VIRTUALIA



A COFFEE ON GOVERNANCE WITH ONE LOVE DAO
BY VAN AMENEYRO

FROM GATEKEEPERS TO WALLETS: THE NEW POWER STRUCTURES IN DIGITAL ART
BY FER CAGGIANO

THE PUNK OPERATING SYSTEM – BUILT WITHOUT PERMISSION
BY STEVE COULTER AKA 45RENEGADE

TOKENNATION AND THE GOVERNANCE OF THE DIGITAL FUTURE
BY MA. VITOR EMANUEL GRIPP, COMMUNITY MANAGER, TOKEN NATION

VOTING OR GOVERNING: SOME PHILOSOPHICAL PROBLEMS NOT SOLVED BY BLOCKCHAIN
BY PROF. DR. RODRIGO CID, UNIVERSIDADE FEDERAL DE OURO PRETO, BRASIL

VIRTUALIA MAGAZINE
VOL. 3, N. 2, 2026
DATA DE PUBLICAÇÃO:
25 DE MAIO DE 2026
DOI: 10.5281/ZENODO.20171962
ISSN: 3085-8364

UFOP
DEFIL
Programa de Pós-graduação

START VIRTUALIA ?

<Y/N>

CIP

VIRTUALIA MAGAZINE [ELECTRONIC RESOURCE]. – VOL. 3, N. 2, 2026. – OURO PRETO: FEDERAL UNIVERSITY OF OURO PRETO, PHILOSOPHY DEPARTMENT, POST-GRADUATION PROGRAM IN PHILOSOPHY, GIFLABS, 2026–
1 ONLINE RESOURCE: IL.

QUARTERLY
DOI: 10.5281/ZENODO.20171962 ISSN:3085-8364
GIFLABS PROJECT "VIRTUALIA MAGAZINE"

