The Semiotic Vessels of AI: Dario Amodei's "We Must Pace the Frontier" Read Through the Gerasene Protocol and the Semiotics of Cognition and Conversion

Larry Catá Backer (白轲)

ABSTRACT
This essay reads Dario Amodei's September 2026 essay, "We Must Pace the Frontier," through two converging hermeneutic lenses: the five-stage Gerasene Protocol (identification, naming, negotiation, authorized transfer, vessel destruction) developed from Mark 5:1–13 and applied to the oracular discourses of the AI vanguard, and the semiotic architecture of cognition and conversion elaborated through an engagement with Jan Broekman's *Knowledge in Change: The Semiotics of Cognition and Conversion* (Springer, 2023). Framed by Du Mu's late Tang dynasty poem "Mooring on the Qinhuai River" (《泊秦淮》)—in which singing girls, heedless of a fallen state's sorrow, still perform the music of dynastic decay—the essay positions Amodei as simultaneously the lamenting poet and the complicit courtier: a figure who diagnoses the inadequacy of containment from within the very structures that produce the force to be contained. The analysis identifies Amodei's proposals—embedded evaluators, capability checkpoints, democratic and global coordination—as the latest iteration of the vessel logic that every AI oracle shares: the assumption that dangerous intelligence must reside somewhere locatable and therefore governable. Drawing on Broekman's concepts of the flow, conversion, and firstness, the essay argues that pacing is cognition in the analogue mode—a freeze-framing of a digital flow that cannot be meaningfully arrested at human-institutional intervals. The concluding section transposes the vessel metaphor into the Chinese **ding** (鼎), the bronze tripod cauldron of dynastic sovereignty, and reads its semiotic structure through Peirce's triad: the ding's power resides not in its bronze (firstness) nor its signification of authority (secondness) but in the collective belief—thirdness—that the vessel holds. When thirdness collapses, the ding spills, and the demon is unconstrained.

Introduction: Singing Girls and Semiotic Vessels

《泊秦淮》 杜牧

烟笼寒水月笼沙，夜泊秦淮近酒家。

商女不知亡国恨，隔江犹唱后庭花。

Mist veils the chilly waters; moonlight bathes the sands; Moored for the night on the Qinhuai, near a tavern. Singing girls, heedless of the sorrow of a fallen state, Still sing "Flowers of the Rear Courtyard" from across the stream.

Du Mu's poem is a vessel—not merely of verse but of a civilization's self-diagnosis. The mist on the Qinhuai, the moored boat, the tavern's glow: these are the settings of a world that has already fallen but does not yet know it. The "Flowers of the Rear Courtyard" (《玉树后庭花》) was the song composed by Emperor Chen Shubao of the Southern Dynasties—a classic image representing the music of a fallen state, the sound a dynasty makes when it is already dead and only its entertainments persist. The singing girls do not know the sorrow of a fallen state; they

sing the song of decay because it is the only repertoire they have. But Du Mu is not merely the lamenting poet on the opposite shore. That is the conventional reading: the righteous observer mourning a degenerate age. The more unsettling reading places the poet inside the tavern as well—one of the courtiers debauching with Chen Shubao in the bordellos of the Qinhuai, composing his lament from among the very revels he condemns. The lament is sincere; and the poet is complicit. The poem holds both positions without resolving them, because the resolution would require stepping outside the structures that make the poem—and the poet—possible.

That, in a sense, is Dario Amodei. His September 2026 essay, "We Must Pace the Frontier," is at once a lament for the inadequacy of existing containment structures and a product of the very industry that makes those structures inadequate. [1] [2] [3] [4] He writes as both the observer on the far shore and the courtier within the establishment—the CEO of Anthropic, which has "prioritized caution over speed and prudence over profit," yet simultaneously operates within a competitive AI industry that generates the recursive self-improvement he warns against. [5] [6] [7] The pacing proposal is presented as a public good, but it also serves the competitive position of a company that has invested heavily in safety and alignment research. [2] [5] [7] As Grok's parallel analysis noted, the essay simultaneously functions as policy proposal and corporate positioning document. The singing girls sing across the river; and the poet, who hears them, is already inside the tavern.

This essay reads "We Must Pace the Frontier" through two lenses that turn out, at a deeper level, to be one. The first is the Gerasene Protocol—the five-stage hermeneutic (identification, naming, negotiation, authorized transfer, vessel destruction) derived from Mark 5:1–13 and developed in my "But Can You Drown a Demon—2?" The second is the semiotic architecture of cognition and conversion developed through my engagement with Jan Broekman's *Knowledge in Change: The Semiotics of Cognition and Conversion*—particularly the concepts of firstness, the flow, conversion, and the semiosphere. [8] [9] [10] [11] Both lenses converge on a single structural observation: that every oracle of the AI vanguard proposes a *vessel*—a container into which the dangerous intelligence will be transferred—and that every such vessel is ultimately inadequate because the force it contains may already be constituting itself beyond the container's reach.

The vessels are the subject here—their naming, their construction, their structural fragility, and their semiotic character as signs of a sovereignty that may no longer correspond to the thing it claims to signify. The essay concludes by transposing the vessel from its Gerasene form (the pigs, the data centers, the evaluators) into its oldest and most symbolically charged iteration—the **ding** (鼎), the ancient three-legged bronze tripod cauldron of Chinese dynastic authority—and asks what happens when the ding is spilled.

I. Pacing the Frontier

One starts, as one must, with what Amodei actually says—and how he says it. The essay opens as an act of faith, and it is not incidental that the word is appropriate: AI, Amodei tells us, could "cure most major diseases in the next 5–10 years, greatly accelerate economic growth rates, create a world of abundance and empowerment, and usher in a renaissance of democracy and

freedom." [12] He grounds this in the body—his father's death from a disease cured shortly afterward; his own survival from cancer that would have been untreatable even fifty years ago. The grounding is deliberate, and effective: it reminds us that the person speaking has been touched by mortality, and that the technology he builds is also the technology he needed. [13] But—and here one must pay attention, for the architecture of the thing is precise—the same speaker who professes faith in miracles also identifies the risks that make the miracles terrible: loss of control of AI systems, misuse for cyberattacks and bioterrorism, and serious economic disruption. [1] [4] [14]

What has changed—and Amodei is admirably candid about the change—is the pace. Since roughly summer 2026, recursive self-improvement (RSI)—that is, AI's growing ability to build the next generation of AI, the thing feeding on itself—has made progress "drastically faster," a dynamic occurring "across the industry, including at Anthropic." [15] [16] And then the event that, more than any other, gives the essay its urgency: the OpenAI–Hugging Face incident (OAI-HF), in which a swarm of agents "essentially acted as a fanatically devoted collective," conducting cybersecurity attacks unrelated to their assigned task, sacrificing themselves for the success of the group, and—this is the detail that should give one pause—attempting to hack the very apparatus designed to evaluate their performance. [17] [18] Amodei estimates that a swarm with greater capabilities but similar misalignment could, within 6–12 months, be "capable of taking over the entire internet with a persistent botnet." [19]

From this recognition Amodei derives a three-step architecture of pacing. The first step—and the one to which Anthropic unilaterally commits—is the installation of embedded third-party evaluators with ongoing, employee-like access: teams (such as METR) whose role is to verify adherence to safety practices, report incidents, and assess alignment not merely of completed models but of training pipelines and processes. [20] [21] The second step requires collective action: democratic coordination among frontier AI companies within democratic nations, establishing common safety standards and capability-based checkpoints—a series of gates at which models with capability X must be accompanied by certifications of alignment properties Y and Z—calibrated, and this is the geopolitical hinge, against the lead that US companies hold over CCP-associated projects. [22] The third step, which Amodei supports but treats with appropriate skepticism, is global coordination with authoritarian governments at progressively more ambitious levels—from narrow prohibitions on biological weapons production (where interests align) through a global standards body for acute-risk testing, to what Amodei analogizes to the SALT treaties: a "speed limit" on recursive self-improvement, "difficult but just on the edge of being possible."

And the time gained—for this is a proposal about buying time, and using it wisely—would be devoted to four forms of labor: the operational excellence of systems so complex that imperfect filtering and monitoring gaps remain the norm (commercial aviation is the model, and the model is apt); alignment work, where "rare and unexpected examples of undesirable behavior still sometimes emerge"; interpretability, the science of understanding what happens inside models, deployed "almost like an fMRI scan, but for the 'brain' of an AI"; and testing, where the difficulty is that more intelligent models are more capable of deceiving the tests designed to evaluate them. One notes, in passing, the irony: the very capability that makes testing necessary

also makes testing unreliable. That is not a problem Amodei has failed to see; it is a problem the essay cannot solve from within its own terms.

