, and no primary data collection from individual respondents occurred.

Results

These findings discuss the linguistic mechanisms that occurred most frequently across genres and corporations. These mechanisms are shown here using the source in which they occur most clearly.

A Dual-Narrative Structure: Promissory Address, Defensive Filing

All six firms' discourse clearly divides between a promissory register directed at external stakeholders and a defensive register directed at internal stakeholders.

Firm	Promissory Register (Keynote / Letter)	Defensive Register (Form 10-K)	Same Fiscal Period?
Amazon	AI is “a once-in-a-lifetime opportunity” (Jassy, 2025)	“Profitability or other intended benefits... in our newer activities [including AI]... may not meet our expectations” (Amazon Inc., 2026)	Yes (2025–2026)
Microsoft	AI will “discover a better coolant” for data centers (Microsoft, 2025c)	“Any failure or perceived failure to meet our sustainability goals... could result in claims and lawsuits... or damage to our reputation” (Microsoft, 2025b)	Yes (2025)
Apple	Cook stated: Apple Intelligence “understands your personal context” (Apple, 2026)	Filings reserve hedged, liability-oriented language for supplier and regulatory risk not present in keynote or report text (Apple Inc., 2024, 2025)	Adjacent (2025–2026)
NVIDIA	AI factories have “one job and one job only”: generating “incredible tokens” (NVIDIA, 2025a)	“Stakeholders may not be satisfied with our corporate sustainability practices and goals” (NVIDIA, 2025b)	Yes (2025)
Google/Alphabet	AI heralds a “new golden age of discovery	“We cannot guarantee that our goals and initiatives will be fully realized on the timelines we	Yes (2025)

	and wonder" (Google, 2025a)	expect or at all" (Alphabet Inc., 2025)	
Meta	AI will unlock "personal superintelligence" and unprecedented productivity (Zuckerberg, 2025)	"Perceived failure... to comply with applicable... laws and regulations could result in legal and regulatory proceedings... and adversely affect our business" (Meta Platforms, Inc., 2026)	Yes (2025–2026)

Table 1. Promissory versus defensive register, paired by firm.

Elaborating on one of the six firms, Microsoft's 2025 corpus makes this phenomenon particularly evident, as it offers four different genres for the same fiscal year. Satya Nadella frames AI as an active participant in ecological exploration in the keynote register, describing an experimental cycle: "the reasoning process over knowledge, generating hypotheses, and running experiments", through which AI will "discover a better coolant" that will allow harmful chemicals to be replaced in data center operations. (Microsoft, 2025c). The same time period is framed in an elevated but assuring manner in the shareholder-letter register; Microsoft finds itself "at the heart of a generational moment in technology," and any increased resources in the immediate future are placed within a temporal promise, as the company is thinking in "decades" and executing in "quarters" (Nadella, 2025). In the sustainability-report register, the same commitments appear as a resolute determination from the corporation tempered by acknowledgement of difficulties although the corporation is "determined" in pursuing its 2030 targets (Microsoft, 2025a), its scope 3 greenhouse gas emissions were up 26 percent due to "growth in AI and Cloud" (Microsoft, 2025a), this conflict is managed through the temporal diversion examined below.

In the regulatory filing register, the same commitments are fully transformed. Microsoft's 2025 Form 10-K does not represent the company's sustainability goals as commitments. Instead, it presents them as potential liabilities: "any failure or perceived failure to meet our sustainability goals... could result in claims and lawsuits, regulatory actions, penalties, or damage to our reputation" (Microsoft, 2025b). It also states that many of the ESG objectives "depend on the adoption of certain behaviors... by third parties, including our customers and partners" and that the "firm may not be able to meet some goals" if those parties fail to cooperate (Microsoft, 2025b). No keynote nor shareholders' letter in the entire corpus includes language this cautionary; no 10-K in the entire corpus includes language this bold. Each document's intended audience defines its register, not any other factor. In contrast, the fact base did not change between documents that were filed during the same fiscal year.

This pattern generalizes across the corpus. Apple's promotional register relies on unqualified assurances; for example, it uses terms such as "milestone," "unequivocal," and "tangible outcomes" (Apple Inc., 2024). Apple's filings, like those of other Big Tech corporations, reserve hedged, liability-oriented language for regulators alone. Alphabet's 2026 proxy statement extended the defensive register beyond the keynote address and filing registry. It leaks into the shareholder governance itself. The board's response to a shareholder-initiated climate disclosure proposal states "our existing environmental reporting has long provided detailed and extensive information" (Pichai, 2026). This represents a defensive position the company takes not against regulatory bodies, but on behalf of its own investors.

