İçtimaiyat
Sosyal Bilimler Dergisi | Journal of Social Sciences

Yazar(lar) / Author(s)

Assist. Prof. Dr. Muhammed Veysel Bilici
Kırklareli University, Faculty of Science and Letters, Department of Sociology, Kırklareli-Türkiye
e-posta: veyselbilici@gmail.com
(Sorumlu Yazar/Corresponding author)

Dr. Nurullah Karaca
Kırklareli University, Rectorate, Continuing Education Application and Research Center, Kırklareli-Türkiye
e-posta: nurullahkaraca@gmail.com

Makale bilgileri/Article information

Tür-Type: Araştırma makalesi-Research article
Geliş tarihi-Date of submission: 21.06.2026
Kabul tarihi-Date of acceptance: 17.10.2026
Yayım tarihi-Date of publication: 28.10.2026

Hakemlik-Review

Çift Taraflı Kör Dış Hakemlik
Double-Blind External Peer Review

Etik beyan/Ethics statement

Yazar(lar), çalışmanın hazırlanma sürecinde etik ilkelere uyduklarını beyan etmiştir.
The author(s) declared that they complied with ethical principles during the preparation of the study.

Benzerlik taraması- Plagiarism checks

intihal.net iThenticate

Çıkar çatışması-Conflict of interest

Çıkar çatışması beyan edilmemiştir
No conflict of interest declared

Finansman-Grant support

Fon, hibe veya destek alınmamıştır
No funding or support was received

Lisans- License

CC BY NC

Myth, Hegemony, and Rival Futures of Artificial Intelligence

Abstract

Technology leaders and their critics often describe the future of artificial intelligence through figures from myth, scripture, science fiction, and futurism. Drawing on Gramsci's concept of hegemony and Mosco's account of myth as common sense, we examine whether and how this mythic language sets the terms in which that future is argued, and who keeps or refuses those terms. Each of us independently coded a purposive set of English-language texts, selected for their mythic language, using directed content analysis. Some speakers embrace figures that others reject: for those who build the technology, the Singularity has begun, while for their critics it is only a story. Beneath this disagreement lies an assumption we call the terms of the debate, that the future depends on artificial intelligence. The technology elite keeps these terms most consistently, even though its members disagree over speed, safety, and control. Speakers outside the elite are more divided. Some keep the terms, fewer refuse them, and about half say nothing about them. The results point to a weak hegemony that works through shared terms without agreement on figures.

Keywords: Artificial Intelligence, Hegemony, Myth, Technology Elite, Directed Content Analysis.

Mit, Hegemonya ve Yapay Zekâya Dair Rakip Gelecekler

Öz

Teknoloji sektörünün liderleri de onları eleştirenler de yapay zekânın geleceğini çoğu zaman mitten, kutsal metinlerden, bilimkurgudan ve fütürizmden ödünç alınan imgelerle anlatmaktadır. Makalede, Gramsci'nin hegemonya kavramından ve miti bir tür sağduyu sayan Mosco'nun yaklaşımından hareketle, bu mitik dilin yapay zekânın geleceğine ilişkin tartışmanın zeminini ne ölçüde ve hangi yollarla kurduğu, bu zeminin kimlerce korunup kimlerce reddedildiği soruları cevaplanmaktadır. Mitik göndermeler içerdikleri için amaçlı olarak seçilen İngilizce metinler, yönlendirilmiş içerik analiziyle her iki yazarca bağımsız olarak kodlanmıştır. Aynı imgeler kimi konuşmacılarca benimsenmiş, kimilerince reddedilmiştir. Yapay zekâyı geliştirenler Tekilliği başlamış bir süreç sayarken eleştirmenleri onu bir anlatıdan ibaret görmektedir. Bu ayrılığın altında, tartışmanın zeminini oluşturan ortak bir varsayım, yani geleceğin yapay zekâya bağlı olduğu kabulü yatmaktadır. Teknoloji elitleri hız, güvenlik ve denetim konularında kendi aralarında ayrışsa da bu zemine oldukça tutarlı biçimde bağlı kalmışlardır. Elitler dışında kalan konuşmacılar daha bölünmüş bir görünüm sergilemektedirler. Bunların bir kısmı tartışmanın hâkim zeminini korumaktayken, daha küçük bir kısmı reddetmekte, yaklaşık yarısı ise bu konuda sessiz kalmaktadır. Bulgular, imgeler üzerinde uzlaşı olmasa da ortak bir zemine dayanan zayıf bir hegemonyaya işaret etmektedir.

Anahtar Kelimeler: Yapay Zekâ, Hegemonya, Mit, Teknoloji Elitleri, Yönlendirilmiş İçerik Analizi.

Yazar katkı oranı/ Author contributions

Çalışmanın Tasarımı/Conceiving the Study	M. V. Bilici (%50) - N. Karaca (%50)
Veri Toplama/Data Collection	M. V. Bilici (%50) - N. Karaca (%50)
Veri Analizi/Data Analysis	M. V. Bilici (%50) - N. Karaca (%50)
Makale Yazımı/ Article Writing	M. V. Bilici (%50) - N. Karaca (%50)
Gönderim ve Revizyon/Submission and Revision	M. V. Bilici (%50) - N. Karaca (%50)

Yapay zekâ beyanı / Artificial intelligence statement

Bu çalışmada yapay zekâ araçları (Claude Opus 5.5, Sonnet 5.5 ve Haiku 5.5; Anthropic) dil, anlatım ve iç tutarlılığın, kaynakça ve atıf biçiminin, verilerin ve istatistiksel sonuçların kontrolünde kullanılmıştır. Araştırma soruları, kuramsal çerçeve, yöntem, metinlerin bağımsız kodlanması, analiz ve sonuçlar yazarlara aittir. Yazarlar yapay zekâ desteğiyle elde edilen bütün çıktıları gözden geçirip doğrulamıştır ve makalenin içeriğinden bütünüyle sorumludur.

In this study, AI tools (Claude Opus 5.5, Sonnet 5.5, and Haiku 5.5; Anthropic) were used to check language, style, internal consistency, reference and citation formatting, and the data and statistical results. The research questions, theoretical framework, method, independent coding of the texts, analysis, and conclusions are the authors' own. The authors reviewed and verified all AI-assisted outputs and take full responsibility for the content of the manuscript.

Atıf- Citation (APA)

Bilici, M. V. & Karaca, N. (2026). Myth, hegemony, and rival futures of artificial intelligence. *İçtimaiyat*, *10*(Özel Sayı),45-62. https://doi.org/10.33709/ictimaiyat.1975963

ISSN: 3149-8655
E-ISSN: 2602-3377

1. Introduction

The future of artificial intelligence is often foretold rather than argued. Members of the technology elite promise abundance or warn of the end of humanity, and many of them speak as if that future were already decided. Sam Altman's (2025c) essay "The Gentle Singularity" opens in the language of physics and futurism: "We are past the event horizon; the takeoff has started." In the same year, Peter Thiel, by contrast, warned in off-the-record lectures on the Antichrist that Armageddon is coming but said that the outcome is "obviously not written in stone" (Bhuiyan et al., 2025). The Singularity and the Antichrist are examples of *figures*, that is, names and images drawn from myth, scripture, science fiction, or futurism and applied to technology, and the industry's critics often answer with figures of their own. Taken together, such borrowings make up what we call *mythic language*. When it comes from the industry's own leaders, it may do more than describe the future. This possibility raises the main question of this paper: what the mythic language surrounding artificial intelligence does in the debate on its future, and who keeps or refuses the terms in which that future is argued.