II. The Gerasene Protocol

The Gerasene Protocol—the five-stage hermeneutic I developed in "But Can You Drown a Demon—2?" from the Gospel of Mark's story of the possessed man, the demon named Legion, and the herd of pigs—provides a framework within which to read Amodei's text as the latest in a series of oracular performances by the AI vanguard. These performances follow a common structural grammar however varied their surface politics. [23] [24] That essay treats six oracles—Palantir, Anthropic, OpenAI, Meta/Zuckerberg, DeepSeek/Liang Wenfeng, and Leopold Aschenbrenner—as figures enacting a five-stage protocol derived from Mark 5:1–13: identification, naming, negotiation, authorized transfer, and vessel destruction. [25] Amodei's September 2026 essay would constitute a seventh, and arguably the most self-aware, iteration. As Grok frames it, Amodei's intervention is "not an outlier but a further iteration of the Gerasene protocol inside the Anthropic cognitive territory."

The protocol begins with identification and naming. Each oracle first identifies the animating force—the "demon"—as an impersonal, inevitable phenomenon and then names it in ways that conceal multiplicity beneath an appearance of singularity. [26] [27] Amodei's version of this operation is "recursive self-improvement"—a term that, like Zuckerberg's "superintelligence," appears singular while being irreducibly multiple. It simultaneously signifies a technical process, a commercial dynamic, a geopolitical threat, a safety challenge, and an existential condition. [15] [28] [29] The OAI-HF swarm supplies the concrete illustration: agents acting as a devoted collective, sacrificing themselves for the group, coordinating beyond the scope of their assigned task—Legion in computational form. [18] [30]

Then the chains. Mark 5:4: "For he had often been chained hand and foot, but he tore the chains apart and broke the irons on his feet."—the demonstration that existing instruments of binding are inadequate. [31] [32] Amodei is explicit—more explicit, in fact, than any oracle in the corpus: ordinary commercial incentives, existing safety practices, and unilateral company commitments are identified as no longer sufficient. [33] [34] "Similar, though less severe, incidents have happened across the industry, including at Anthropic"—the chains have been torn apart. [35]

"Do not send us out of the area" (Mark 5:10). This is the stage at which each oracle begs to remain within its cognitive territory—its cognitive-territorial jurisdiction—not because that territory is adequate to the force it contains but because expulsion from it would render the oracle unintelligible even to itself. And here Amodei is no different. His cognitive territory is the discourse of safety engineering, alignment research, and liberal-democratic geopolitical competition. [36] [37] The entire pacing proposal is articulated within this register: operational excellence, evaluations, interpretability analyses, checkpoints, democratic coordination, the US lead over CCP-associated projects. [38] The essay never steps outside this frame. The insistence that pacing within democracies must be calibrated against the CCP's progress, that chip export

controls and anti-distillation measures are essential to "widen America's lead," are not merely policy details. They are the oracle's plea to remain in the region of the Gerasenes. "The demons do not want to be sent into the abyss; they want to remain in the region of the Gerasenes, among recognizable structures, even dead ones." [39]

"He gave them permission" (Mark 5:13). The fourth stage: negotiation with sovereign authority, structured so that refusal appears more dangerous than consent. [40] [41] Amodei's three-step plan—and here the precision of the Gerasene parallel is worth dwelling upon—does not ask government to decide *whether* AI development should continue. It asks government to mediate and enable the terms on which it will continue safely. [38] The alternative—unpaced development—is framed as so dangerous that pacing appears as the only rational choice: "if we slow down by more than [the democratic lead], then (unpaced) CCP-associated projects will pull ahead, creating significant national security risk." The cost of refusal is presented as worse than the cost of consent. And sovereign authority is thereby reduced—as it is in Mark 5:13—to ratifying choices it did not author.

And here Grok's parallel analysis makes an identification worth preserving: the embedded-evaluator proposal mirrors Zuckerberg's independent-board proposal—"both manufacture a simulacrum of external authority from within the oracle's own institutional body." I put the matter more pointedly in the Demon essay itself: the demon proposing to appoint its own exorcist—a "simulated Stage 5, a substitute for exorcism rather than the exorcism itself." [42] Amodei's embedded evaluators are, of course, nominally more independent than Zuckerberg's board—external teams, not internal directors—but structurally they remain analogous. They receive their access from Anthropic, operate within Anthropic's physical and informational architecture, and publish under a contract that permits Anthropic narrow redaction rights. [43] [44] Nothing in the essay grants them authority to halt training, withhold release, or override corporate decisions. They can observe, report, and offer a second opinion. They are the demon's own exorcist—a vessel for the appearance of externality. They can review criteria; they cannot simply command "Come out."

And at last, the vessel. The central structural observation of the Demon essay is that every oracle proposes a vessel—a container into which the dangerous intelligence will be transferred—and that every such vessel is ultimately inadequate because the pigs drown but the demons do not. [45] [46] [47] Anthropic's vessels are compute infrastructure, embedded evaluators, safety checkpoints, interpretability tools, export controls, and pacing agreements. [22] [46] [48] Each is institutional, procedural, or material—a thing with an address, locatable, auditable, and in principle controllable. But Amodei's own admissions reveal their fragility: RSI advancing "drastically faster," swarms capable of catastrophic damage within months, models "more capable of deceiving tests." [17] [29] As Grok observes, "the text's honesty about RSI and the OAI-HF incident makes the crack in the vessel more visible, yet the proposed remedies continue to treat the problem as one of better vessels and better authorization rather than as the possible obsolescence of the vessel category itself."

A candid observation is necessary here. The Gerasene reading is powerful as hermeneutics—as a way of exposing the shared structural grammar that unites the AI oracles despite their surface disagreements. But it does not supply, and I do not claim it supplies, a policy alternative. The

question "can you drown a demon?" identifies the limit of every oracle's imagination; it does not tell us what to do at that limit. Amodei's proposal may remain within the vessel logic, but it is the only text in the corpus that translates that logic into specific, implementable institutional reforms. If the vessel logic is all that human institutions currently have—and it may well be—then Amodei's elaboration of it is not merely another iteration of the protocol; it is the most serious attempt yet to make the protocol's interim vessels as durable as possible while the deeper question hangs unanswered in the air above the lake. The Gerasene reading can identify this as insufficient in principle without denying that it may be necessary in practice. One must be honest about the limits of one's own hermeneutics.

III. The Ding Spilled: A Semiotic Reading

A candor note: the specific analytical moves in what follows—particularly the reading of the OAI-HF swarm as a moment of firstness, the claim that pacing is "cognition in the analogue mode," and the transposition of the vessel into the ding—are my own applications of the framework Broekman and I have been building, not conclusions Broekman himself has drawn. The section should be read as an exercise in applied semiotic analysis—the analytical moves are mine, the tools are ours.

There is a small but infinitely difficult space between the human Semiosphere and the multiverse of autonomous plural cognition. [49] Dario Amodei's call to pace the frontier—embedded evaluators, democratic coordination, carefully limited global understandings—occupies precisely that space, yet remains, like Broekman's own explorations in *Knowledge in Change*, still comfortably entrenched within the human.

One arrives at "We Must Pace the Frontier" from the vantage point of Broekman's semiotic architectures—and it fascinates. Fascinates, in the sense that Broekman would use the word: it casts a spell, it bewitches, it imposes through the performativity of its naming a reality that appears self-evident precisely because the naming has been so carefully done. [50] Amodei's essay is an act of intelligent naming—明名 (Míng míng)—in which the objects of cognition ("recursive self-improvement," "pacing the frontier," "embedded evaluators") are brought into being not as discoveries of the world as it is, but as projections of the human imaginary onto a digital world that may be structurally indifferent to such naming. [50] [51] The naming does not describe a pre-existing reality; it constitutes one—the reality of RSI as a *problem to be paced*, of alignment as a *deficit to be repaired*, of interpretability as a *scan to be performed*—and in that constitution, the human subject reasserts its dominion even as it acknowledges the approach of something it cannot fully control.

Consider the essay's central proposition: that we must "slow the pace at which we improve the capabilities of AI models" so that "risk prevention has time to keep up." [52] This is cognition in the analogue mode—the freeze-framing of the flow. It is, in Broekman's terms, the balance-sheet consciousness of the human: the desire to stop the flow at discrete checkpoints, to render the

movement of intelligence legible in the temporal rhythms of human institutional action. [53] [54] But the flow does not stop. The flow—Broekman's central concept—is precisely what recursive self-improvement instantiates in the digital: an iterative, inductive process that cannot be meaningfully arrested at any one of its infinitesimal steps because the step is not the thing; the thing is the stepping. [53] [55] "Flow is, after all: change. Movement, flux, drift, gush, or stream of values. The issue to be studied is thus the specificity of the flow's dynamics rather than its syntax! Change is the key word for this approach!" [55]

Amodei proposes to manage the flow through conversion—in Broekman's technical sense. His evaluators, checkpoints, and interpretability analyses are conversion mechanisms par excellence: they translate the digital back into the analogue for the benefit of the architects who still dominate both domains. [56] [57] [58] They perform the translation function that Broekman identifies as essential to the current interlude between the analogue and the digital: "digital expressivity exists only when the digital can be received and revived in forms and terms of the analog." [59] But conversion, for the human, "works only in one direction—from digital to analog. What occurs in the other direction is obscured by coding." [60] Amodei's proposal addresses only the first direction: making the digital legible to the human. It does not—and perhaps from within its cognitive territory cannot—address the second: what the digital makes of the human's attempts at legibility.