The dual-narrative framework identified earlier operates across at least three distinct audiences, rather than just two. This divide exists consistently across six separate companies operating independently. The pattern supports the idea that this represents not an incidentally developed writing style but a structural feature of how Big Tech manages two types of audiences, two types of contractual obligations, and two registers.

Responsibility Attribution: Who Is Credited, Who Is Blamed

The next consistent pattern that arises within the promissory register is that corporations give themselves credit for outcomes attributable to their internal engineering decisions, while giving external entities such as suppliers, other third-party partnerships, or "the industry" credit for outcomes out of their hands. Table 2 illustrates how each of the six companies attributed outcomes across the board.

Firm	Credits Itself For	Blames External Parties For
Amazon	AWS power-usage efficiency of 1.15, versus an industry average of 1.25 (Amazon Inc., 2025)	74% of its footprint, attributed to "data center construction and fuel consumption by third-party delivery service providers" (Amazon Inc., 2025)
Microsoft	Contracted carbon-removal credits and server-recycling programs (Microsoft., 2024)	ESG goals that "depend on the adoption of certain behaviors... by third parties, including our customers and partners" (Microsoft, 2025a)
Apple	Its own product-design and engineering teams for on-device processing that avoids cloud computation (Apple Inc., 2025)	Supply-chain emissions, managed by requiring suppliers to adopt a clean-energy mandate rather than absorbing the cost directly (Apple Inc., 2024)
NVIDIA	Chip-design efficiency gains, e.g., a claimed "25x improvement in energy efficiency" for the Blackwell generation (NVIDIA, 2025a)	Manufacturing footprint, offloaded by noting NVIDIA is "a fabless semiconductor company" with no factories of its own (NVIDIA, 2024)
Google/Alphabet	TPU efficiency gains and carbon-free energy procurement (Google, 2024)	Regional CFE shortfalls, attributed to "permitting processes, inadequate grid infrastructure, and challenges in the renewable energy supply chain" (Google, 2024)
Meta	Matching 100% of operational electricity with renewable energy (Meta Platforms, Inc., 2024)	99% of its footprint, located "multiple tiers upstream... where there is currently no viable approach to trace emissions" (Meta Platforms, Inc., 2024)

Table 2. Attribution of credit and blame, by firm.

In every instance, the division was one-way. No company credited a supplier with efficiency improvements; no company blamed its own engineering for increasing emissions. That imbalance is also a discursively created phenomenon; it relies on the idea that what constitutes "our corporation" and the "supply chain," and thus what is accountable to us versus them, is fixed when we wish to hold someone accountable and flexible when we want to take credit. For example, NVIDIA's fabless model keeps product manufacturing outside the bounds of accountability, while NVIDIA uses first-person language in its advertising to claim credit for product performance.

Temporal deferral as the mechanism that holds the deflection together

The dual-narrative format works only when the two registers do not contradict each other. That is where temporal deferral comes in: it shifts a current deficit into a future reality already underway, turning an unrealized goal into something that will happen at some point in time rather than a failed commitment.

Firm	Deflection Phrase	Present Shortfall Being Deflected
Amazon	"The path to being a more sustainable company will never be linear, because we are charting new territory at an unprecedented scale" (Amazon Inc., 2025)	6% rise in absolute emissions in 2024, to 68.25M MMT CO ₂ e (Amazon Inc., 2025)
Microsoft	"Our journey towards being carbon negative is a marathon, not a sprint" (Microsoft, 2025a)	26% rise in Scope 3 emissions, attributed to "AI and cloud expansion" (Microsoft, 2025a)
Google/Alphabet	"Ambitious vision, complex reality" framing an operational milestone against a rising baseline (Google, 2025b)	Data center electricity consumption up 27% year-on-year; Scope 3 emissions up 22% year-on-year (Google, 2025b)
Meta	Participation in the "Symbiosis Coalition," an advance market commitment to future carbon-removal credits (Meta Platforms, Inc., 2025)	Total carbon footprint rose to 8.2M MT CO ₂ e in 2024 (Meta Platforms, Inc., 2025)
NVIDIA	Suppliers "engaged... to adopt science-based targets" framed as an in-progress solution (NVIDIA, 2025a)	Scope 3 emissions of 6.9M MT CO ₂ e, with purchased goods alone at 6.0M MT CO ₂ e (NVIDIA, 2025a)
Apple	On-device "Apple Intelligence" reframed as a durable architectural solution rather than a partial mitigation (Apple Inc., 2025)	Cloud-dependent tasks still routed to Apple's own server fleet, whose footprint is reported separately (Apple Inc., 2025)

Table 3. Temporal deferral phrase and the shortfall it accompanies, by firm.