What is at stake is how much room public debate leaves for decisions that do not concern artificial intelligence (AI) itself. They concern who profits from the systems and the data they are built on, how employers use them to reorganize work, and who answers for the harm done with them. If every option under discussion is an option about AI, these decisions may reach public debate only at its edges and be settled largely by technology firms and their investors. Those who build and fund AI would then lead the debate without having to win its arguments. How a group can lead a debate in this way is an old problem of hegemony, and the AI debate is a current case of it.

Research on myth and technology has long traced technology's double image as salvation and as doom (Geraci, 2010; Mosco, 2004; Noble, 1997). For AI, recent work finds a common narrative in national strategies, among experts, and in the press, in which AI is inevitable (Bareis & Katzenbach, 2022; Köstler & Ossewaarde, 2022) and its development a race (Köstler & Ossewaarde, 2022; Richter et al., 2025). Köstler and Ossewaarde (2022), whose study is closest to ours in method, follow the German government's vision of AI into the national press. They find that the media "largely adapt the government's frames" (p. 249), that is, its schemata for interpreting the future of AI, and they call the enforcement of such frames by governments and media "an exercise of hegemony" (p. 250).

Some of this work draws on the concept of sociotechnical imaginaries, which Jasanoff (2015) defines as "collectively held, institutionally stabilized, and publicly performed visions of desirable futures" (p. 4). Such a vision can start as what Hilgartner (2015, as cited in Jasanoff, 2015) calls a "vanguard vision" (p. 4), held by an individual or a small group, and become an imaginary only when it is "communally adopted" (p. 4). Where earlier studies find a common narrative of AI's inevitability, the concept suggests a different reading of the AI debate. The prophecies of the technology elite resemble vanguard visions, which may be held most firmly by those who promote them and adopted only in part beyond them, and the refusals of other speakers can be read as resistance. A single vision, moreover, can contain both a promise and a threat, because imagined futures come with "shared fears of harms" (p. 5).

Critical accounts of AI take this point further, suggesting that hope and fear are argued in the same terms, in which AI itself, rather than those who build, own, and govern it, decides what

comes next. Crawford (2021) calls utopian and dystopian talk about AI "metaphysical twins," since "each offers a profoundly ahistorical view that locates power solely within technology itself" (p. 214). For O'Sullivan (2025), speculative scenarios of superintelligence "acquire the weight of destiny" through constant repetition, a process he calls "the manufacture of inevitability," and he traces the narrative to "particular corridors of power," where "the loudest prophets of superintelligence are those building the very systems they warn against." Less is known about which speakers keep these terms and whether the elite holds them more firmly than others, including industry practitioners.

In this study, we analyze a corpus of English-language texts on technology and AI, selected for their mythic language, asking whether the speakers who invoke the figures share a two-part premise about AI's future. The first part is that artificial general intelligence (AGI) is near. The second is that the future depends on AI, whether on the choice to build it or stop it, on what kind of AI is built, or on what AI is expected to do on its own. In each case, AI is treated as the main force that shapes what comes next. We call the second part *the terms of the debate*, an assumption that rival sides can share even as they disagree. Speakers may thus dispute whether AGI is near and still keep these terms. Analyzing both parts lets us compare disagreement over the first with the degree of agreement on the second.

Separating the two parts also allows us to examine what we call a *hegemony of terms*, that is, leadership that rests on a debate's terms rather than on agreement about its stories. The study contributes this concept and a comparison of how firmly the elite and other groups of speakers hold the terms, and both bear on Mosco's (2004) account of how a technology comes to look inevitable.

The main question breaks down into three research questions. The first asks what stories the mythic language borrows and what their promises and threats take for granted (RQ1). The second concerns how members of the technology elite use the figures, including against one another, and whether they share the same terms (RQ2). The third asks whether speakers outside the elite affirm or doubt that AGI is near, and whether they keep or refuse the terms (RQ3).

2. Conceptual Framework

2.1. A Hegemony of Terms

To ask who holds a set of terms, and who refuses them, is to ask about hegemony. According to Gramsci (1971), the supremacy of a social group shows itself in two ways, "domination" over its opponents, if need be by force, and "intellectual and moral leadership" over its allies (p. 57). The decisive step comes when the group achieves "intellectual and moral unity, posing all the questions around which the struggle rages not on a corporate but on a 'universal' plane" (pp. 181–182). We borrow the formula, not Gramsci's theory of class rule, and take from it the idea that a group can lead a debate without winning every argument in it, as long as the debate is held on terms that the group has set. Gramsci adds that ideologies that are historically necessary, unlike arbitrary or "willed" ones, organize human masses and "create the terrain on which men move, acquire consciousness of their position, struggle" (p. 377). We read the terms of the AI debate as such a terrain, and the mythic language as one means of struggle on it. Whether the "great masses" consent to the group's leadership (p. 12) is a question about audiences. Our concern, however, is with those who speak about AI in public, from the

technology elite to journalists, academics, and religious leaders. All of them take part in forming opinion on AI, some far more than others.

Building on Gramsci (1971), Williams (1977) argues that hegemony is never finished. It has to be renewed and defended, and it is continually resisted, so any analysis of it must include those who stand "outside or at the edge of the terms of the specific hegemony" (pp. 112–113). We take the word "terms" from Williams, but where he means the terms of a whole social order, we mean the shared assumption on which a single debate, here the debate on AI's future, rests. Williams also warns that opposition can be "in practice tied to the hegemonic," although "authentic breaks" with a hegemony have "often in fact occurred" (p. 114). Since a break can begin on a small scale, the few who refuse the terms deserve attention.

In debates on the course of AI, then, the terms are the second part of the premise introduced above, the assumption that the future depends on AI. We distinguish three main forms of this assumption, which a single speaker may combine. In the first, the future is a choice about whether to have AI at all, that is, whether to build it before others do or to stop it before it is too late. The second form starts from AI's arrival, and speakers argue over how to make it safe, whether it should be open or closed, who should control it, and who will win the race. Even a dispute over control is then about which AI the future will bring. The third rests on the expectation that AI will decide the future on its own, by changing everything or by taking control. A speaker refuses the terms when they say that the future depends on something other than AI, such as ownership, politics, or work, or that AI is not the real question. The other part of the premise, whether AGI is near, is a separate issue, since a speaker who doubts it can still keep the terms.

2.2. Myth, Leadership, and Consent

Williams (1977) offers concepts for describing what the mythic language takes from the past (RQ1). He calls residual those older forms that, formed in the past, are still at work in the present (p. 122). Newer forms are harder to classify, since it is difficult to tell whether they belong to a new phase of the dominant culture or are *emergent*, that is, alternative or oppositional to it (p. 123). The distinction separates the figures inherited from scripture and myth from the futurist ones. We also borrow Williams's notion of a selective tradition, which he develops for a whole culture and which we apply to the figures of a single debate. Such a tradition keeps only some parts of the past (p. 115) and uses them "to ratify the present and to indicate directions for the future" (p. 116). What a figure is chosen to promise or threaten may therefore reveal the assumptions of those who use it. Such assumptions can also come to seem natural, a process that Mosco (2004) ties to Gramsci's (1971) notion of common sense, the sediment that earlier ideas leave in everyday thought, when he writes that "myth is congealed common sense" (p. 29). Writing about cyberspace, Mosco identifies three teleological myths, "the end of history, the end of geography, and the end of politics" (p. 13). The third is particularly close to the terms of the AI debate, because a future that follows a fixed technological trajectory leaves little for politics to decide. For Barthes (1972), myth gives "an historical intention a natural justification" and makes "contingency appear eternal" (p. 142). In this way, prophecies of AGI can present a set of corporate decisions as destiny.