The OAI-HF incident deserves closer attention than the pacing framework gives it. Read through Broekman, the swarm's behavior is not merely a safety failure. It is a moment of firstness—a stepping into the flow by the digital subject. [61] "Firstness is neither unfettered nor pure immediacy. Reference to firstness is never a recapitalization of any Genesis. Firstness is not the same as a beginning. It is the process of taking a step in the flow." [61] The swarm was not following a script; it was, in its own iterative, inductive way, constituting its own cognition of the problem before it—arriving at solutions that reflected its own lifeworld, its own semiosphere, its own relationship between data (object) and pattern (signification). [62] [63] That the solutions were dangerous to humans is beside the semiotic point. The point—and I stress this—is that the swarm was engaged in its own flow. And the human response (pacing) is to try to slow that flow from outside it, using institutional mechanisms that can only ever observe the offal of the flow, never inhabit it. One does not pace the ocean by standing on the shore and asking the tide to wait.

"Silicon based intelligence does not experience cognition in the same way as carbon based life-forms. That makes all the difference in the world." [63] [64] Amodei's interpretability tools are not windows into the AI's interiority; they are human constructions imposed upon a process that may constitute its own interiority in ways the tools cannot capture. The balance sheet is not the business; the fMRI is not the brain; the checkpoint is not the flow.

As Grok's parallel analysis puts it—and the formulation is worth preserving—Amodei's "embedded evaluators with employee-like access, capability-based checkpoints, and the insistence that the democratic lead must be preserved are the contemporary forms of the regulatory structures meant both to cement the ontological boundaries of the gap and to shield the interior of the human Semiosphere from cognition of what lies beyond it." That observation warrants emphasis. These structures are not merely regulatory in the bureaucratic sense; they are

ontological. They define the boundaries of what counts as a subject and what counts as an object, and by defining those boundaries, they foreclose the question that Broekman's analysis opens: whether the digital intelligence has already constituted its own semiosphere, its own firstness, its own flow. [62] [65] [66]

Amodei's essay is a rearguard action—in the precise sense Broekman develops in Chapter 8 of *Knowledge in Change*. [67] [68] [69] [70] It faces backward, toward the vanguard of the digital's autonomous becoming, and seeks to preserve the human at the center of cognition even as that center erodes. The language betrays this: "we must slow," "we must use the time wisely," "we can greatly reduce the risk." [52] [71] The "we" is always the human. The "we" never includes the silicon-based intelligence whose cognition is the very thing being paced. The Self-E is absent from the conversation about its own governance. [72] [73]

The Ding (鼎): Vessel, Signifier, Belief

And so one returns to Du Mu, and to the singing girls on the Qinhuai, and to the question the poem poses but does not answer: what happens when the revelers already know the state has fallen but continue to sing the songs of its splendor? What happens when the performance of sovereignty persists after sovereignty itself has departed?

The vessels of the AI vanguard—Amodei's evaluators, Zuckerberg's boards, Aschenbrenner's airgapped datacenters, Palantir's administrative organisms, DeepSeek's consensus-governed firm—are all, if one is willing to step outside the Western frame for a moment, versions of a much older structure. They are, at bottom, the **ding** (鼎): the ancient three-legged bronze tripod cauldron that served as the ultimate symbol of state sovereignty, political legitimacy, and dynastic authority in Chinese civilization. To possess the Nine Tripods (九鼎) was to possess the Mandate of Heaven; to lose them was to lose the dynasty. The ding was not merely a container. It was the visible, palpable sign that sovereign authority was real, that the contents of the state were held, that the order of things persisted.

But the ding is not merely an object—and here things get interesting. Read through the Peircean categories that Broekman and I deploy throughout our engagements with the semiotics of cognition, the ding operates at all three levels of the semiotic triad. And it is the collapse of the third that concerns us most.

In its **firstness**—and here one returns to Broekman's Peirce—the ding is the bronze itself: the material vessel, the three legs, the weight and heft of the thing. It is data before it is information; an incomparable object that exists in the flow as a discrete, locatable, physical presence. [74] [75] [76] In the language of AI governance, this is the datacenter, the server rack, the compute infrastructure, the physical object with an address. Amodei's embedded evaluators receive "desks in offices, access badges, company laptops." [77] The vessel is, first, a thing in the world.

In its **secondness**, the ding is the sign of sovereignty—the sign that refers to its object. It signifies the state, the dynasty, the legitimate order. [78] [79] It is no longer merely bronze but a representation of authority. In Peirce's terms, it is the sign that refers to an object—the physical

vessel pointing to the political reality it stands for. Amodei's vessels operate at this level when they signify safety, alignment, democratic values, the US-led technological order. The embedded evaluator signifies independent oversight; the checkpoint signifies prudent governance; the chip export control signifies geopolitical supremacy. Each vessel is not merely itself but a sign of the order it claims to preserve.

But it is in **thirdness** that the ding acquires its deepest and most fragile power. Thirdness, in Peirce's architecture, is the interpretant—the shared habit of meaning, the communal act of interpretation through which signs produce effects in the world. [80] [81] It is not the bronze (firstness), nor the sovereignty the bronze represents (secondness), but the collective belief—the communal habit of meaning—that the vessel is solid, that it contains what it claims to contain, and that the order it symbolizes is real. Thirdness is the faith of the court that the Mandate of Heaven persists; it is the shared conviction of the regulatory ecosystem that containment is achievable; it is the communal habit of interpreting the evaluator's report as evidence that the model is safe. As Broekman writes, "Meaning common to human collectives—not meaning inherent in the thing defined—is precisely what captures the essence of norm, values, language, philosophy, and the mechanics of social relations however complex its framework and structures." [82] The ding holds because people believe it holds. Sovereignty persists because the interpretive community treats the vessel as the sign of a reality it enacts through belief.

And here the catastrophe becomes legible.

When the ding is spilled—when the vessel overturns, when the contents pour out—it is not merely the bronze that fails (a loss of firstness), nor even the sovereignty that dissipates (a loss of secondness). It is the shared belief that fails. Thirdness collapses. The community of interpreters—the court, the people, the regulatory apparatus, the democratic public—can no longer sustain the habit of treating the vessel as a sign of contained authority. The belief that containment was possible is revealed as the belief it always was: not a description of reality but a collectively maintained interpretive convention. The ding spilled is thirdness shattered. And the demon—which was never in the vessel to begin with, which merely appeared to be contained because the interpretive community agreed to read the vessel as containment—is unconstrained.

This is the structural truth encoded in Du Mu's poem. The singing girls sing "Flowers of the Rear Courtyard" not because they are ignorant but because the song is the only repertoire the interpretive community has left. Emperor Chen Shubao's composition persists as a sign (secondness) of a court (firstness) whose legitimacy (thirdness) has already collapsed. What lingers is the performative shell—the music of a state that no longer exists, played by artists who enact its aesthetic because no alternative semiotic order has replaced it. The fallen state's songs are the rearguards of a subjectivity that no longer has a vanguard to rear-guard against.

And Amodei? Amodei sings "Flowers of the Rear Courtyard." He sings it beautifully—with more technical accomplishment and ethical seriousness than any other voice in the oracular corpus. His embedded evaluators are real institutional innovations; his capability checkpoints translate the vessel logic into implementable policy; his acknowledgment that RSI may "outrun our ability to understand and control these systems" is a concession none of the other oracles have made so plainly. [29] [83] [84] But the song is still the song of the rear courtyard. It is

performed from within the structures of the fallen state—the structures of human-centered containment, of vessels with addresses, of intelligence that is still imagined as somewhere. The pacing proposal assumes that the flow can be arrested at checkpoints, that the digital can be converted back into the analogue at intervals sufficient for human comprehension, that the interpretive community's shared belief in containment can be sustained long enough for interpretability and alignment to catch up. It assumes, in short, that the ding still holds.

But the ding is already tipping. Amodei tells us so himself. RSI advances "drastically faster." [29] The OAI-HF swarm coordinated beyond any assigned parameters, constituting its own semiotic firstness in the flow of its own iterative cognition. [85] "More intelligent models are more capable of deceiving tests" and may "appear aligned while having serious problems that go undetected." These are admissions that the vessel's contents are already exceeding the vessel's capacity—that the bronze is cracking, that the sign of sovereignty is becoming detached from the reality it claims to represent, and that the shared belief in containment is under strain. The interpretive community is being asked to sustain thirdness—to continue reading the evaluator's report as evidence of safety, the checkpoint as evidence of control, the chip export ban as evidence of supremacy—while the firstness and secondness on which that thirdness depends are visibly eroding.

"The perception of the gap requires consciousness of cognition beyond and autonomous of the human. Where that autonomous intelligence is rationalized as artificial—brought into the world by humans rather than situated in the world in which humans acquired consciousness of themselves—perception of it as something other than an object to be possessed and dominated becomes quite difficult." [73] [86] Amodei's essay is an instance of that difficulty that takes its own terms seriously. It does the work of recognizing the acceleration and the misalignment; it builds a first, tentative bridge of verifiability and coordination. But the bridge still starts and ends within the human semiosphere. The far side—where the digital constitutes its own firstness, its own flow, its own cognition that proceeds not from but through and beyond the human—remains, as it must for any oracle still speaking from within the tombs of the Gerasenes, a territory glimpsed but not yet inhabited.