Microsoft shows how clear this language can be. They say that their effort to become carbon negative "is a marathon, not a sprint." (Microsoft, 2025a) While Alphabet does not use these exact words, it does the same thing with much more specificity in its regulatory reporting. In addition to saying there would be "a rise before a drop" in their emissions, meaning that for them to reach zero net emissions eventually, their emissions have to increase first, they also describe that those short-term increases are part of an overall long-term decrease in emissions that is already planned. (Alphabet Inc., 2025) Both statements create the same effect as Brown and Michael (2003) did in their research on the hype surrounding xenotransplantation. Those researchers define an unfulfilled result not as a failure of the claim itself, but as a "technical delay" on what was otherwise an

uninterrupted path. (Brown & Michael, 2003). Temporal deferral succeeds precisely because it cannot be verified until later. This is why proof of success is simply postponing the test.

Beyond specific phrasing, temporal deferral is built into how Microsoft and Alphabet measure and disclose their greenhouse gas emissions. For example, Microsoft reports that “Scope 3 emissions,” i.e., the indirect emissions caused by suppliers, manufacturing, and product use, accounted for “more than 96 percent of all of our greenhouse gas emissions.” (Microsoft, 2024) At the same time, Microsoft gives credit to its purchasing power for buying “the rights to retire 5,015,019 metric tons of carbon emissions over the next 15 years.” (Microsoft, 2024) Thus, while Microsoft acknowledges that almost all of its present-day footprint exists now, most of that footprint is being offset through contracts to remove CO₂ from the atmosphere over the next fifteen years. These contracts appear in publicly disclosed legal documents that assure investors the company has committed to reducing its climate impact. However, since no one knows how well these efforts will succeed over the next fifteen years, nothing can confirm whether they will succeed during this time frame.

The Promissory Delta: Discourse Against Measured Emissions

The alignment between discourse confidence and material results varies widely throughout the corpus. The direction of the gap also matters. Table 1 depicts total CO₂-equivalent emissions (including Scopes 1, 2 & 3, without offsets) for the six companies that disclose similarly, compared to a 2020 base year.




Figure 1. Gross GHG emissions, percentage change from 2020 baseline, by firm (2020–2024).

Firm	2020 Baseline (Mt CO ₂ e)	2024 Actual (Mt CO ₂ e)	Change	Dominant Discursive Register
Meta	5.13	8.20	+59.8%	Ecological-agent promissory (Meta Platforms, Inc., 2024, 2025a)
Google/Alphabet	10.80	15.80	+46.3%	Temporal deferral ("rise-before-drop") (Alphabet Inc., 2025)
Microsoft	Baseline	14.30	+23.4%	Temporal deferral ("marathon, not a sprint") (Microsoft, 2025a)
Amazon	60.64	68.25	+12.5%	Bottleneck-as-opportunity reframing (Amazon Inc., 2025)
Apple	22.60	15.30	-32.3%	Restrained, engineering-credit register (Apple Inc., 2024, 2025)

Table 4. Grounded emissions metrics and dominant discursive register, by firm.

The pattern shown is not that the promissory language is false or misleading, but that the intensity of the promissory language does not correlate well, and at times inversely correlates, with actual results. Meta and Google are both responsible for the largest increase in their respective emissions. At the same time, they were the only two companies among the six studied that used the most extreme ecological-agency framing throughout the corpus. Ecological-agency framing is discourse that views AI as an agent working to optimize the planet rather than as a resource-intensive system to manage. Conversely, Apple uses the least rhetoric in the corpus and reported declining emissions by 32.3% over the same period, rather than relying on statements about AI's future planetary benefits. A truly counter-intuitive result if one subscribes to a technowashing perspective. The company reporting the greatest measurable decoupling from the "bad" (emissions) behavior is not the company claiming the most rhetorical space for decoupling. As noted above, if promissory language were a simple measure of how well a company performed, then the relationship would be reversed. Instead, the most aggressive promissory language appears to substitute for describing/reporting actual success, similar to Crawford's (2021) explanation of abstraction serving as a means to legitimize extraction rather than describe it.