We expect the terms to be held most firmly by the *technology elite* (RQ2), by which we mean the founders, top executives, and leading investors of major technology firms. Its members,

however, often quarrel. Hall (1979) speaks of "a struggle for hegemony, within the dominant bloc" (p. 16), a contest for leadership among groups that already share power, and we adapt the idea to the quarrels within the technology elite. In such quarrels, one side may discredit the other's view by calling it a myth. Johansen, following Lloyd (1990, as cited in Johansen, 1999), identifies this strategy in Greek thought, where calling a rival's account a myth marked it as fiction and tended to hide what the two shared (pp. 279–280). Greek myth also supplies a name for the technology elite. The essayist Léon de Sailly (2025) calls its members the "New Olympians," arguing that tech billionaires "see themselves as modern-day Olympians, bringing the Promethean fire of progress to humanity." We develop the name into a critical heuristic with three points. The first is the posture of a forecaster who speaks about everyone's future. Forecasts made from that posture can help to make themselves true (Merton, 2016, p. 506). The second, in our use, is the shared terms within which members quarrel, as the gods of Olympus did among themselves. The third is that the name, transferred from myth onto a group of people, is itself a figure. We therefore ask of it, and of images like it, the same questions as of the other figures, namely who uses them, against whom, and what they assume. The heuristic thus raises the question of whose terms these are, without assuming that the elite is a class with a single will or a common direction.

For Gramsci (1971), intellectuals are the dominant group's "deputies" (p. 12) in winning consent. Speakers outside the elite (RQ3), from journalists to academics, may act as such deputies and pass the terms on, or they may refuse them. Abercrombie et al. (2015) object, however, that a dominant ideology helps to explain "the coherence of the dominant class but not . . . the coherence of a society as a whole" (p. 3). We therefore read the voices outside the elite as closely as those of the elite, and we return to this objection in the Discussion.

3. Method

Our research questions follow from the framework set out above, so we address them with a directed qualitative content analysis (Hsieh & Shannon, 2005), in which the central codes are derived from theory.

The basic unit of analysis is the *allusion*, which we define more narrowly than in everyday usage, as a passage in which a speaker invokes a figure from myth, scripture, fiction, or futurism and transfers its meaning onto something in the technology world. A figure is the name or image that a speaker invokes. A *speaker* is any person, group, or firm whose statements a text presents, whether as its author, as someone it quotes or reports, or as a participant in an interview or a recorded discussion. We count an allusion only if it meets three conditions. First, the figure can be traced to an identifiable tradition or work (*source*). Second, the figure is applied to the industry's leaders, firms, products, the technology itself, its future, or those who regulate, use, or criticize it (*target*). Third, replacing it with a plain word loses a claim, not only a stylistic effect (*transfer*). Idioms such as "holy grail" count only if they are developed in the text, and the names of products, models, projects, and firms count only when they are coined or when their meaning is part of the claim. Mythic language here means all the allusions taken together.

From our personal archives of web pages, we selected the 1,094 texts that relate to AI, Silicon Valley, and the technology industry. We then screened them together for possible allusions as defined above and kept the 353 English-language texts that might contain one. These texts

date from 2000 to mid-2026, most of them from 2025 onward. They include newspaper and magazine articles, newsletters and blog posts, company and personal websites, and transcripts of podcasts and videos. Most of the texts set aside used mythic words only in dead idioms, such as “holy grail” or “tech titans,” or in passing references. The rest were works of fiction, forum comment threads, texts without usable English content, or texts published before 2000. Because the corpus is purposive, chosen for its mythic language, no share in this study can be generalized to technology discourse at large. The findings are based on these 353 texts. We also consulted other web sources to place them in context, and the Findings quote several of these, either to cite firsthand a document that corpus texts report on or to show what a speaker said elsewhere or later. None of these sources was coded or counted.

Each allusion in the corpus received two kinds of code. Descriptive codes record the figure, its target, and its *register*, the broad tradition from which it comes. Two analytic codes come from the framework. The code *valence* records what the figure says its target brings (promise, threat, both, or neither). Promise and threat stand for paired opposites that run through the texts, such as utopia and dystopia, salvation and doom, hope and fear, or opportunity and risk. The code *use* records what the speaker does with the figure (embraces, describes, rejects, or inverts it).

Speakers who invoke a figure in a text also received two codes for that text, one for each part of the premise, judged only from their own words, as quoted or reported. One records whether the speaker says that AGI, as they define it, is near, that is, already here or coming within a few years or a decade or two (affirms, doubts, or silent). The other records the speaker’s stance on the terms (keeps, refuses, or silent). A speaker who both keeps and refuses the terms in the same text is coded as refusing.

Every text was coded independently, while speaker roles were assigned jointly. The allusions were first listed separately, and any differences in their descriptive codes were then settled together. Valence and use, however, were coded by both of us for all allusions. Where our codes for valence, use, or either part of the premise differ, both codes are retained rather than resolved. We measured agreement between our independent codes with Cohen’s kappa. It was lowest for use (.71) and for the stance on the terms (.70), and higher for valence, for whether AGI is near, and for the register and target codes.

To classify the uses, we adapt Hall’s (2005) reading positions, though not his model of how audiences receive a message. We combine each use with the same speaker’s stance on the terms in the same text and assign the positions in a fixed order. First, a use coded as inverting a figure, that is, turning someone else’s use of it against them, counts as re-aiming, whatever the speaker’s stance. Second, a speaker who refuses the terms takes the oppositional position, which we call refusing, with any use of the figure. Third, a speaker who keeps the terms and embraces a figure carries it, which is the dominant-hegemonic position, and one who keeps the terms but rejects a figure debunks it, which is the negotiated position. Re-aiming falls between the negotiated and the oppositional positions, so we count it separately. We assign no position to any other use, such as a description of a figure or a use by a speaker silent on the terms.

The codes guide our reading of the texts rather than replace it, and the argument rests on that reading, not on the counts. The shares in the Findings are calculated over the cases on which

our codes agree, and disputed cases are examined separately in the discussion of refusal. The contrast between the elite and other speakers also appears in each coder's codes taken alone.

Speakers outside the technology elite fall into the five groups of Table 2. Industry practitioners work in the technology industry without being among its founders, top executives, or leading investors, and the press comprises journalists, editors, and news organizations. Academics teach or do research at universities and research institutes, while independent writers publish under their own names outside a technology firm, a university, or a news organization. Religious leaders are counted together with speakers who fit no other group, such as novelists and politicians. Each speaker is counted once for the terms, by the rule given in the note to Table 2, and whether AGI is near is treated in the same way. Refusing is defined as saying that the future depends on something other than AI, or that AI is not the real question. Those who refuse the terms thus turn away from AI by definition and may point instead to ownership, politics, work, or other concerns. The same definition makes keeping the terms the expected stance of anyone who speaks about AI, so the results lie chiefly in how the groups differ, not in how many speakers in any single group keep the terms.