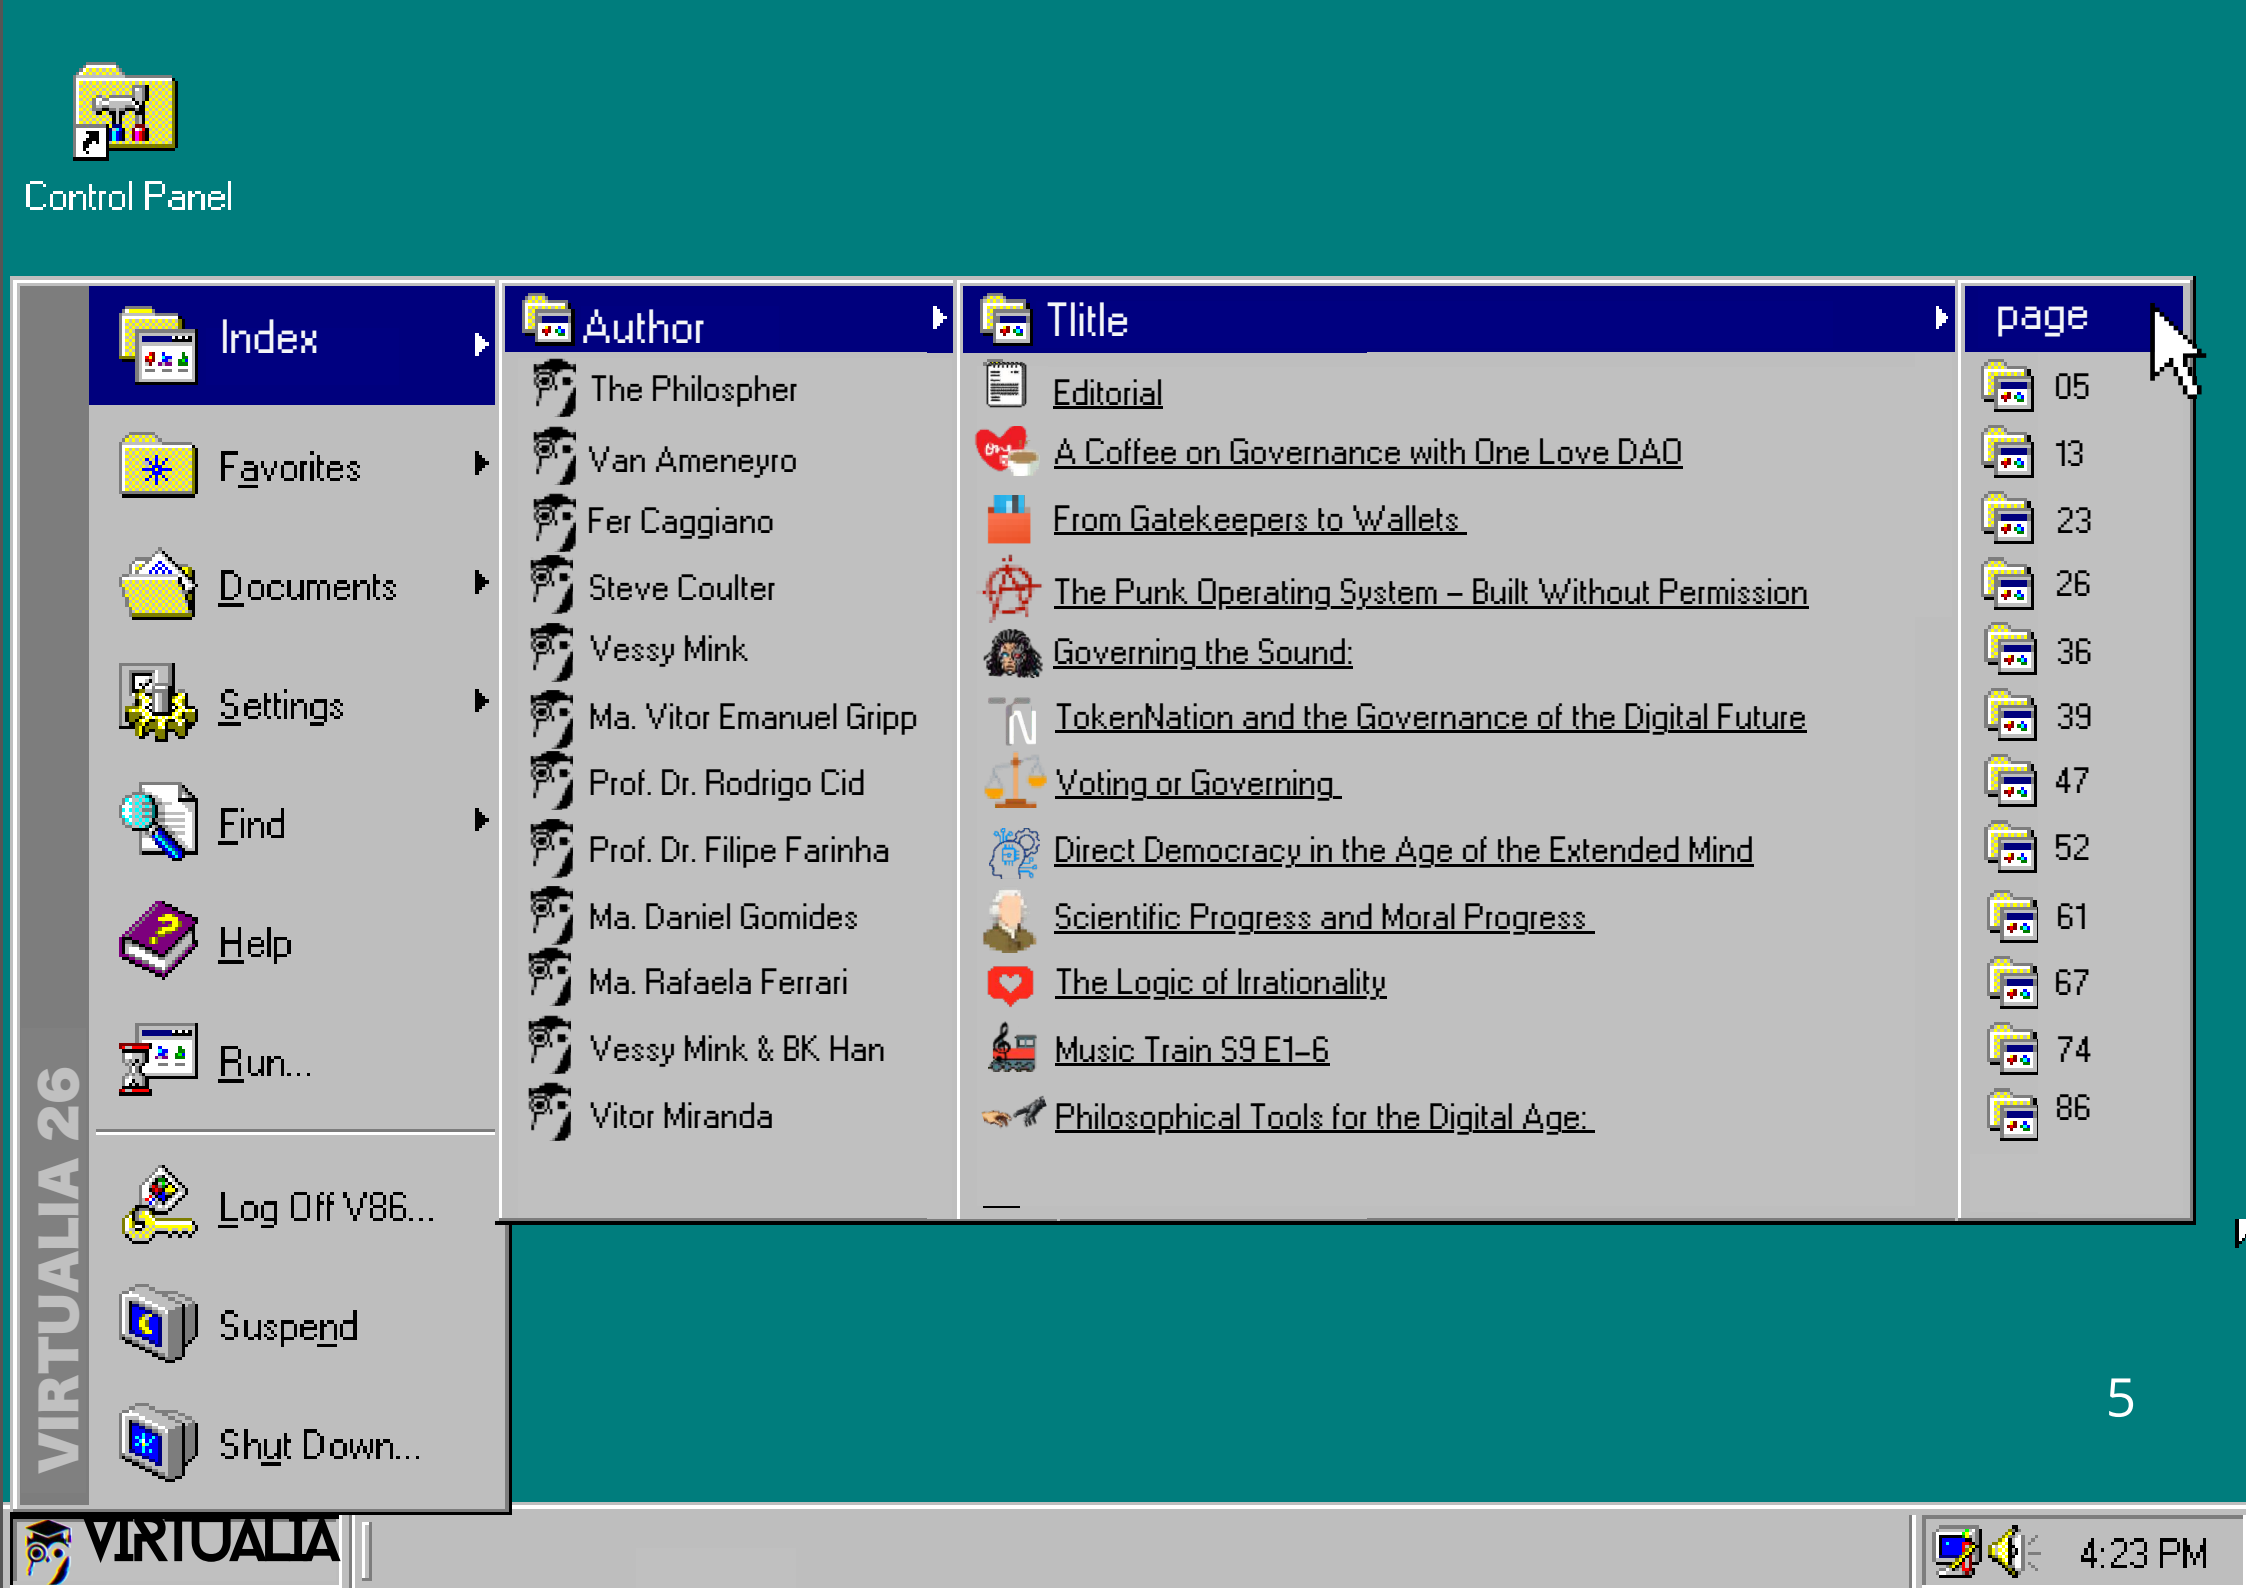
1. TECHNOLOGY. 2. WEB3. 3. PHILOSOPHY – RESEARCH GROUP PROJECT 4. CRYPTOART. I. FEDERAL UNIVERSITY OF OURO PRETO, PHILOSOPHY DEPARTMENT, POST-GRADUATION PROGRAM IN PHILOSOPHY, GIFLABS. II. TITLE.