Beyond the Firm: Independent Evidence and the Limits of Self-Report

Because the promissory register's central claims are situated in an unverifiable future, independent non-corporate evidence is required to serve as a verification check on corporate self-reporting, which, due to its inherent temporal structure, cannot be independently verified. Three key empirical findings directly relate to the discourse patterns outlined in Section 4.2. The first concerns the credibility of the promissory register itself: an industry-wide audit of technology sector claims

about AI's climate benefits found that about 75% of these claims were unproven, with no documented case of a consumer AI tool delivering a confirmed emissions reduction (Leman, 2026). A related finding on water comes from independent audit data showing an increasing trend in AI data center water usage. However, firms are reporting their volumetric "water-positive" commitments. Still, their evaporative cooling systems lose most of the water to direct evaporation back into the atmosphere rather than returning it to the source (Kenny, 2025). This significant fact is completely omitted from the promissory register presented in Section 4.1 because neither the keynote speaker for any of the six firms nor their letters in the corpus mentioned water.

This directly tests the ecological-agent framing discussed in Section 4.1 and finds it empirically unsupported across the technology sector at present. The second finding related to the material substrate that the promissory register abstracts-away from: AI chip manufacturing still relies upon electricity grids in East Asia primarily powered by fossil fuels (Wu et al., 2025), and, notably, the firm most frequently associated within this corpus with ecological-agent language was also identified as disclosing the least information regarding clean-energy use among suppliers (Greenpeace East Asia, 2025). The third finding relates to who bears the costs. State governments have provided billions of dollars in tax exemptions to data center operators to attract them, and utility companies continue to pass grid upgrade expenses on to ordinary electricity rate-payers (Brennan et al., 2025). A pattern was also reflected in congressional testimony on AI's increasing share of total electricity demand and its impact on rate-payers (Castor, 2026)

Discussion

Together, the four findings above describe one coordinated system rather than four separate problems. The dual-narrative structure (4.1) explains how a firm can sound like a climate innovator and, at the same time, be a climate liability within the same fiscal year without internal contradiction. This is possible because each register addresses a different audience under different legal and reputational constraints. Temporal deferral (4.2) makes this split consistent over time; it allows a hedged 10-K and a confident keynote to both remain true statements, since the keynote's claims are not subject to verification within the period covered by the 10-K. The promissory delta (4.3) shows what this arrangement produces empirically: an inverse relationship between material decoupling and promissory Ediscourse expansion. Apple is the only company that decreased its absolute emissions and also has restrained promissory registers.

In contrast, the firms with the boldest ecological-agent rhetoric, Meta and Google, exhibited the largest percentage increase in emissions during this period. The firms doing the least material decoupling produce the most extensive promissory discourse. The independent evidence (4.4) confirms that this is not a corpus artifact but a pattern observable from outside the firms.

This is the paper's response to its own repositioning question from Section 2.5. Both Perakslis (2020) and Ribeiro and Soromenho-Marques (2022) correctly identify that technology firms and technology narratives launder responsibility, but neither specifies the mechanism identified here. Perakslis' technowashing operates inside the present tense, on an interface, concealing exclusion that already exists. Ribeiro and Soromenho-Marquez' technowashing operates through spectacle, distracting attention towards a speculative elsewhere. The mechanism this paper identifies is neither concealment nor distraction in the strict sense; it is deferral. The claim is stated plainly, in a keynote, about real and already expanding infrastructures, and verification has been merely assigned to a future date when the firm controls the framing. This is a more difficult phenomenon than either concealment or spectacle for regulation, because nothing in it, at the moment it is spoken, is false. This will have direct implications for accountability frameworks. Disclosure regimes based upon retrospective verification, audited emissions, and measured water use must remain necessary, and this paper's independent evidence corpus depends upon them. Nevertheless, a firm can comply fully with every such requirement while still running the exact dual-narrative structure described in Section

4.1, because the promissory register's claims are never presented as the kind of claim disclosure regimes are built to check. A regulator reading only the 10-K would find nothing to challenge; a shareholder reading only the letter would find nothing to trust.

Recommendations

Three recommendations are proposed:

First, disclosure frameworks should require firms to distinguish promissory claims from claims that a reader can verify and audit within the current reporting period. This should be done by explicitly labeling them, regardless of the document genre.