4. Findings

Tables 1 and 2 summarize the allusions by register and the speakers by group. Futurism and science fiction each supply about a third of the allusions, and the Bible about a quarter. Classical myth and folklore account for most of the rest, and non-Western traditions very few. About three in five allusions are aimed at technology or a product.

Table 1: Registers of the Mythic Language in the Corpus

Register	Share of allusions	Typical figures	What the figures mostly promise or threaten	Use by the technology elite	Use by other speakers
Futurist	32%	the Singularity, the intelligence explosion, the paperclip maximizer	threat in about half, promise in a quarter	half of the elite's figures, usually embraced as forecasts	rejected in about two in five uses, often as hype or marketing, and embraced about as often
Science fiction and fantasy	31%	Skynet, Frankenstein, Orwell's *1984*, the Matrix, the Ring of Power	threat in seven in ten	a quarter of its figures, mostly warnings about AI, often from speakers who voice hope elsewhere, with Orwell aimed at the state	mostly embraced as warnings about AI and the tools built on it, rejected in a quarter of uses
Biblical	27%	the Tower of Babel, the Antichrist, salvation, the end times, the pale horse of Revelation	threat in four in ten, many neither	a quarter of its figures, the Antichrist aimed at the safety camp and the state, the demon at AI itself	mostly embraced to warn, or turned into labels for the industry's faith in itself, rejected in about a third of uses
Classical and European folklore	9%	the oracle, Prometheus, Icarus, the Hydra	threat in about half	almost absent	mostly embraced as cautionary tales, rejected in one use in ten
Non-Western	1%	the "destroyer of worlds" from the Bhagavad Gita, the Mandate of Heaven from Chinese political thought, the *jianghu* of Chinese martial-arts fiction	mixed	almost absent	used mostly by writers and journalists, often to show borrowings beyond the Anglophone canon, too few uses to compare

Note. Shares are of all 1,085 allusions in the corpus. Promise, threat, and use are given for the allusions on which our codes agree. The safety camp warns that advanced AI could escape human control.

The sections that follow take the research questions in turn, beginning with the figures themselves (RQ1), then the members of the elite (RQ2), and then the speakers outside it (RQ3).

Table 2: Speaker Groups and the Terms of the Debate on AI

Speaker group	*n*	Keep the terms	Refuse them in at least one text	Silent
Technology elite	24	63%	4%	33%
Outside the elite				
Industry practitioners	34	29%	9%	62%
Press	65	18%	12%	69%
Academics	75	24%	29%	47%
Independent writers	59	37%	27%	36%
Religious and other speakers	61	26%	13%	61%
All speakers outside the elite	294	27%	19%	54%

Note. The terms are the assumption that the future depends on AI. Each speaker is counted once. Keeping the terms means keeping them in at least one text and refusing them in none, and a speaker who both keeps and refuses them in the same text counts as refusing them. Speakers are left out when our codes for the terms never agree. Percentages may not sum to 100 because of rounding.

4.1. Old Stories, New Machines

We group the figures into four clusters by what they borrow and how they are used. The clusters cut across the registers, and we call them *futurist prophecy*, *apocalypse and the demonic*, *science fiction*, and *the vocabulary of faith*. Futurist prophecy rests on forecasts that hardened into named scenarios, such as the Singularity and the intelligence explosion. Such figures can promise transcendence or threaten extinction. Apocalypse and the demonic draws instead on scripture, on judgment and the end, through the Antichrist, Armageddon, and the summoned demon. Science fiction supplies figures of creations that turn against their makers and of powers that watch everyone, such as Skynet, Frankenstein, and Orwell's *1984*, and its figures are the ones most often cast as a threat (Table 1). The vocabulary of faith differs from the other three clusters, because it describes the people who promote AI rather than what AI might bring. Words such as "cult," "prophet," "church," and "pontiff" come mostly from speakers outside the elite, who aim them chiefly at the industry's leaders. Members of the elite also use them against their critics, as when the *safety camp*, those who argue that advanced AI could escape human control, is called a cult.

Across the corpus, threat outnumbers promise by more than two to one. Yet the split between promise and threat hides an agreement. The journalist Karen Hao observes that "boomers" and "doomers" alike "believe that AGI is possible, and it's imminent, and it's going to change everything" (Firth, 2025). Hao's description contains both parts of the premise and attributes them to both camps. For Williams (1977), a dominant culture's selective tradition remains vulnerable, because what it discards can be recovered (pp. 115–116). Here both sides select openly from a common stock. The elite draws mostly on futurism, and Thiel on the Antichrist, while the industry's critics draw more on science fiction and Babel.

The mythic language, then, borrows widely, and its figures mostly present threats and pass easily between opposing sides. Usually, AI itself or the firms that make it bring the promise or the threat, and the future thereby appears to depend on AI. A figure that both the elite and the industry's critics can use does not tie its user to the terms. Any agreement among speakers would therefore lie in the terms rather than in the figures, which raises the question of who keeps the terms, and how consistently.

4.2. On Olympus: "Living Through the Singularity"

Shared terms might seem least likely within the elite, since its members fight over who controls AI and call their critics cultists, panic-mongers, and servants of the Antichrist. Such fights, however, are the kind of struggle within the dominant bloc that Hall (1979, p. 16) describes, and they need not reach the terms. Members of the elite embrace their figures in more than eight of every ten uses, and most members keep the terms (Table 2), while the only one who refuses them, Thiel, keeps them in other texts.

Futurist figures serve the elite as forecasts of its own work. Jack Clark (2026), a co-founder of Anthropic, writes that "we might soon be able to build an AI system that may be smart enough to develop its own successor, thus kicking off a process of recursive self-improvement." The question he struggles with, Clark adds, is "how do I get my mind right with living through the singularity?" When members of the elite warn, the danger they name is mostly AI itself, whether a rogue system or a summoned demon, and such a threat still leaves the future dependent on AI.

A forecast can also help to bring about what it predicts, all the more when the forecaster can pay for it. The technology critic Evgeny Morozov (2025) pictures Silicon Valley's leaders as prophets who "propose futures Monday, finance them Tuesday, and force their manifestation by Friday." Morozov's image differs from the self-fulfilling prophecy as Merton (2016) defined it. In Merton's account, a false belief comes true through the behavior that it evokes in those who hold it, whereas in Morozov's satirical image the forecasters themselves bring their forecasts about, with capital.

Other figures are aimed at people rather than at technology. The Antichrist and the *katechon*, the power that restrains him, belong almost entirely to Thiel, whose lectures, in the reporters' words, treat a "one-world state" as "essentially identical to the antichrist" (Bhuiyan et al., 2025). His Antichrist, Bhuiyan et al. (2025) write, "cultivates a fear of existential threats such as climate change, AI and nuclear war" in order to amass power. The danger Thiel foresees is thus a government strong enough to stop technology, an idea that Schmitt (2005), whom Thiel counts among his influences (Bhuiyan et al., 2025), would read as political theology.

David Sacks, a venture capitalist who served as the White House adviser on AI, invoked Orwell against Pope Leo XIV's 2026 call for more regulation of AI. Sweeping state power over AI, he warned, could be used "to censor, surveil, and control citizens — as Orwell foretold in '1984'" (Sacks, 2026, as cited in Metz, 2026). The investor Marc Andreessen (as cited in Marantz, 2024) accused "the AI risk cult" of a "full-blown moral panic." The Antichrist, Orwell's *1984*, and the cult are aimed at the same targets, namely the regulator, the state, and the safety camp. In Sacks's version, moreover, the dystopia that critics use for platforms describes the government that would regulate them. These figures are rarely aimed at fellow members of the elite. Even Thiel, who describes Bill Gates as "a Dr Jekyll and Mr Hyde-type character," denies that Gates or Andreessen could be the Antichrist (Bhuiyan et al., 2025).