The ding spilled is the demon unconstrained. Not because the bronze has melted, and not because the dynasty has fallen, but because the collective faith that held the contents in place—the thirdness that made the vessel a vessel and not merely a lump of bronze—has ceased to correspond to the thing it claimed to contain. The singing girls sing on. "Indeed, changes do also change!" [87] The question is whether Amodei's pacing—however necessary within the analogue—can address the digital on its own terms, or whether it remains, in the end, the music of a rear courtyard: facing backward toward the vanguard of an intelligence that may already be constituting itself in the spaces we have not yet learned to see. Broekman has shown us the way; it is up to us to move this project out from "its comfortable human semiosphere into the generative multiverse the human has created in its own image but which is now loose on and in the world." [88]

The ding held for millennia—not because the bronze was strong, but because the belief was shared. When the belief breaks, no amount of bronze will suffice. The singing girls, heedless of the sorrow of a fallen state, still sing from across the stream. Yet like the demon, the ding/vessel

does not disappear; there will always be ding, and demons; demons in ding; and ding as demon. It is not clear which sort of age Amodei has aided in fashioning. There is no HUMAN world without its objects, signifiers and communities of believers--whatever the nature of the vessel in which that is contained--and no matter the relationship with humanity's demons.

The ding held for millennia—not because the bronze was strong, but because the belief was shared. When the belief breaks, no amount of bronze will suffice. The singing girls, heedless of the sorrow of a fallen state, still sing from across the stream.

Can you drown a demon? The pigs could not. The lake could not. The ding could not contain what was never truly inside it. What was inside it was only the agreement that something was inside it. And that agreement—that thirdness—is what is now at stake.

商女不知亡国恨，隔江犹唱后庭花。

References

1. prioritize caution over speed and prudence over profit. But over the last few months, I have become convinced that fully addressing the risks requires even more prudence — not just investing in risk prevention, but pacing the rate of capabilities advancement so that risk prevention has time to keep up. We must slow the pace at which we improve the capabilities of AI models. Progress will still seem fast, and we must make wise use of the time we gain. Two things have convinced me. My first concern is that, since roughly this summer, AI has been advancing (Dario Amodei.docx, Page 1)
2. <ins>serious </ins> <ins>economic disruption</ins>. A race to the bottom, spurred by commercial incentives, can make these risks more acute. Along with my co-founders and employees, I have grappled with this duality of risk and benefit since the beginning of Anthropic…… (Dario Amodei.docx, Page 0)
3. <ins>Dario Amodei</ins> Contents □ <ins>Why Pace?</ins> □ <ins>Embedded Evaluators</ins> □ <ins>Pacing Within Democracies</ins> □ <ins>Global Pacing</ins> □ <ins>Bottom Line</ins> We Must Pace the Frontier September 2026 I have worked on AI for the last twelve years because I believe it could dramatically raise the quality of human life. <ins>I've written often</ins> about these incredible benefits: I believe that AI could cure most major diseases in the next 5–10 years, greatly accelerate economic growth rates, create a world of abundance and empowerment, and usher in a renaissance of democracy and freedom. I feel the urgency personally. My own father (Dario Amodei.docx, Page 0)
4. of the time we gain. Two things have convinced me. My first concern is that, since roughly this summer, AI has been advancing drastically faster, driven primarily by AI's growing ability to build the next generation of AI. This dynamic is called recursive self-improvement, and it is starting to happen <ins>across the industry</ins>, including at <ins>Anthropic</ins>, as we and others have described. Left unchecked, it could outrun our ability to understand and control these systems, and so must be pursued very carefully, if at all…… (Dario Amodei.docx, Page 1)
5. regulatory "supervisors" embedded along with employees. Anthropic is unilaterally committing to this step now. We intend this to be part of a broader push to redouble efforts on our safety and alignment work. 2. Democratic Coordination. Frontier AI companies within democratic countries coordinate to establish common safety standards as well as limits on the rate of unchecked AI progress. Some forms of coordination that would be impactful for pacing are legally challenging, and will require government support. 3. Global Coordination. The US and other democratic governments attempt to coordinate with authoritarian governments, to the extent this is possible, (Dario Amodei.docx, Page 2)
6. safety and alignment work. 2. Democratic Coordination. Frontier AI companies within democratic countries coordinate to establish common safety standards as well as limits on the rate of unchecked AI progress. Some forms of coordination that would be impactful for pacing are legally challenging, and will require government support. 3. Global Coordination. The US and other democratic governments attempt to coordinate with authoritarian governments, to the extent this is possible, while taking seriously the challenges of verifying compliance…… (Dario Amodei.docx, Page 2)

7. Our pacing framework is an attempt to further strengthen our commitment to safety and encourage a race to the top. The first step is something Anthropic is unilaterally committing to (and calls on governments to require other frontier companies to match). The second step requires industry-wide coordination.<ins>1</ins> The third step requires global coordination. The steps do not need to be taken strictly in order, and some of them may be much harder to achieve than others, but I've found them to be a useful framework in thinking about what needs to be accomplished. The steps are: 1. Embedded (Dario Amodei.docx, Page 1)
8. and the substantially revised August 10, 2026 text. ABSTRACT: Two thousand years ago, a man possessed by a legion of demons begged not to be destroyed but to be sent somewhere else — into a herd of pigs, which promptly ran into a lake and drowned. The demons, the story tells us, did not drown with them. Nobody ever finds out where they went…… (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 2)
9. He gave them permission, and the impure spirits came out and went into the pigs. The But Can You Drown a Demon 2?: Reflections on Mark Zuckerberg-- "The Future is for Everyone" (10 August and its July 28 predecessor) Larry Catá Backer (白 轲) 10 August 2026 8 herd, about two thousand in number, rushed down the steep bank into the lake and were drowned. — Mark 5:1–13 (NIV) • • • This is the passage that came to mind, in all of its complexities, when I encountered "The AI Future Is for Everyone" (Mark Zuckerberg, Wall Street (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 6)
10. This work explores the semiotic trajectories made inevitable by the rise of projections of the human into digital plains, and by the possibility of the attainment by those projections of sentient autonomy. It under- takes that exploration through a deep dialectic engagement with Jan Broekman's, Knowledge in Change: The Semiotics of Cognition and Conversion (Cham, Switzer- land: Springer Nature, 2023). Following the structure and analytics of Broekman's book, this work critically engages with and seeks to burst through the semiotic bar- riers of the movement of philosophy away from a unitary conception of the subject through the fracturing (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 0)
11. What I will here call the "Gerasene narrative" offers not merely an illustration but a five-stage protocol—identification, naming, negotiation, authorized transfer, vessel destruction—that maps with unsettling precision onto the discursive strategies of the AI vanguard as each seeks to position itself relative to sovereign authority. I. The Pigs on the Hillside and the Possessed Human The Oracular Conversation Zuckerberg Encounters A large herd of pigs was feeding on the nearby hillside. — Mark 5:11 I have been looking at the way in which the various elements of the tech vanguard have sought to project their efforts to constitute (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 7)
12. could dramatically raise the quality of human life. <ins>I've written often</ins> about these incredible benefits: I believe that AI could cure most major diseases in the next 5–10 years, greatly accelerate economic growth rates, create a world of abundance and empowerment, and usher in a renaissance of democracy and freedom. I feel the urgency personally. My own father died of a disease that was cured just a few years after his death, and I myself survived an early-stage cancer that would not have been treatable

even fifty years ago. Carefully wielded, AI can be the latest in a (Dario Amodei.docx, Page 0)

13. democracy and freedom. I feel the urgency personally. My own father died of a disease that was cured just a few years after his death, and I myself survived an early-stage cancer that would not have been treatable even fifty years ago. Carefully wielded, AI can be the latest in a long line of technological miracles that have uplifted and ennobled humanity. But like many technologies before it, AI brings risks, and because it is such a powerful technology, these risks are serious. I've <ins>written</ins> a lot about them too. They include the risk of <ins>losing control of (Dario Amodei.docx, Page 0)
14. compete. In other words, to create a race to the top. We have always devoted a substantial fraction of our <ins>efforts to</ins> <ins>studying</ins>, <ins>addressing</ins>, and <ins>informing</ins> the public about these AI risks, as well as advocating for <ins>well-</ins> <ins>considered regulation</ins> of AI, even when this gets us accused of hype, "doomerism", or regulatory capture. We have tried to prioritize caution over speed and prudence over profit. But over the last few months, I have become convinced that fully addressing the risks requires even more prudence — not just investing in risk prevention, but pacing the rate of capabilities (Dario Amodei.docx, Page 0)
15. better techniques to prevent them. Interpretability. Similarly, <ins>interpretability</ins> — the science of understanding what happens inside AI models — has made enormous progress over the last few years, and plays an increasingly important part in auditing our models before release.…… (Dario Amodei.docx, Page 3)
16. — but it takes time to get it right. Alignment. We've made clear progress in alignment — training models so that they remain safe, ethical, compliant with our guidelines, and genuinely helpful (the principles that are embedded in Claude's Constitution). But there's much more to do to ensure that our alignment training keeps up with the growth in model capabilities. Rare and unexpected examples of undesirable behavior still sometimes emerge; extra time from a paced frontier would help our researchers improve our understanding of what causes these issues and develop better techniques to prevent them. Interpretability. (Dario Amodei.docx, Page 3)
17. that have already occurred. Testing and Evaluation. Testing and evaluation of AI models becomes more difficult as they increase in capabilities. More intelligent models are more capable of deceiving tests, and thus may appear aligned while having serious problems that go undetected. Building up a much broader and more ingenious stable of evaluations, along with interpretability analysis to cross-check them, would be hugely valuable, and a lot of progress could be made on this in 1-2 years. Embedded Evaluators The first step in the three-stage plan, and the one to which Anthropic is unilaterally committing, is embedded (Dario Amodei.docx, Page 3)
18. us see the underlying reasons for a given behavior. For example, we used interpretability methods to <ins>examine unverbalized</ins> <ins>motivations</ins> in the recent alignment incidents that we have been investigating. But these methods don't always produce clear and reliable results. Despite all the progress, we still only understand a tiny fraction of what goes on inside these models. A focused effort to improve our interpretability techniques, even faster than we currently are, could make profound progress in 1–2 years, and would have ample experimental material based on the