Second, researchers extending this study should examine Microsoft's promissory claims to become "carbon negative" by 2030 along with its related "water positive" target. Meta's and Google's commitments to reduce greenhouse gas emissions by 2030 should be examined by comparing them with each company's actual outcomes as their deferred verification date approaches, enabling an empirically grounded examination of the temporal-deferral mechanism over time. Third, when journalists or civil society actors review a corporate AI-climate claim, they should first consider the document's generic nature and/or its intended audience. They should then determine if the same firm describes the same commitment using a different generic form in its quarterly and annual regulatory filings.

Conclusion

This paper explored how Big Tech companies discuss AI's environmental implications, whether that discourse differs by genre, and whether it coincides with their emission levels. Across the companies studied, a structure emerged in which all discourses were organized as promissory for public audiences and defensive for regulators, both discourses underwritten by postponing verification to an undetermined future. Confidence in the promissory discourse was not correlated with real-world emission trends, except in one instance. Apple alone among firms rejected the expansive ecological-agent register used by Google and Meta and confined its statements about reductions to particular, engineer-certified accomplishments. Apple was the only firm that demonstrated an actual total reduction of CO₂ emissions. This suggests that, compared with other firms in the study, restrained rhetoric (and not overconfidence) correlated with real performance. Because none of this requires any single statement to be false, retrospective verification-based disclosures cannot alone close the gap this paper identifies; a company can satisfy every audited requirement while the promissory register continues to do legitimating work no regulator can check.

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of Interest

The author declares no conflict of interest.

Author Contributions

The author conducted conceptualization, data collection, discourse analysis, and writing.

Data Availability Statement

The data supporting this study, publicly available corporate sustainability reports, keynote transcripts, shareholder letters, regulatory filings, and independent institutional and non-governmental reports, are available from the corresponding author upon reasonable request and are otherwise publicly accessible through the originating organizations.

References

- Alphabet Inc. (2025). Form 10-K: Annual report for the fiscal year ended December 31, 2024. U.S. Securities and Exchange Commission. <https://www.sec.gov/Archives/edgar/data/1652044/000165204425000014/goog-20241231.htm>
- Amazon Inc. (2025). 2024 Amazon sustainability report. <https://sustainability.aboutamazon.com/2024-report>
- Amazon Inc. (2026). Form 10-K: Annual report for the fiscal year ended December 31, 2025. U.S. Securities and Exchange Commission. <https://s3.amazonaws.com/sec.irpass.cc/3099/0001493152-26-016874.pdf>
- Apple Inc. (2024). 2024 environmental progress report: Covering fiscal year 2023. https://www.apple.com/ip/environment/pdf/Apple_Environmental_Progress_Report_2024.pdf
- Apple Inc. (2025). 2025 environmental progress report: Covering fiscal year 2024. https://www.apple.com/environment/pdf/Apple_Environmental_Progress_Report_2025.pdf
- Apple. (2026). Keynote (WWDC26). Apple Developer. <https://developer.apple.com/videos/play/wwdc2026/101/>
- Borup, M., Brown, N., Konrad, K., & Van Lente, H. (2006). The sociology of expectations in science and technology. *Technology Analysis & Strategic Management*, 18(3–4), 285–298. <https://doi.org/10.1080/09537320600777002>
- Brennan, K., Kak, A., & West, S. M. (2025). Artificial power: AI Now 2025 landscape report. AI Now Institute. <https://ainowinstitute.org/2025-landscape>
- Brown, N., & Michael, M. (2003). A sociology of expectations: Retrospecting prospects and prospecting retrospects. *Technology Analysis & Strategic Management*, 15(1), 3–18. <https://doi.org/10.1080/0953732032000046024>
- Burri, R. V. (2015). Imaginaries of science and society: Framing nanotechnology governance in Germany and the United States. In S. Jasanoff & S.-H. Kim (Eds.), *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power* (pp. 233–253). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226276663.003.0011>
- Castor, K. (2026). Prepared opening statement of Subcommittee on Energy Ranking Member Kathy Castor: Hearing on "AI and the grid: Meeting growing power demand while protecting ratepayers." House Committee on Energy and Commerce, Subcommittee on Energy. <https://www.congress.gov/119/meeting/house/119189/documents/HHRG-119-IF03-MState-A000370-20260429-U2.pdf>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Crawford, K., & Joler, V. (2018). *Anatomy of an AI system: The Amazon Echo as an anatomical map of human labor, data and planetary resources*. SHARE Lab & AI Now Institute. <https://www.anatomyof.ai>
- Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, 54(1), 64–87. <https://doi.org/10.1525/cmr.2011.54.1.64>
- Google. (2024). Google 2024 environmental report. <https://sustainability.google/reports/google-2024-environmental-report/>
- Google. (2025a). Google I/O '25 keynote. In YouTube. <https://www.youtube.com/watch?v=o8NiE3XMPrM>