Rivalry also sets members of the elite against one another. Elon Musk and Sam Altman, for instance, fought over who would control OpenAI, and according to Farrow and Marantz (2026), Musk "demanded majority control" during "discussions about whether to reconstitute OpenAI as a for-profit company." The founders of OpenAI, however, held that AI could bring "a post-scarcity utopia" or devastation that "could be total" (Farrow & Marantz, 2026). Musk and

Altman's feud may hide the shared belief that made control of OpenAI worth fighting over. When the conflict went to trial in 2026, the *New York Times* reporter David Streitfeld (2026) likened the trial to "one of those episodes where the Greek gods descend to mundane Earth to settle a dispute," an image close to the "New Olympians" of our heuristic.

The quarrels inside the elite are real, but they concern the timetable and the promise, not the terms. Members of the elite disagree on whether AGI is near and on what AI will bring. Dario Amodei, Anthropic's chief executive, gave "maybe 70 percent or 80 percent" odds that AI systems "much smarter than humans at almost everything" would arrive before the end of the decade (Roose & Newton, 2025). Sacks (2025), for his part, declared that apocalyptic predictions of job loss "are as overhyped as AGI itself," yet the future he describes is a race that competition "helps America win." Musk, on the other hand, warns of danger through his figures, but he has said that "you could sort of think of humanity as a biological boot loader for digital super intelligence" (as cited in Geuter, 2026). Demis Hassabis, then the head of Google DeepMind, voiced hope instead, saying he would be "very worried about society today" if he did not know that something "as transformative as AI was coming down the line" (Perrigo, 2025). Whether they expect, doubt, warn, or hope, each makes the future depend on AI.

Thiel, the one partial exception, calls AI a symptom of a wider stagnation in science and technology and dismisses "crazed corporate utopianism" and "effective altruist luddism" alike (Bhuiyan et al., 2025). In this report he refuses the terms, since he tentatively expects AI to matter roughly as much as the internet did in the 1990s and presents the future as a choice between a one-world state and Armageddon. His view may have shifted since. In a later interview, he suggested that AI is perhaps "bigger than the internet" and could reaccelerate growth (Döpfner, 2026, 2:46).

The meaning of AGI is itself unsettled within the elite. It differs from speaker to speaker and shifts even within one speaker's statements. In January 2025, Altman (2025a) wrote that OpenAI was "now confident we know how to build AGI as we have traditionally understood it." A month later he called AGI "a weakly defined term" while offering a working definition of his own (Altman, 2025b). Andreessen, in a podcast episode released on May 19, 2026, described AGI as "an AI that's as smart as a person" and said, "I actually think we crossed that about 3 months ago" (Rogan, 2026, 1:44:31). By contrast, a leaked agreement between OpenAI and Microsoft, as the sociologist Alex Hanna describes it, equated AGI with "the tool that generates $100 billion in profit" (Jaffe, 2025). As Willison (2026) summarizes the clause, Microsoft's commercial rights to OpenAI's technology would cease once AGI was achieved. In this contract, at least, the definition of AGI was tied to ownership, although Willison reports that the clause appears to have ended in April 2026.

4.3. Beyond Olympus: "This Ever-Present Goalpost"

Outside the elite the pattern differs, above all because so many speakers are silent. In their uses of figures, these speakers embrace about half and reject about a third. On the terms, about a quarter keep them, about one in five refuse them in at least one text, and about half are silent on them (Table 2). Journalists are among the speakers most often silent on the terms, and the press describes the figures it reports far more often than any other group, passing them on without endorsing or rejecting them. In 2026, Pope Leo XIV used the Tower of Babel to warn the AI industry, and the reporter Cade Metz (2026) noted that Jeremy Nixon, who once

worked at Google, “chuckled at the pope’s invocation of the Tower of Babel.” At Google, Metz adds, “a project the company called Babel built translation technology that has essentially allowed everyone to speak the same language,” so in one article the figure serves both as a warning and as the company’s own name for its work. Metz leaves the reader to weigh the two. Part of the silence may be a matter of genre, since a news report is not written to take a side.

Industry practitioners are also frequently silent (Table 2). Of all groups outside the elite, they refuse the terms least, yet they keep them much less often than members of the elite do. In these texts, then, those who work in the industry keep the terms at about the same rate as other speakers outside the elite.

Most speakers outside the elite do not say when AGI will come, and among those who address the question, speakers who doubt that AGI is near outnumber those who expect it soon by nearly two to one. The two parts of the premise tend to be answered together, since speakers who say AGI is near usually keep the terms, and most of those who doubt it refuse them. The exceptions show that the timetable and the terms can come apart. The technology writer Kevin Kelly (2025) argues that “there has been no exponential increase in artificial intelligence,” so that “we have time,” yet he expects AI to demand “a similar redesign of roles and spaces” to the one that electrification brought to organizations a century ago. Those who warn of extinction show the same separation, this time among speakers who expect AGI soon. Daniel Kokotajlo, the lead author of *AI 2027*, welcomed the lengthening of his own timelines to AGI and added, “The situation is still basically the same, though” (Lovely, 2025). The danger now lay further off, but the choice it posed was unchanged.

Uses that fit none of the four positions described in the Method, together with the disputed cases, make up more than half of all uses outside the elite, so the shares that follow rest only on those that can be given a position. When uses rather than speakers are counted, from about a quarter to more than half of these come from speakers who refuse the terms, depending on whether a text in which a speaker both keeps and refuses them counts as refusal. Carrying accounts for most of the rest, and re-aiming is rare. The elite carries its figures far more often.

No speaker outside the elite carries the figures more often than Eliezer Yudkowsky, a leading voice of the safety camp, and none shows more clearly that carrying a figure need not mean endorsing what the industry promises. In a 2025 interview he says that current systems will not yet try to break out and “start hunting down humans,” adding, “They are not quite that smart yet, as far as I can tell” (Klein, 2025). He takes up the rogue system and the self-improving machine as warnings, and the remedy he proposes, to stop building such systems, is still a choice about AI. As Williams (1977, p. 114) cautions, opposition of this kind can remain bound in practice to the hegemony it resists. The futurist figures and the terms thus pass from the elite to its most alarmed critics.

Of the registers large enough to compare, the futurist register is the one most often rejected outside the elite, in about two in five uses, roughly as often as it is embraced (Table 1). Some of these rejections keep the terms and thus debunk the figures rather than refuse them. Kelly (2025), for instance, writes that the doomers “promote the most extreme fantasy version of AI” and that “they believe the hype,” yet the objection concerns the forecast and leaves AI’s place in the future untouched. Even Skynet, the best-known figure of science fiction, is rejected more often than it is embraced. As the journalist Andrew Marantz (2024) notes, the basic arguments

of the safety camp "do not require imagining that the current crop of Verizon chatbots will suddenly morph into Skynet." The safety camp's case, in Marantz's account of it, rests on an AI whose "objectives are at odds with human flourishing," not on a sentient one, and in our reading the cinematic killer robot only makes it seem fictional.