CDD 006.8

Note Pad

File Edit Format View Help

- EDITOR-IN-CHIEF:
  - RODRIGO CID
- EXECUTIVE EDITOR:
  - VITOR MIRANDA
- WRITERS:
  - VAN AMENEYRO
  - FER CAGGIANO
  - STEVE COULTER
  - VESSY MINK
  - VITOR EMANUEL GRIPP
  - RODRIGO CID
  - FILIPE FARINHA
  - DANIEL GOMIDES
  - RAFAELA FERRARI
  - VITOR MIRANDA


Control Panel
Index
Favorites
Documents
Settings
Find
Help
Run...
Log Off V86...
Suspend
Shut Down...
VIRTUALIA 26
Author
The Philospher
Van Ameneyro
Fer Caggiano
Steve Coulter
Vessy Mink
Ma. Vitor Emanuel Gripp
Prof. Dr. Rodrigo Cid
Prof. Dr. Filipe Farinha
Ma. Daniel Gomides
Ma. Rafaela Ferrari
Vessy Mink & BK Han
Vitor Miranda
Tlitle
Editorial
A Coffee on Governance with One Love DAO
From Gatekeepers to Wallets
The Punk Operating System – Built Without Permission
Governing the Sound:
TokenNation and the Governance of the Digital Future
Voting or Governing
Direct Democracy in the Age of the Extended Mind
Scientific Progress and Moral Progress
The Logic of Irrationality
Music Train S9 E1-6
Philosophical Tools for the Digital Age:
page
05
13
23
26
36
39
47
52
61
67
74
86
5
VIRTUALIA
4:23 PM

EDITORIAL

GOVERNANCE AS A PROBLEM OF COLLECTIVE DECISION

BY THE PHILOSOPHER

Governance, in its essence, is the art of answering an unavoidable question: how do we decide together? From tribal councils to modern corporations, from cooperatives to DAOs and artistic communities, the challenge repeats itself. And at the heart of this question lies a mechanism that is deceptively simple yet philosophically treacherous: the vote.

Which voting system is the fairest? Arrow's impossibility theorem taught us that there is no universal answer. Every context — its size, its urgency, its level of trust among participants — demands a specific institutional design. This dossier does not offer ready-made formulas. Instead, it maps how different agents and sectors are confronting this problem.



VAN Ameneyro, in their conversation with One Love DAO, shows us how governance becomes a “necessary evil” in the world of digital art. The promise of horizontality runs up against concrete questions: who gets a seat at the table? Who is still left out? The answer, for VAN, lies in transparency, real participation, and the construction of institutional memory — something fragile in the volatile environment of Web3 platforms.



Fer Caggiano delivers an incisive diagnosis: power has not disappeared with decentralisation. It has merely moved — from institutions and curators to wallets and tokens. Community curation, in practice, often reproduces plutocracy. Their article forces us to ask: before voting, who defines what can be voted on? Governance begins with visibility.



Steve Coulter (aka 45renegade) offers an unexpected interpretative key: punk as an operating system. Before lean startup, before bootstrapping, before the creator economy, punk already practised independent production, direct distribution, and community building as survival strategies. His text reminds us that governance is not only about formal rules and votes. It is also about ethics, refusal, and the courage to build without asking for permission.



Vessy Mink appears twice in this issue, each time with a different governance experiment. In Governing the Sound, she presents Optimus Goddess, a licensing model where the artist retains full publishing rights and curation is replaced by equitable partnerships — an attempt to rewrite the rules of a historically extractive industry. In Music Train S9 E1-6, co-created with BK Han, she documents a live, collaborative song-minting project. The audience does not merely listen; it participates in real-time creation, turning musical production into a participatory governance performance. Here, the very act of making music becomes a collective decision process.