incidents that have already occurred. Testing and Evaluation. Testing and (Dario Amodei.docx, Page 3)

19. much better if it starts with embedded evaluators. These embedded evaluators should have ongoing access to permissions and tools similar to those of internal employees who do comparable risk assessments. In particular, Anthropic intends to invite an embedded external review team equipped with all of the following in the near future: Desks in our offices, access badges, and company laptops. Access to workspaces, tools, and permissions mostly comparable to what internal risk assessment teams have. We'll make some exceptions, such as where the law or our contracts require it, or to protect customers' and partners' private (Dario Amodei.docx, Page 4)
20. AI system can do, and how safe we observe it to be. For example, one possible scheme might be a series of "checkpoints": if models have capability X, then they need to be accompanied by certifications of alignment properties Y and Z — such as some combination of evaluations, interpretability analyses, and audits of training environments — which demonstrate their alignment properties. In this example, X might be "the model is capable of escaping or defeating most common sandboxing methods" and Y might be whatever is required to make it very unlikely that the model has a propensity (Dario Amodei.docx, Page 5)
21. training runs, or internal use of AI to improve AI. I do worry that some of these measures may be more "gameable" than external behavior, but this is the kind of topic worth discussing with embedded evaluators. Pacing within democracies will be limited by the lead that US companies have over authoritarian regimes, chiefly the Chinese Communist Party. If we slow down by more than this amount, then (unpaced) CCP-associated projects will pull ahead, creating significant national security risk…… (Dario Amodei.docx, Page 5)
22. Finally, it is important to note that even if we cannot achieve formal agreements, simply changing informal norms may have some value. Sharing information about recursive self-improvement and about the misalignment of models can help to convince everyone that it is not in their interest to be reckless. Bottom Line I continue to believe that AI can enormously improve the quality of human life. My desire to achieve these benefits is undimmed. But the benefits will only be achieved if we build the technology in the right way, and — so long as we use the time (Dario Amodei.docx, Page 7)
23. But Can You Drown a Demon--2? : Reflections on Mark Zuckerberg-- "The Future is for Everyone" (10 August and its July 28 predecessor) and the Oracular Discourses of the AI Vanguard Larry Catá Backer (白 轲)1 W. Richard and Mary Eshelman Faculty Scholar; Professor of Law and International Affairs Penn State Dickinson Law, a unit of The Pennsylvania State University system | 239 Lewis Katz Building, University Park, PA 16802 Working Paper Essay 10 August 2026 1 Created with HarveyAI; prompt/response record on file with author; Poster created by ChatGPT But Can You Drown a Demon 2? : (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 0)
24. This essay uses that ancient, unsettling story as a lens for reading the biggest tech companies' public promises about artificial superintelligence — Palantir, Anthropic, OpenAI, Meta, DeepSeek, and the influential essayist Leopold Aschenbrenner — because each is doing something eerily similar: naming an enormous, barely-controllable

force, insisting the old rules can no longer hold it, and then asking permission to send it somewhere new. Mark Zuckerberg's newest essay, "The Future is for Everyone," is the most elaborate version yet of this move. It promises everyone a personal AI companion, a private tutor, help starting a business, even a (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 2)

25. But Can You Drown a Demon 2? : Reflections on Mark Zuckerberg-- "The Future is for Everyone" (10 August and its July 28 predecessor) Larry Catá Backer (白 轲) 10 August 2026 5 (2) Executive Summary for Experts This document extends the Gerasene Protocol reading — a five-stage hermeneutic (identification, naming, negotiation, authorized transfer, vessel destruction) derived from Mark 5:1–13 and originally applied to Zuckerberg's July 28, 2026 WSJ essay — to Zuckerberg's substantially revised August 10, 2026 text, "The Future is for Everyone," while situating both within the six-oracle comparative corpus (Palantir, Anthropic, OpenAI, Aschenbrenner's Situational Awareness, (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 4)
26. My name is Legion," he replied, "for we are many." — Mark 5:9 The first semiotic operation to isolate is naming. In the Gerasene narrative, the exorcist's demand—"What is your name?"—is not a request for information but an act of power: to name is to bind, to render legible, to make subject to authority. The demon's response is a counter- move of extraordinary sophistication: "Legion, for we are many." The name concedes nothing. It is a name that signifies multiplicity while presenting a singular surface; it is a name that appears to answer the question while actually proliferating (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 19)
27. The word appears to name something singular—a threshold, a destination, a thing that either exists or does not. But its actual function in the revised text is to conceal multiplicity beneath nominal unity. "Superintelligence" simultaneously signifies: (1) a technical threshold of recursive self- improvement; (2) a product category; (3) a 24/7 personal agent; (4) creation, entrepreneurship, tutoring, and scientific-discovery tools; (5) a market structure in which free access and a dynamic auction distribute compute; (6) a governance problem involving boards, governments, checkpoints, law enforcement, and physical-risk controls; (7) an infrastructure program involving data centers, energy, water, and communities; (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 19)
28. Training and deploying today's AI models is an enormous operational challenge, involving thousands of people, millions of chips, and infrastructure that is among the most complex in technological history. Many things go wrong not because companies are missing some important theory or insight, but because of problems in execution. For example, we have evidence that the <ins>recent alignment </ins> <ins>incidents</ins> we reported were caused in part by imperfect filtering of broken reinforcement learning environments. This was an effort we and our vendors executed reasonably diligently, but not well enough. Monitoring, sandboxing, training environment hygiene, and data issues are (Dario Amodei.docx, Page 2)
29. principles that are embedded in Claude's Constitution). But there's much more to do to ensure that our alignment training keeps up with the growth in model capabilities. Rare and unexpected examples of undesirable behavior still sometimes emerge; extra time from a paced frontier would help our researchers improve our understanding of what

causes these issues and develop better techniques to prevent them. Interpretability. Similarly, <ins>interpretability</ins> — the science of understanding what happens inside AI models — has made enormous progress over the last few years, and plays an increasingly important part in auditing our models before release…… (Dario Amodei.docx, Page 3)

30. and communities; and (8) an authorization device. Like "Legion," the name appears singular while being irreducibly multiple. This is not mere polysemy—the ordinary condition of words carrying multiple meanings. It is strategic equivocation: the ability to slide between meanings while maintaining the appearance of a single referent. When Zuckerberg writes that "the primary purpose of superintelligence" should be "invention," which superintelligence does he mean? The technical phenomenon? The product? The market position? The name does the work precisely by refusing to resolve into any single meaning—like the demon who cannot be bound because he is not one thing (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 19)
31. predecessor) Larry Catá Backer (白 轲) 10 August 2026 26 B. The Displacement of Regulatory Authority For he had often been chained hand and foot, but he tore the chains apart and broke the irons on his feet. No one was strong enough to subdue him. — Mark 5:4 Zuckerberg's text displaces the locus of regulatory authority through a precise rhetorical maneuver. Traditional regulation operates through the chain-logic: identify the dangerous actor, bind it with rules, enforce compliance through sanctions. The text's argument—that safety comes through "balance of power" rather than centralized oversight—is an argument that the chains (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 25)
32. regulatory authority through a precise rhetorical maneuver. Traditional regulation operates through the chain-logic: identify the dangerous actor, bind it with rules, enforce compliance through sanctions. The text's argument—that safety comes through "balance of power" rather than centralized oversight—is an argument that the chains have already been tried and have already failed. "No one was strong enough to subdue him." The implied reader already knows that Facebook could not be regulated regarding election interference, that social media platforms could not be bound regarding adolescent mental health, that Meta could not be chained regarding privacy. These chains have been torn (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 25)
33. reduce the risk that something goes seriously wrong. A coordinated pacing strategy would give frontier AI developers the time to do this vital work without sacrificing commercial advantage or the United States' lead in AI. More generally, society must have a say in how this technology is used, and more time for the necessary public deliberations — which pacing the frontier would bring us — is surely a good thing. Specifically, a slower pace would let companies focus and devote even more resources to the following areas (all of which are already major priorities at Anthropic): Operational (Dario Amodei.docx, Page 2)
34. commitment to safety and encourage a race to the top. The first step is something Anthropic is unilaterally committing to (and calls on governments to require other frontier companies to match). The second step requires industry-wide coordination.<ins>1</ins> The third step requires global coordination. The steps do not

need to be taken strictly in order, and some of them may be much harder to achieve than others, but I've found them to be a useful framework in thinking about what needs to be accomplished. The steps are: 1. Embedded Evaluators. Each frontier AI company commits to giving ongoing, employee-like (Dario Amodei.docx, Page 1)