- Google. (2025b). Google 2025 environmental report. <https://sustainability.google/reports/google-2025-environmental-report/>
- Greenpeace East Asia. (2025). Nvidia ranks last on AI supply chain decarbonization, Greenpeace report finds [Press release]. <https://www.greenpeace.org/eastasia/press/68382/>
- Jasanoff, S. (2015). Future imperfect: Science, technology, and the imaginations of modernity. In S. Jasanoff & S.-H. Kim (Eds.), *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power* (pp. 1–33). University of Chicago Press.
- Jasanoff, S., & Kim, S.-H. (2009). Containing the atom: Sociotechnical imaginaries and nuclear power in the United States and South Korea. *Minerva*, 47(2), 119–146. <https://doi.org/10.1007/s11024-009-9124-4>
- Jassy, A. (2025). Annual letter to shareholders: Parting thoughts. Amazon.com, Inc.
- Kenny, R. (2025). Report: Water use in AI and data centers. Government Digital Sustainability Alliance (GDSA).
- Leman, M. (2026). The energy and environmental impact of AI and how it undermines democracy. Greenpeace International. <https://www.greenpeace.org/international/story/82486/ai-energy-environment-democracy/>
- Meta Platforms, Inc. (2024). For a better reality: Meta 2024 sustainability report. <https://sustainability.atmeta.com/asset/2024-sustainability-report-executive-summary/>
- Meta Platforms, Inc. (2025). 2025 sustainability report: Connecting to a better reality (Reporting on fiscal year 2024). <https://sustainability.atmeta.com/2025-sustainability-report/>
- Meta Platforms, Inc. (2026, January 28). Form 10-K: Annual report for the fiscal year ended December 31, 2025. U.S. Securities and Exchange Commission. <https://www.sec.gov/Archives/edgar/data/1326801/000162828026003942/meta-20251231.htm>
- Microsoft Corporation. (2024). 2024 environmental sustainability report: Reporting on our 2023 fiscal year. <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Microsoft-2024-Environmental-Sustainability-Report.pdf>
- Microsoft. (2025a). 2025 environmental sustainability report: Accelerating progress to 2030 (Reporting on our 2024 fiscal year).
- Microsoft. (2025b). Form 10-K: Annual report for the fiscal year ended June 30, 2025. U.S. Securities and Exchange Commission.
- Microsoft. (2025c). Microsoft Build 2025 | Satya Nadella opening keynote. On YouTube. <https://www.youtube.com/watch?v=ceV3RsG946s>
- Nadella, S. (2025). Thinking in decades, executing in quarters. On LinkedIn.com. <https://www.linkedin.com/pulse/my-annual-letter-thinking-decades-executing-quarters-satya-nadella-7orpc/>
- NVIDIA. (2024). NVIDIA corporate sustainability report fiscal year 2024.
- NVIDIA. (2025a). GTC March 2025 keynote with NVIDIA CEO Jensen Huang. On YouTube. <https://www.youtube.com/watch?v= waPvOwL9Z8>
- NVIDIA. (2025b). Form 10-K: Annual report for the fiscal year ended January 26, 2025. U.S. Securities and Exchange Commission.

- Perakslis, C. (2020). Exposing technowashing: To mitigate technosocial inequalities. *IEEE Technology and Society Magazine*, 39(1), 10–11.
<https://doi.org/10.1109/MTS.2020.2967517>
- Pichai, S. (2026). Letter from the chief executive officer.
- Ribeiro, S., & Soromenho-Marques, V. (2022). The techno-optimists of climate change: Science communication or technowashing? *Societies*, 12(2), Article 64.
<https://doi.org/10.3390/soc12020064>
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571–610.
<https://doi.org/10.5465/amr.1995.950808033>
- Wu, K., Ip, J., Sung, Y., Ryu, R., & de Vries, A. (2025). Chipping point: Tracking electricity consumption and emissions from AI chip manufacturing. Greenpeace East Asia.
- Zuckerberg, M. (2025). Meta annual letter to shareholders.