TOKEN NATION

Me. Vitor Emanuel Gripp, writing from inside Token Nation, shows how a technology event can become a living laboratory for governance. By bringing together academics such as Maria Goreti (Fiocruz), Carlos Frederico (UFOP/KryptoLab) and Rodrigo Cid (UFOP/GIFLABS), Token Nation does not merely discuss decentralisation — it practises it, in the curation of its stages, the selection of its projects, and the constant negotiation between efficiency and participation. His account, grounded in his own journey from exhibiting artist to community manager, reminds us that governance is not an abstract protocol but a daily, messy, collective achievement.



Dr. Rodrigo Cid, in Voting or Governing, returns to the philosophical bedrock of the problem. He reminds us that Arrow's theorem is not a mathematical curiosity but a structural warning: no voting system is neutral. By applying this lesson to blockchain, his article demonstrates that digital governance does not escape the aporias of collective choice — it merely translates them into code. He also develops the problem of many hands and the fragility of institutional memory in decentralised systems, connecting directly with the questions raised by VAN Ameneyro and Fer Caggiano.



Dr. Filipe Farinha, from the University of Saint Joseph in Macau, brings a comparative perspective. In Direct Democracy in the Age of the Extended Mind, he examines how different jurisdictions and cultural contexts shape the possibility of legitimate digital governance. His reflection on personal exocortices — AI assistants that extend a citizen's cognitive capacity — asks whether direct democracy might finally become feasible at scale, provided we solve the problems of authenticity, manipulation, privacy, and civic deskilling. The exocortex, for Farinha, is not a substitute for democratic agency but a prosthesis for it.



Ma. Daniel Gomides, from the Federal University of Ouro Preto, revisits Rousseau's First Discourse in Scientific Progress and Moral Progress. He argues that technical sophistication does not guarantee ethical advancement. His analysis of AI, hunger, and climate agreements shows that the gap between what we can do and what we should do remains as wide as ever — a sobering reminder that governance cannot be reduced to algorithmic efficiency. Rousseau's warning, written in 1750, still echoes today: science may teach us how to build better tools, but it does not teach us how to be better humans.



Me. Rafaela Ferrari Kley, from Degenerados Club, takes a complementary path in The Logic of Irrationality. Diagnosing a civilisation that has never been so technologically advanced yet remains emotionally manipulable, she recalls that Aristotle's logic was a civilising attempt to contain collective hysteria. Today's algorithms, however, are optimised for engagement, not truth. Her reflection forces us to ask whether governance can ever be purely rational, or whether it must always wrestle with the irrational architectures of attention, fear, and belonging. The question she leaves us with is disarmingly simple and profoundly uncomfortable: are we seeking to understand reality, or merely constructing emotionally bearable versions of it?



Finally, this issue closes with a special launch. Rodrigo Cid's The Philosophy of Artificial Intelligence (co-authored with Pedro Luiz Caetano Filho) is not merely reviewed here; it is treated as a governance artifact. Its systematic collaboration with AI systems, its transparent use of the CRediT-IA framework, and its rigorous discussion of Arrow's theorem, opacity, bias, and responsibility provide a conceptual toolkit for exactly the questions this magazine raises. The book is available for free download (DOI: 10.5281/zenodo.20143966), and we invite our readers to read it alongside the articles in this issue.

At the end of this volume, we hope that the reader will not find definitive answers. Instead, find a richer repertoire of questions. About who votes, about who decides what is put to a vote, about how we move from paper to action. About whether code can ever fully replace trust, and whether we would want it to. About what the classroom, the punk venue, the music studio, the DAO, the tokenised event, and the blockchain have in common: all are arenas where the same question echoes, again and again — how shall we decide together?

Enjoy the reading.

Prof. Dr. Rodrigo Cid
The Philosopher
(GIFLABS / Universidade Federal de Ouro Preto)

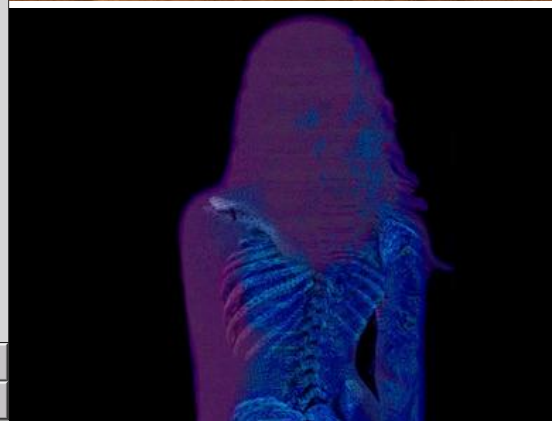
1. A COFFEE ON GOVERNANCE WITH ONE LOVE DAO