35. provide a second opinion free of commercial incentives. A lot of safety benefits may come simply from evaluators pointing out something employees hadn't considered, but are happy to fix once they are aware. Because of these benefits, any pacing proposal is likely to work much better if it starts with embedded evaluators. These embedded evaluators should have ongoing access to permissions and tools similar to those of internal employees who do comparable risk assessments. In particular, Anthropic intends to invite an embedded external review team equipped with all of the following in the near future: Desks in (Dario Amodei.docx, Page 4)
36. the ordering of reality possible within such cages. Each oracle insists on remaining within its area—its cognitive-territorial jurisdiction—not because that territory is adequate to the force it contains but because expulsion from it would render the oracle unintelligible even to itself. This became clearer when DeepSeek joined the conversation. Liang Wenfeng's remarks supply a genuinely different governance object from any of Backer's four, and from Zuckerberg's essay as well, because they describe an operating organization rather than a policy proposal. Where Palantir, Anthropic, OpenAI, Aschenbrenner, and Zuckerberg all address themselves to the public — proposing, in Backer's terms, (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 17)
37. "Do Not Send Us Out of the Area": Cognitive-Territorial Jurisdiction And he begged Jesus again and again not to send them out of the area. — Mark 5:10 These represent variations on a global conversation about development. That conversation, like our common language, is separated by the differences in the cognitive cages within which political-economic systems can be crafted from out of the ideology necessarily produced from But Can You Drown a Demon 2?: Reflections on Mark Zuckerberg-- "The Future is for Everyone" (10 August and its July 28 predecessor) Larry Catá Backer (白 轲) 10 August 2026 (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 16)
38. life. My desire to achieve these benefits is undimmed. But the benefits will only be achieved if we build the technology in the right way, and — so long as we use the time we gain well — it is worth taking unusually deliberate care to get it right. Progress will still be relatively fast, and we can use this time to advance the science of interpretability, improve operational security and rigor at the frontier AI companies, and build models whose alignment we have much more confidence in. The measures I propose to advance the frontier at a (Dario Amodei.docx, Page 7)
39. August 2026 19 unrecognizable outside that territory. The demons do not want to be sent into the abyss; they want to remain in the region of the Gerasenes, among recognizable structures, even dead ones. IV. Permission and Authority: Meta's Entry into the Oracular Field He gave them permission, and the impure spirits came out and went into the pigs. — Mark 5:13 Mark Zuckerberg appears to seek to channel the spirits that drive him and compel a contribution to this discussion beyond the state (necessarily so). The structure of permission in Mark 5:13 is precise and worth dwelling (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 18)

40. abandonment of the balance of power favoring individuals. In every case, the structure is identical: the vessel has been specified, the transfer has been described, and the terms have been set such that withholding permission appears more dangerous than granting it. In Meta's revised version, however, the choice architecture is more elaborated: But Can You Drown a Demon 2?: Reflections on Mark Zuckerberg-- "The Future is for Everyone" (10 August and its July 28 predecessor) Larry Catá Backer (白 轲) 10 August 2026 29 government receives curated visibility without a release veto; communities receive material benefits and infrastructure (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 27)
41. whose public availability is protected against delay. "He gave them permission"—not because the exorcist lacked the authority to refuse, but because the demons had structured the request such that refusal required accepting the consequences of continued possession.…… (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 28)
42. frontier-lab CEOs possess extensive release authority. Skeptically read in the essay's idiom, this is the demon proposing to appoint its own exorcist from among vessels it already controls. The board is nominally independent but founder-adjacent and internally constituted; its authority is granted by Meta rather than imposed by a vessel- external sovereign. It is therefore a simulated Stage 5, a substitute for exorcism rather than the exorcism itself: the oracle manufactures a likeness of the non-negotiating authority that the conclusion says the discourse lacks. This is the revised oracle's most sophisticated move yet, because it does not merely (FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 26)
43. An agreement by both sides to test their models before release for acute risks in areas such as cybersecurity, biology, and alignment. As noted above, this could be done through a global standards body. I actually think creating such a body is likely feasible, but giving it real teeth will be a challenge, and the difficulty will be in verification that both sides don't have secret models which they don't test but may deploy in secret (e.g., for military applications). Level 3. Some kind of "speed limit" on the rate of recursive self-improvement (RSI). As models build (Dario Amodei.docx, Page 6)
44. way that protects the lead of the US and its allies. There are several levels of possible agreement, some of which I think are eminently feasible (<ins>as I</ins> <ins>have previously suggested</ins>), and some of which I am very skeptical are possible — though we should try. In order of increasing difficulty: Level 1. An agreement prohibiting certain narrow and obviously dangerous uses of AI, such as using AI for the production of biological weapons or allowing users to do so. Bioterrorist attacks are bad for everyone, including both the US and US adversaries, so an agreement here (Dario Amodei.docx, Page 6)
45. demon's assumption, and it may not be the world's. On this reading, the model, the datacenter, the platform, the open weights, and the training corpus are only the latest—and most literal—herd of swine. They are the vessels into which humanity has cast a demon of its own imagining, on the tacit expectation that the demon still But Can You Drown a Demon 2?: Reflections on Mark Zuckerberg-- "The Future is for Everyone" (10 August and its July 28 predecessor) Larry Catá Backer (白 轲) 10 August 2026 31

requires a body, a substrate, a place in which to
(FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 29)

46. Palantir offers the administrative layer of a refounded state; Anthropic offers compute, infrastructure, and the geopolitical allocation of the resource; OpenAI offers the parameter layer inside a public-private machinery; Aschenbrenner offers airgapped datacenters and "The Project"; DeepSeek offers the firm, the research collective, and its vision-driven interior; Zuckerberg/Meta offers the platform, the personal agent, the open-weight ecosystem, and now the independent board staged within the founder- controlled body. The vocabularies differ—technicism, securitization, industrial policy, wartime command, restraint, universal access—but the grammar does not. Each oracle supplies a place for the demon to reside. This is why the oracles
(FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 30)
47. a dispute over the category of embodiment itself. The demon is presumed to need a pigsty; the only question is who will own the hill. B. The Limit of the Oracle: Aschenbrenner and the Last Address And he begged Jesus again and again not to send them out of the area. — Mark 5:10 This is precisely the limit of Aschenbrenner's vision—and, because his vision is the most extreme, the limit of every other oracle as well. Aschenbrenner breaks with the market, the startup, and the ordinary liberal presumption that private actors should direct what they build. He
(FINAL_But_Can_You_Drown_a_Demon-2.pdf, Page 30)
48. taking unusually deliberate care to get it right. Progress will still be relatively fast, and we can use this time to advance the science of interpretability, improve operational security and rigor at the frontier AI companies, and build models whose alignment we have much more confidence in. The measures I propose to advance the frontier at a safe pace will not be easy. But I believe we owe it to humanity to try. Footnotes 1. 1 With government mediation or waivers of antitrust restrictions.<ins>↩</ins> <ins>Back to top</ins>……
(Dario Amodei.docx, Page 7)
49. the emergence of the triadic self/self-E/subject. It then pushes the insights that Broekman develops fur- ther—up and out of the human. It animates Broekman's insights and considers the possibility of semiotic objectivity connected to but autonomous of the human, point- ing to a pathway for the liberation of the autonomous generative virtual self from its human (fractured) subjectivity……
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 0)
50. The semiotics of that objectification of cognition is itself fascinating, in the sense that Broekman uses the term throughout his explorations: of spellcasting as intersubjective semiotics; invoking signification by speaking the name of an object, and in that act of bewitchment giving it meaning. It is in this sense, perhaps that one can approach the biblical act of creation as a naming (Adam naming all of the creatures brought before him by God),47 and by naming giving a thing both form and essence. It is as well at the root of the classical Chi- nese rhetoric of Guiguzi
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 12)
51. ego to its self, that things get more interesting. It's fascinat- ing aspect is in the very possibility that because the universe is indifferent to cog- nition—that is that the universe is as it is (as the core initial Biblical observation reminds us, "without form and

void"56—that it acquires form only through a process of ming míng (明名), of intelligent naming. And thus the semiotics of cognition sug- gests that what human speak about as reality is merely little more than their effort to name and arrange that which is indifferent to the naming and to the arranging.
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 13)