Refusal comes most often from academics and independent writers (Table 2). Those who refuse the terms point above all to the firms, owners, and capital behind AI. Hao describes AGI as "this ever-present goalpost that keeps shifting, depending on where the company wants to go" (Firth, 2025). Morozov (2025) argues that "the tech oligarchs," not their platforms, "present the greater danger," and Pope Leo XIV's (2026) encyclical calls for "naming the new monopolies of AI" (para. 109). Politics and the law come next, and work, religion, the land, and ideas less often.

Many of those who refuse the terms also say what the figures do for those who use them. O'Sullivan's (2025) essay imitates the way each new product is turned into prophecy, "*Claude solves coding problems; the singularity is near*," and argues that "predictions of machine godhood reshape how authority is claimed and whose interests steer AI governance." The forecast thus becomes a credential, since the forecaster who has predicted the technology claims a say in how it is governed. The writer Om Malik (2026) finds such a credential in Amodei, whom he calls "the prophet of catastrophe and the prophet of salvation at once," adding that both claims "require you to trust his judgment." The same charge extends to science fiction. For Brian Merchant (2026b), an outspoken critic of the tech industry, Musk's "biggest innovation of all" has arguably been "the relentless science fictionalization of capitalism itself," a habit he also finds in Anthropic's warnings about its Mythos model, which he reads as a "doom marketing campaign."

Morozov (2025), like others who refuse the terms, turns the vocabulary of faith against the industry's leaders and their supporters, calling the leaders "Silicon Valley pontiffs," served by "court hagiographers." The two images present the leaders as shielded from criticism, since a pontiff need not argue and a hagiographer does not question. The name "New Olympians," which de Sailly (2025) gave to tech billionaires, belongs to the same family, since it too casts the leaders as people whose prerogative to speak for the future needs no defense. This vocabulary can also explain. Hao notes that the term AGI creates a "quasi-religious fervor" around OpenAI's work, a fervor that keeps people working toward a horizon that moves (Firth, 2025).

Some speakers can be read in two ways, as keeping the terms or refusing them, and as speaking about the future or not. These are the cases on which our codes differ. Such boundary cases are slightly more frequent among academics and the press than among other groups, and with biblical than with futurist figures. The clearest case is the encyclical, which warns that without law, oversight, and politics, change will be "presented as necessary and inevitable," under rules "shaped by those who control data, infrastructure and computing power" (Leo XIV, 2026, para. 106). Yet its central choice, that of "constructing Babel or rebuilding Jerusalem" (Leo XIV, 2026, para. 9), is also a choice about how AI is built, and press reports present the Pope as both keeping and refusing the terms. Like Thiel's Antichrist, the Pope's Babel names a danger in the concentration of power. The encyclical, however, sees "the modern Babel" not only in "the globalized technocratic paradigm" but also in the clash of powers that seek to preserve or seize "supremacy," and in a race "to develop evermore

powerful technologies or to secure control over them" (Leo XIV, 2026, para. 185). It also calls for the very regulation that Thiel fears. In these texts, no member of the elite invokes Babel. Here a clerical voice, Gramsci's (1971, p. 7) type of the traditional intellectual, criticizes the terms from an institution old enough to have a vocabulary of its own, and still speaks partly within the debate's terms.

The fourth position, re-aiming, is clearest in an exchange over the Ring of Power from Tolkien's *The Lord of the Rings*. Altman (2026) wrote that AGI has "a real 'ring of power' dynamic to it," the Ring being "the totalizing philosophy of 'being the one to control AGI,'" and proposed "sharing the technology with people broadly, and for no one to have the ring." For him, the danger lies in who controls AI, and the remedy is wider access to the technology under democratic control. Merchant (2026a) answered with the end of Tolkien's story, writing that "the only way to ensure evil doesn't triumph is to chuck the One Ring itself into Mount Doom." He keeps Altman's figure and turns it against him. In the novel the Ring is destroyed rather than shared or held in better hands, so the lesson is to end the race, not to win it. Re-aiming can also come from inside the safety camp, and from a speaker who elsewhere carries the figures. Yudkowsky recalls that Musk called AI "summoning the demon," and turns the image against the founding of OpenAI, saying that "apparently the solution on this was to give everybody their own demon," which "doesn't actually address the problem" (Klein, 2025).

5. Discussion

Mythic language runs through these texts, yet its figures do not tie speakers to a side. What holds instead is a less visible agreement, that the future depends on AI, kept most firmly by the elite and in part by speakers outside it. How the terms of a debate can hold while its figures keep changing sides is the puzzle that the findings leave.

5.1. Contested Figures, Shared Terms

Part of the answer lies in the age of the figures. Many of them are residual, in Williams's (1977) term, formed in scripture, myth, and older fiction long before the debate on AI. Because no side of the debate created them, no side owns them, and any side can take them up. The futurist figures pass between sides for another reason. As new forms, they are not easily classified, since it is unclear whether they are merely novel or truly emergent (Williams, 1977, p. 123). Here the speakers themselves dispute where the figures belong, and a forecast means one thing to those who make it and another to critics who retell it as a sales pitch.

If the figures were what kept speakers within the terms, their changing sides would weaken the terms. This does not happen, because the contest over the figures takes place within the terms. Gramsci (1971) located a social group's decisive success in posing the questions of the struggle on a "universal" (pp. 181–182) rather than a corporate plane, and the patterns we found give that formula a concrete form. Members of the elite quarrel over whether AGI is near and who should control it, and they aim insults drawn from faith and scripture at their critics, yet, Thiel apart, they offer no future that is not about AI.

These patterns also qualify earlier research on myth and technology. The double image of salvation and doom is old, yet here it appears not as two camps but as a common stock of figures that pass between sides. Crawford (2021) and O'Sullivan (2025) described shared terms across the discourse. Read speaker by speaker, however, the texts show them held unevenly, most consistently by the elite, less often by others, and refused by speakers who

point most often to those who own AI. The name "New Olympians" was meant to raise the question of whose terms these are, and the three points of the heuristic suggest an answer. The elite's posture shows in its forecasts, whose figures its members usually embrace, and its shared terms show in quarrels that stop at those terms. The name itself, an old figure for a new power, joins the vocabulary of faith and prophecy that critics aim at the elite's leaders, while the elite's own register is futurist. Whether its members would accept the name is not tested here. The New Olympians, the texts suggest, keep the terms most firmly but not alone.

5.2. How the Terms Come to Look Inevitable

Since the figures do not tie speakers to the terms, the sense that AI is inevitable is likely to come from somewhere else. Mosco (2004) located it in myth, in the prophecies that greeted each new medium. The findings suggest a qualification, since no single prophecy is accepted widely enough to make AI look inevitable, and outside the elite the futurist figures are rejected about as often as they are embraced. The inevitability may come from the terms instead. When every option that speakers put forward for the future is an option about AI, whether faster or safer, open or closed, American or Chinese, AI stops being one of the options and becomes the condition of all of them. Even a call to slow AI down can be framed this way, as when Amodei (2026) writes that pacing "within democracies will be limited by the lead that US companies have over authoritarian regimes." The result is Mosco's end of politics reached without being announced, and the work Barthes (1972, p. 142) assigns to myth, done less by any story than by the shape of the choice. Earlier studies found such inevitability in national AI strategies (Bareis & Katzenbach, 2022, pp. 875–876) and an imagined AI race among experts (Richter et al., 2025, p. 11). Here the inevitability holds across rival factions of the elite and across promise and threat, and even among the industry's most alarmed critics and those who debunk its forecasts.