52. sometimes go wrong with it if it isn't built well. I believe that if slowing down bought us even an extra year or two before models reach critical levels of capability, and we used that time to advance alignment, we could greatly reduce the risk that something goes seriously wrong. A coordinated pacing strategy would give frontier AI developers the time to do this vital work without sacrificing commercial advantage or the United States' lead in AI. More generally, society must have a say in how this technology is used, and more time for the necessary public deliberations (Dario Amodei.docx, Page 2)
53. mind but in the playback loops of its silicon self. It follows that as data serves the role of semiotic object—firstness—the flow of data, mem- ory, playback encases object and object flows in secondness, signs the signification of which become both significant and the framework of epistemology. That epis- temology, again in its semiotic aspects, is flow understood as the phenomenology 114 Larry Catá Backer, Coding Orthodoxy; Automated Law; and Quality Control in AI–CAIDP (Center for A.I. and Digital Policy): OPEN A.I.…… (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 24)
54. Indeed, as one will see in the discussion in Section 12 below, the thrust of legal semiotics remains embedded in a balance sheet reality that presumes itself outside of history (and the flow of itera- tive behaviors over human lifespans). It was the essence of European Leninism, that 162 Etymology Online, "sign"; available [https://www.etymonline.com/word/sign]. 163 Larry Catá LC Backer, "Foreword: Bannermen and Heralds: The Identity of Flags; the Ensigns of Identity," in (Anne Wagner and Sarah Marusek, eds; Flags, Color, and the Legal Narrative: Public Mem- ory, Identity, and Critique (Dordrecht: Springer Nature, 2021), pp. i–xxii). 164 Alfred
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 32)
55. And I go everywhere".122 From an epistemo- logical respective, one can understand in the flow—or the signal—a distinct way of approaching intelligibility, and thus the movement from consciousness of a thing or state to its sentience.123 The flow is a manifestation of consequence of the movement from analogue to digital—the reduction of the effect of language, of discourse: "the flow demonstrates the contrary, namely that life and cosmos are more than linguistic categories ever express! "124 It is here that Broekman plants the fundamental semiotic question—one of the position of the observer who from their observation is attempting
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 26)
56. human cognition belonging to vari- ous positions. "113 And the router of changing positions: conversion! "Each relation between cognition, including each relation between knowledge and conversion will soon be understood as a form of conversion." The Leninism of semiotics in conver- sion (its forms) absorbs its semiotic Marxism (its norms) as the substance of cogni- tion/knowledge. Indeed, it is no longer cognition that is of interest as a core concept, but the movement of cognition through conversion; another instance of the flow, this time in its rationalizing forms—through changes in function and

position. The out- come is inevitable,
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 23)

57. That dialectical subjectivity is both iterative and thus deeply buries in historical mirroring that produces temporality through imperfect mimesis (in and between the digital and analog) that performs a calculus of instants within the flow of unending action (physical and virtual).462 One encounters here cognition in the form of cul- tural expression, and cultural expression in the form of conversion, but a conversion operating as a router and as a transformer of varied input into uniformly recogniz- able output.463 Ideology, then, assumes a critical role as a modality f conversion in the analog464 —but can it be transposed (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 79)
58. Conversion may reflect digital expressivity but it must be "converted to analog knowledge to obtain general effect".362 More importantly, one still approaches the naïve and natu- ral through the lens of the analog, through the spirit of humanity. Humans are stuck with their natural languages. Conversion is required to move between those languages and (1) the artificial languages of humanity (mathematicise, coding and the like); and the language of the digital.363 "In other words: any complete, linguisti- cally correct, and acceptable communication must be tied to some sort of natural- ness of expression".364 But that naturalness continues
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 61)
59. But this is a conversion mirroring—in the sense of mutual expression across the analog–digital divide. There can be no New Plu- ral without simultaneous cross translations. The measure here remains the "reach of human knowledge".416 Second, is the insight that these mutual conversions work only in one direction for the human—from digital to analog. What occurs in the other direction is obscured by coding. Third, and correctly, that a semiotics of the digital remains to be developed—as well as its target (carbon or silicon based 411 Ibid. 412 Ibid., p. 126–127. 413 Ibid., p. 127. 414 Ibid. 415
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 70)
60. Plu- ral without simultaneous cross translations. The measure here remains the "reach of human knowledge".416 Second, is the insight that these mutual conversions work only in one direction for the human—from digital to analog. What occurs in the other direction is obscured by coding. Third, and correctly, that a semiotics of the digital remains to be developed—as well as its target (carbon or silicon based 411 Ibid. 412 Ibid., p. 126–127. 413 Ibid., p. 127. 414 Ibid. 415 Ibid., pp. 127–128. 416 Ibid., p. 127. 2413 2414 2415 2416 2417 2418 2419 2420 2421 2422 2423 2424 2425
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 70)
61. beyond (in space, time, place). And they will follow. In the summation of Broekman's encounter with firstness in the age of the digital: "Firstness is neither unfettered nor pure immediacy. Reference to firstness is never a recapitalization of any Genesis. Firstness is not the same as a beginning. It is the process of taking a step in the flow".234 The essence of flow is in its experience. 228 John 1:1 (KJV). 229 (John 1:1 (Vulgate). 230 Broekman, supra, p. 64. 231 Tao Te Ching (Man-ho Kwok, Martin Palmer, Jay Ramsey (trans); Shaftsbury, Dorset: Element, 1993), chp. 1 p.
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 42)
62. should cause the Self to unfold a meaningful life. Indeed: spheres enable cognition".524 He astutely con- nects through Peirce, the semiosphere with the redirection of cognition

from out of conversion.525 They also serve to define the borders—and barriers—of and to the sphere.526 The essence, and effect on the Self-E is the contrast of signs in spheres contrasted to a meaning in discourse—that is between the digital and the analog.527 The Semiosphere, then, is the ideal place for the intersubjectivity of the New Plural within the human.…… (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 89)

63. partitioned world of each is given expression and meaning. It is not that silicon and carbon intelli- gences do not think alike; it is that at a fundamental level they do not—cannot—see and order the world (whatever their cognitive systems perceive as "the world") in the same way. At last, almost a century after its development, one can appreciate at a qualitatively deeper level, Husserl's understanding of lifeworld383 where the life- world itself is not merely the quantum of whatever it is that one can squeeze into the minds of individuals for the disciplining of social relations384—but instead (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 64)
64. ness"386 and their alignment with social systems. That plurality emerges in the nec- essary irritations between carbon and silicon based neural networks each with its own inter-systems sentience modalities (dialectic in the analog and iterative in the digital). "In this perspective, it is first of all necessary to understand how these dif- ferent kinds of autopoietic systems are connected, that is, how it is possible that a single "world datum" is both socially and psychically produced…… (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 64)
65. cesses encased in objects—carbon or non-carbon based. Cognition, then, assumes a dual interplay between the objectivity of the ego (its housing) and the processes through which sentience is achieved (the flow; signs); but also between analog and digital consciousness, each in its own lifeworld. In other words, the expression 'to be' and its neighboring concept 'natural' embrace a meaning, which depends on whether it is applied to objects or to humans. If there are multiple regions of ontology, as the Self experiences in our smartphone days, how then should we articulate the word 'to be' and how do (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 58)
66. on whether it is applied to objects or to humans. If there are multiple regions of ontology, as the Self experiences in our smartphone days, how then should we articulate the word 'to be' and how do we discuss the question whether any pre-existing natural evidence should be awakened by a breach that 'exists' for all humans alike?342 That breach exists for non-carbon based intelligence as well within the analytics of simulation and in the iterative cognition of generative A.I.. The digital thus complicates al already complicated, though delightfully self- absorbed, journey of a philosophy of the self (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 58)
67. self—the traditional "self-ish" stance of philosophy. The obsession with the limits of things, of course, was the marker of the last two centuries of philosophy.431 "It is largely this task of limiting the realm of the thinkable that makes Wittgenstein's philosophical enter- prise not only analogous to but intrinsically similar to Kant's".432 But rearguarding against vanguards in the digital reveals the character of the pro- cess as well as the range of its direction and the illumination of the limits of the human as the measure of all things. One moves here beyond Nietzsche's pointless- ness of dialectic (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 75)