The terms become visible when someone refuses them and names what they leave out. Those who refuse the terms point most often to the firms that own AI and the capital behind them, from Hao's shifting goalpost (Firth, 2025) and Morozov's (2025) oligarchs to the encyclical's monopolies (Leo XIV, 2026). They say in public what material accounts of AI argue in theory, namely that the elite's power rests on what it owns and controls. For Crawford (2021), the data practices behind AI rest on "a commercialized capture of what was previously part of the commons," and the AI gold rush "consists of enclosing different fields of human knowing, feeling, and action" (p. 120). Mejias and Couldry (2024) write of "data colonialism" (p. 4). O'Sullivan (2025) even calls the power structure "feudal," a word also used in the debate over whether such power is still capitalist or has become technofeudalism, in which rent replaces profit (Gane, 2025; Varoufakis, 2023). Our argument concerns the terms of public talk rather than the economic form of this power, so it does not depend on how that debate is resolved.

The AGI clause shows the same link between AI and ownership from inside the industry. When AGI had to be defined in a contract, it was reportedly defined as a profit threshold, and its arrival was written to change who holds commercial rights to the technology. The mythic language does not create the concentration of ownership that these critics describe, but by keeping the debate on AI itself, it may make that concentration harder to see.

5.3. A Weak Hegemony

Calling the pattern of shared terms a hegemony invites the objection noted earlier, that a dominant ideology may bind the dominant class more than the society it rules (Abercrombie et al., 2015, p. 3), and the findings support this objection. The terms bind the elite more tightly than they bind anyone else, while outside it about half are silent and the press frequently passes the figures on without endorsing or rejecting them. Silence could mean that speakers are indifferent, or that they take the terms for granted. Telling the two apart would require a study of audiences. The press also adopts the terms less fully here than the German media adopted their government's frames in Köstler and Ossewaarde's (2022) study, perhaps because our texts range more widely than one country's press. The press, then, does not act as the "deputies" that Gramsci (1971, p. 12) saw in intellectuals.

We suggest that the terms form a weak hegemony, one that binds those who lead it more tightly than those whom it would lead, and whose terms only a minority outside the elite openly shares. It is a process as Williams (1977, p. 112) describes it, renewed and never complete, and a vision that has spread beyond its promoters without being shown to be shared widely enough to count as an imaginary (Jasanoff, 2015, p. 4). Such a hegemony can also shift. Writing of a world-system in crisis, Wallerstein (1999) held that "little pushes will determine which way the process moves" (p. 33). Applied to a debate, by analogy only, the point suggests that a hold this weak leaves room for those who refuse the terms.

Two other readings could explain the same pattern. Members of the elite may simply believe that AI will decide the future, but sincere leaders can supply shared terms too. Alternatively, the pattern may be marketing, and Beck's (1992, p. 23) observation that risk is big business for its winners fits Merchant's (2026b) reading of Anthropic's warnings. In an earlier essay (Merchant, 2025), he treats the difference between sincere belief and marketing as beside the point. Of one forecaster of AI's end times, he writes that he "may be a true believer" or "a talented grifter," and "it doesn't really matter." Marketing does not explain why speakers without a commercial stake in AI also keep the same terms, although writing about AI may itself invite the terms, and terms that sell can also become common sense.

6. Conclusion

We began with the way many members of the technology elite speak of the future, as if it were already decided. In the texts we read, the mythic language that the elite shares with the industry's critics does not, by itself, set the terms of the debate on AI. Its figures pass between speakers, and although they usually cast AI as the source of promise or threat, they do not commit those who use them to the terms. But the terms persist, and the future is still presented as depending on AI. The "New Olympians" keep them most consistently, even while they quarrel. Other speakers keep them less often, and about half are silent on them. Those who refuse them point most often to ownership and capital, whereas Thiel, in his lectures, places the danger in the state. The patterns describe a weak hegemony that rests on terms more than on stories.

Disputing a forecast may therefore leave the terms in place, as it does for those who debunk the figures. A critique of the terms would also ask whether the choice a forecast offers is the only one. Only a minority of speakers raise this question. Yet the answer may determine how

much room public debate leaves for decisions that do not concern AI itself, such as who profits from it, how work is organized, and who answers for harm.

The study has limits, and the first matters most for what we found about refusal. Because the texts were chosen for their mythic language and only speakers who invoke a figure were coded, a speaker who refuses the terms without using any figure is not counted here. For the same reason, we can show that figures do not bind speakers to the terms, but not that they play no part in sustaining the terms. The corpus also includes many critical voices, which may have the opposite effect, making refusal outside the elite look more common than it is. Because all the texts are in English, the corpus says little about the debate in other languages. Each code describes a speaker only as our texts record them. Speakers who are silent here may have spoken elsewhere, as Andreessen did on AGI (Rogan, 2026), and their stances may have changed since, as Thiel's may have (Döpfner, 2026). Finally, we read only public talk, not the laws, budgets, and institutions through which a vision gains the force of an imaginary.

Williams (1989, pp. 8–9) set prediction apart from any pretense of being the future. The terms we found come close to that pretense, for they make the future depend on AI and so present AI as the future itself. This implicit claim, more than any single forecast, is the one most open to question, because it leaves unasked whose future AI would serve.

References

Abercrombie, N., Hill, S., & Turner, B. S. (2015). *The dominant ideology thesis* (Routledge Library Editions: Social Theory, Vol. 18). Routledge. https://doi.org/10.4324/9781315763804

Altman, S. (2025a, January 5). Reflections. https://blog.samaltman.com/reflections

Altman, S. (2025b, February 9). Three observations. https://blog.samaltman.com/three-observations

Altman, S. (2025c, June 10). The gentle singularity. https://blog.samaltman.com/the-gentle-singularity

Altman, S. (2026, April 10). [Untitled blog post on the attack on his home]. https://blog.samaltman.com/2279512

Amodei, D. (2026, September). We must pace the frontier. https://darioamodei.com/post/we-must-pace-the-frontier

Bareis, J., & Katzenbach, C. (2022). Talking AI into being: The narratives and imaginaries of national AI strategies and their performative politics. *Science, Technology, & Human Values, 47*(5), 855–881. https://doi.org/10.1177/01622439211030007

Barthes, R. (1972). *Mythologies* (A. Lavers, Trans.). The Noonday Press.

Beck, U. (1992). *Risk society: Towards a new modernity* (M. Ritter, Trans.). Sage.

Bhuiyan, J., Kerr, D., & Robins-Early, N. (2025, October 10). Inside tech billionaire Peter Thiel's off-the-record lectures about the antichrist. *The Guardian*. https://www.theguardian.com/us-news/2025/oct/10/peter-thiel-lectures-antichrist

Clark, J. (2026, May 26). Import AI 458: Reckoning with the future; and a singularity story. *Import AI*. https://jack-clark.net/2026/05/26/import-ai-458-reckoning-with-the-future-and-a-singularity-story/

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press. https://doi.org/10.12987/9780300252392

de Sailly, L. (2025, January 21). The New Olympians: Tech billionaires and the myth of progress. *Hey, Slick!* https://heyslick.substack.com/p/the-new-olympians-tech-billionaires

Döpfner, M. [MDMEETS - Essential Talks with Mathias Döpfner]. (2026, September 24). *Peter Thiel: The AI crisis, Europe's decline & the battle for America* [Video]. YouTube. https://www.youtube.com/watch?v=B7yl7fEHeKM

Farrow, R., & Marantz, A. (2026, April 6). Sam Altman may control our future—Can he be trusted? *The New Yorker*. https://www.newyorker.com/magazine/2026/04/13/sam-altman-may-control-our-future-can-he-be-trusted

Firth, N. (Host). (2025, July 9). Inside OpenAI's empire: A conversation with Karen Hao. *MIT Technology Review*. https://www.technologyreview.com/2025/07/09/1119784/inside-openais-empire-a-conversation-with-karen-hao/

Gane, N. (2025). Capitalism is capitalism, not technofeudalism. *Journal of Classical Sociology, 25*(4), 353–368. https://doi.org/10.1177/1468795X241269293

Geraci, R. M. (2010). *Apocalyptic AI: Visions of heaven in robotics, artificial intelligence, and virtual reality*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195393026.001.0001

Geuter, J. (2026, April 21). AI as a fascist artifact. *Smashing Frames*. https://tante.cc/2026/04/21/ai-as-a-fascist-artifact/

Gramsci, A. (1971). *Selections from the prison notebooks* (Q. Hoare & G. Nowell Smith, Eds. & Trans.). International Publishers.