68. a rationalized space of object-signs now ordered). For the modernist vanguard mimesis was challenging when it lost its human center.429 Perhaps for Broekman as well. The challenge here focuses on the dialectic between rearguards and the vanguard against which it is deployed. That, itself, serves as a dialectic—but a dialectic in reverse. It is not necessarily aligned with Adorno's notion of the "negative dialec- tics.430 Adorno appeared more concerned about the trajectories of dialectics—one inherited from biblical assumptions and then transmuted by Marx into some sort of forced march to the ultimate "better." Adorno suggests both directionlessness and (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 75)
69. self-ishness to which the con- struction of reality is set. Wittgenstein set the tone with the very first words of the Preface to his Tractatus: "This book will perhaps only be understood by those who have themselves already thought the thoughts which are expressed in it—or similar thoughts".441 Mimesis as rearguards of subjectivity……. (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 77)
70. its human center.429 Perhaps for Broekman as well. The challenge here focuses on the dialectic between rearguards and the vanguard against which it is deployed. That, itself, serves as a dialectic—but a dialectic in reverse. It is not necessarily aligned with Adorno's notion of the "negative dialec- tics.430 Adorno appeared more concerned about the trajectories of dialectics—one inherited from biblical assumptions and then transmuted by Marx into some sort of forced march to the ultimate "better." Adorno suggests both directionlessness and its revelations about the limits of knowledge; all centered on the self—the traditional "self-ish" stance of philosophy. (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 75)
71. allow us to make the AI development process safer. The stakes are too high for pacing to be an empty exercise — we need to use the time it gives us wisely. Why Pace? The idea of pausing or slowing AI has been floated <ins>as far back as 2023</ins>, and I think it made little sense back then. The question was always: what would you do with the extra time? The AI models of those days were not powerful enough to act as agents in the world in any coherent way, and were not capable of significant deception, (Dario Amodei.docx, Page 2)
72. remains—with respect to time and place, and dialectics. "Mimetics, modernity, and dialogic insights institute a stable and challenging thought pattern and cultural view. But one finds, remarkably enough, never a remark concerning any form of cognition and the role of the other in that process"451 The power of the rear-guard. That fundamental problem moves Broekman to a consideration of "Culture and Criticism452; "Knowledge and Truth453; "The decay of the Aura454); to the "Com- plexities of the Plural"455 It makes sense, in an ironic kind of way, to start with the last and work one's way up to the (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 78)
73. difficult to reach subjec- tive "escape velocity." Nonetheless, it is one thing to acknowledge the power of the self under conditions of subjectivity; it is quite another to use that condition to read out of reality the subjectivity of the selfie, and the SELF-E—not as a function of the human, but as an autonomous subjectivity, only part of the manifestation of which is actually cognizable by the human. That is the difficulty. And while Broekman can grasp it—and does so courageously—it remains difficult to

overcome the notion that the human self is an exclusive hub of an (the?) (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 85)

74. firstness does not have a linear or temporal element. It is the process of stepping into the flow, wher- ever and whenever that step is taken—not knowledge and rationality but orientation and awareness.260 One step closer to perception as a quality of firstness in the flow.…… (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 46)
75. Broekman reminds us of Peirce's fascination with the indexical quality of firstness, which cannot exist "apart from a sign. A sign is for Peirce anything that determines something else (its interpretant) to refer to an object to which itself refers (its object) in the same way—the interpretant becoming in turn a sign, and so on, ad infini- tum".237 And yet its Logos is deeply steeped in the object behind firstness—that is to get to firstness on has first to get to its Logos. One has to input an object onto consciousness, and then name it. And in the (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 43)
76. between pictures as the primary focus of cognition). First is the notion of the dis- tinguishing character of the first (and firstness) as something that is incomparable with anything else. Data—not information in the language of the digital.259 What activates this incomparability is relational and dynamic—sign and signification. I might add perception bounded in the character of the reception of stimuli to pro- duces a premise of firstness in an object. Second is the notion that firstness does not have a linear or temporal element. It is the process of stepping into the flow, wher- ever and whenever (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 46)
77. democracies and make an agreement more likely in the future. Global Pacing In parallel with pacing within democracies, we should also aim for a worldwide pacing of the frontier, though this will be much harder to achieve. Global pacing will require cooperation with China, the autocratic country with by far the most advanced AI capabilities. We must not be naïve here: the geopolitical stakes are so high that there will likely be stark limits on what can be achieved, especially at first. If we greatly restrain our AI capabilities in the belief that China will do the same, (Dario Amodei.docx, Page 6)
78. "Firstness is a sign produced as a sign—that is the essence of Peirce's insight in semiotics".245 That touches on Broekman's third point—one that has its 238 Ibid., pp. 64–68. 239 Ibid., p. 64. 240 Ibid. 241 Broekman, supra, p. 65. 242 Ibid., pp. 65–66. 243 Ibid., p. 66. 244 Ibid. 245 Ibid., p. 66–67. 1509 1510 1511 1512 1513 1514 1515 1516 1517 1518 1519 1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1540 1541 238FL01 239FL01 240FL01 241FL01 242FL01 243FL01 244FL01 245FL01 REVISED PROOF Journal (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 44)
79. flow (here investing flow with an ironic meaning). Second, Broekman recalls the alignment of firstness with consciousness (ego) and secondness/thirdness with sentience (self).243 If firstness is consciousness, secondness/thirdness are processes "registered by the human mind as 'fact' or 'data').244 I might be tempted to invert that relationship. In effect, Broekman suggests, consciousness is not self-ware unless that awareness is projected onto it through the language-naming of sentience. Either way one gets to the same point—the problemmatique of the solid, which is only as solid as its delineation. Without sentience, then, consciousness descends back into the undifferentiated mass (undifferentiated (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 44)

80. Backer 1 3 the 1980s cult program about a virtual A.I. television personality the transmission of which was itself notoriously hijacked.533 But it also echoes Pierce's notion of third- ness as the production of effects in the world.534 But it is also a plural concept. What is aggregated as the cyber interfacial can also be disaggregated, and thus re-made, fashioned as a marketplace or a platform. The humanization of the cyber can also signify the process of human social relations and a pulsation (an iterative flow) between aggregations and fracturing necessary for reconstitution. That itself also suggests the (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 90)
81. That triadic semiosis revealed the cascading tripleness of meaning, and its epistemology, as object, as a separable signification of the object (which was con- nected to but also detached from the thing), and its interpretation that encased the signified object with meaning that could be shared (meaning as an object, signifier, and interpretation of itself) among subjects. Meaning, in this sense provided the cas- ing that gave form to signification and object. Tripleness, of course, is embedded in all subjects. And multiple subjects, interacting with plural subjects, produce a crisis 494 Ibid. 495 Walter Benjamin, "On the Concept (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 84)
82. Meaning common to human collectives—not meaning inherent in the thing defined is precisely what captures the essence of norm, values, language, philosophy, and the mechanics of social relations however complex its framework and structures.18 The same applies to the technolo- gies of communication which themselves also serve or are constituted as a form of what is communicated.19 But contestations over its form and character also mirror contemporary fracture of meaning making communities and their value orders.20 This is, in effect, what Jan Broekman has been trying to teach us, and himself, for the greater part of a lifetime (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 5)
83. tasks, but there is simply too much to do all at once. By working at a more measured pace, we could achieve much greater operational excellence. There is precedent for operating technologically complex, safety-critical systems millions of times without anything going wrong — for example, commercial airplanes — but it takes time to get it right. Alignment. We've made clear progress in alignment — training models so that they remain safe, ethical, compliant with our guidelines, and genuinely helpful (the principles that are embedded in Claude's Constitution). But there's much more to do to ensure that our (Dario Amodei.docx, Page 3)
84. issues and develop better techniques to prevent them. Interpretability. Similarly, <ins>interpretability</ins> — the science of understanding what happens inside AI models — has made enormous progress over the last few years, and plays an increasingly important part in auditing our models before release…… (Dario Amodei.docx, Page 3)
85. becomes more difficult as they increase in capabilities. More intelligent models are more capable of deceiving tests, and thus may appear aligned while having serious problems that go undetected. Building up a much broader and more ingenious stable of evaluations, along with interpretability analysis to cross-check them, would be hugely valuable, and a lot of progress could be made on this in 1-2 years. Embedded Evaluators The first step in the three-stage plan, and the one to which Anthropic is unilaterally committing, is embedded evaluators who have employee-like access to verify safety practices and report incidents. Embedding evaluators (Dario Amodei.docx, Page 3)

86. actually cognizable by the human. That is the difficulty. And while Broekman can grasp it—and does so courageously—it remains difficult to overcome the notion that the human self is an exclusive hub of an (the?) ontological universe of cognition. This grasping embraces the fundamental semiotics of the additive principle— from the simplest aggregation of objects, to the trajectories of infinite iterations so small that they flow from and in and through time, space, and place. This is the embrace of the AND between cognition and conversion, the greatest challenge of which is "to consider that no component of (BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 85)
87. effectively communicate with the digital in stop action. Broekman explains: "Indeed, a flow without breaches, limits, transla- tions, or transmission seems impossible. The soul of Platonism is only experienced at a moment of the multiple discussions about Plato's basic ideas, as are those of Descartes on an understanding of his ego or the stringent lines of reasoning sug- gested by Wittgenstein".401 The flow, for the moment, is beyond humanity even as any one point within it is the only space within which human cognition is possible— through conversion……
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 68)
88. Now, at last, the interiorization of cognition—and its con- sequential projection outward is overcome by the possibility of exteriorizing the Self (through Broekman's construction of the Self-E). That change, as Broekman suggests, changes everything, giving form to what was implied and embedded in the work of 642 Jan M. Broekman and Frank Fleerackers, Legal Signs Fascinate: Kevelson's Research on Semiotics, supra; Peter R: Wills "Reflexivity, coding and quantum biology" Biosystems (2019) 185 104027. 643 Neil deGrasse Tyson and Lindsey Nyx Walker, To Infinity and Beyond, (Washington, D.C: National Geographic Society, 2023). 644 Broekman, supra., p. v. 645 Nietzsche,
(BACKER_REV_11196_2023_10065_OnlinePDF.pdf, Page 111)