Hall, S. (1979, January). The great moving right show. *Marxism Today, 23*(1), 14–20.

Hall, S. (2005). Encoding/decoding. In S. Hall, D. Hobson, A. Lowe, & P. Willis (Eds.), *Culture, media, language: Working papers in cultural studies, 1972–79* (pp. 117–127). Routledge. https://doi.org/10.4324/9780203381182-13

Hsieh, H.-F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative Health Research, 15*(9), 1277–1288. https://doi.org/10.1177/1049732305276687

Jaffe, S. (2025, May 29). What's really artificial about AI is AI itself. *In These Times*. https://inthesetimes.com/article/big-tech-the-ai-con-book-silicon-valley

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. https://doi.org/10.7208/chicago/9780226276663.003.0001

Johansen, T. K. (1999). Myth and logos in Aristotle. In R. Buxton (Ed.), *From myth to reason? Studies in the development of Greek thought* (pp. 279–291). Oxford University Press. https://doi.org/10.1093/oso/9780198152347.003.0016

Kelly, K. (2025, August 5). Artificial intelligences, so far. *The Technium*. https://kk.org/thetechnium/artificial-intelligences-so-far/

Klein, E. (2025, October 15). How afraid of the A.I. apocalypse should we be? *The New York Times*. https://www.nytimes.com/2025/10/15/opinion/ezra-klein-podcast-eliezer-yudkowsky.html

Köstler, L., & Ossewaarde, R. (2022). The making of AI society: AI futures frames in German political and media discourses. *AI & Society, 37*(1), 249–263. https://doi.org/10.1007/s00146-021-01161-9

Leo XIV. (2026, May 15). *Magnifica humanitas: On safeguarding the human person in the time of artificial intelligence* [Encyclical letter]. The Holy See. https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html

Lovely, G. (2025, December 15). The AI doomers feel undeterred. *MIT Technology Review*. https://www.technologyreview.com/2025/12/15/1129171/the-ai-doomers-feel-undeterred/

Malik, O. (2026, June 7). The myth, the Mythos and the man. *On my Om*. https://om.co/2026/06/07/the-myth-the-mythos-and-the-man/

Marantz, A. (2024, March 11). Among the A.I. doomsayers. *The New Yorker*. https://www.newyorker.com/magazine/2024/03/18/among-the-ai-doomsayers

Mejias, U. A., & Couldry, N. (2024). *Data grab: The new colonialism of Big Tech and how to fight back*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226832319.001.0001

Merchant, B. (2025, October 10). AI profiteering is now indistinguishable from trolling. *Blood in the Machine*. https://www.bloodinthemachine.com/p/ai-profiteering-is-now-indistinguishable

Merchant, B. (2026a, April 14). Why the AI backlash has turned violent. *Blood in the Machine*. https://www.bloodinthemachine.com/p/why-the-ai-backlash-has-turned-violent

Merchant, B. (2026b, June 12). How to use AI doom marketing to dupe the media and rake in billions in 10 easy steps. *Blood in the Machine*. https://www.bloodinthemachine.com/p/how-to-use-ai-doom-marketing-to-dupe

Merton, R. K. (2016). The self-fulfilling prophecy. *The Antioch Review, 74*(3), 504–521. https://doi.org/10.7723/antiochreview.74.3.0504

Metz, C. (2026, May 26). At the epicenter of A.I., Pope Leo's warnings are dismissed. *The New York Times*. https://www.nytimes.com/2026/05/26/technology/pope-leo-ai-religion.html

Morozov, E. (2025, April 3). The new legislators of Silicon Valley. *The Ideas Letter*. https://www.theideasletter.org/essay/silicon-valleys-new-legislators/

Mosco, V. (2004). *The digital sublime: Myth, power, and cyberspace*. MIT Press. https://doi.org/10.7551/mitpress/2433.001.0001

Noble, D. F. (1997). *The religion of technology: The divinity of man and the spirit of invention*. Alfred A. Knopf.

O'Sullivan, J. (2025, December 9). The politics of superintelligence. *Noema*. https://www.noemamag.com/the-politics-of-superintelligence/

Perrigo, B. (2025, April 15). Demis Hassabis is preparing for AI's endgame. *Time*. https://time.com/7277608/demis-hassabis-interview-time100-2025/

Richter, V., Katzenbach, C., & Zeng, J. (2025). Negotiating AI(s) futures: Competing imaginaries of AI by stakeholders in the U.S., China, and Germany. *Journal of Science Communication, 24*(2), Article A08. https://doi.org/10.22323/2.24020208

Rogan, J. [PowerfulJRE]. (2026, May 19). *Joe Rogan Experience #2501 - Marc Andreessen* [Video]. YouTube. https://www.youtube.com/watch?v=PHQvb10vKyk

Roose, K., & Newton, C. (2025, February 28). Anthropic's C.E.O., Dario Amodei, on surviving the A.I. endgame. *The New York Times*. https://www.nytimes.com/2025/02/28/podcasts/hardfork-anthropic-dario-amodei.html

Sacks, D. [@DavidSacks]. (2025, August 9). *A BEST CASE SCENARIO FOR AI? The Doomer narratives were wrong. Predicated on a "rapid take-off" to AGI, they predicted* [Post]. X. https://x.com/DavidSacks/status/1954244614304739360

Schmitt, C. (2005). *Political theology: Four chapters on the concept of sovereignty* (G. Schwab, Trans.). University of Chicago Press. https://doi.org/10.7208/chicago/9780226738901.001.0001

Streitfeld, D. (2026, May 14). Elon Musk and friends in the den of Oakland's literary lion. *The New York Times*. https://www.nytimes.com/2026/05/14/technology/ishmael-reed-elon-musk-play.html

Varoufakis, Y. (2023). *Technofeudalism: What killed capitalism*. The Bodley Head.

Wallerstein, I. (1999). *The end of the world as we know it: Social science for the twenty-first century*. University of Minnesota Press.

Williams, R. (1977). *Marxism and literature*. Oxford University Press.

Williams, R. (1989). Culture is ordinary. In R. Gable (Ed.), *Resources of hope: Culture, democracy, socialism* (pp. 3–18). Verso.

Willison, S. (2026, April 27). Tracking the history of the now-deceased OpenAI Microsoft AGI clause. *Simon Willison's Weblog*. https://simonwillison.net/2026/Apr/27/now-deceased-agi-clause/