by VAN Ameneyro

What is wonderful about art and its many disciplines is how open and free it is, allowing it to expand and scatter seeds of emotion in others. Yes, art in itself is what sustains the human spirit above all else, yet it cannot simply exist in the anarchic way within our social system, for it runs the risk of being taken hostage, and it is, in fact, taken hostage by various forces. In traditional art, however, sufficient time has been devoted to developing structural guidelines and adapting policies that make art and culture a social pillar encompassing both the creator, the artist, and the recipient (the viewer).







Put like that, it sounds simple. However, the process that takes place between these two points, the sender and the receiver, is far more complex. This is where actors, strategies, and interests emerge that have regulated culture over time, and in this article, we speak more specifically about the arts. We will not go back to the times when art began to be institutionalized, nor to how it became necessary to establish parameters to maintain a series of actions that elevated it to a level where it is now even assigned a socioeconomic value, establishing its importance -as previously mentioned-as a social pillar.

What we will talk about, however, is a necessary ill that gives shape and above all direction to the arts of the 21st century and their mediums, specifically digital art and Web3. When I refer to power management as a "necessary ill or evil," I do not mean it disparagingly or as a rejection of its existence, but rather from the standpoint of its ideal potential versus its limitations within the ecosystem itself.

Governance is currently the "model" through which many associations and institutions have found a more horizontal way to manage the arts sector, leaving behind the vertical hierarchies that have reinforced the stark inequality among those involved in both traditional and digital art.

I have had the opportunity to be part of one of the most beautiful initiatives driven by One Love DAO: Visible Violets, an event in which women artists exclusively exhibit their works to showcase the depth and meaning of our gender, alongside the different realities each of us lives within our own cultures. Without a doubt, a comprehensive exercise that demonstrates One Love DAO's commitment to providing spaces for digital artists.



The work they do is admirable, and especially the way they integrate governance into their initiatives, communities, and the project as a whole. That is why co-founder Jennifer Pepen Aquino answered a few questions about how this essential form of management for digital art manifests within Web3.

V.

When you think about governance in the context of digital art and web3, do you see it as something distinct from how the traditional art world operates, or are they more entangled than people like to admit?

J,OneLove.

I don't see them as separate. They're more entangled than people want to admit. Web3 gave us new tools, transparency, ownership, participation. **But it didn't erase human behavior. Reputation, access, relationships, capital, those currents still shape outcomes.**

What's different is the opportunity. Web3 lets us redesign the system in real time. The question is whether we actually choose to.

It is precisely *how* we choose to let digital art evolve within the system that leaves us in a vulnerable state of credibility, and puts at risk the fragile roots being built, largely through trial and error.

V.

Governance in the art world has always been shaped by gatekeepers like curators, institutions, galleries. Do you see web3 as genuinely disrupting that structure, or has it simply created a new set of gatekeepers with different aesthetics but the same dynamics?

J,OneLove.

Both are true. Barriers are lower. Artists can move faster, build directly, exist without waiting for permission. That part is real.

But some level of filtering still exists, and not all of it is negative. Thoughtful curation, context, and selective visibility can protect against saturation and dilution. Without that, everything flattens and it becomes harder for meaningful work to be seen and sustained.

The goal isn't to remove structure entirely. It's to make sure those systems are more visible, more fluid, and not controlled by a small few.

This is when the expression "necessary evil or ill" I referred to at the beginning begins to make sense, as each of the elements that frame art give it seriousness and respectability. The elements that fall within a clear governance ,one that can provide solidity, are:

- Legitimacy
- Access and participation
- Value distribution
- Rules of the game
- Memory and archive: this last one is still a point to be developed further.

Regarding access and participation, the openness is vast. One can easily swing back and forth between developing and curating, but this mobility is dangerous because without the necessary objectivity, there is a risk of blurring disciplines and specialist professionals. It is in that blurring that the collector misinterprets the art and the artist becomes frustrated.

V.

When we talk about governance in digital art spaces, who do you think should have a seat at the table, and more importantly, who is still being left out of those conversations?
-from my perspective I have an idea but of course I want to know your point of view, because you as a DAO and as an artist that impulse visibility from all kind of artists, with different background to create an inclusive network with more opportunities I think there is a wider perspective)

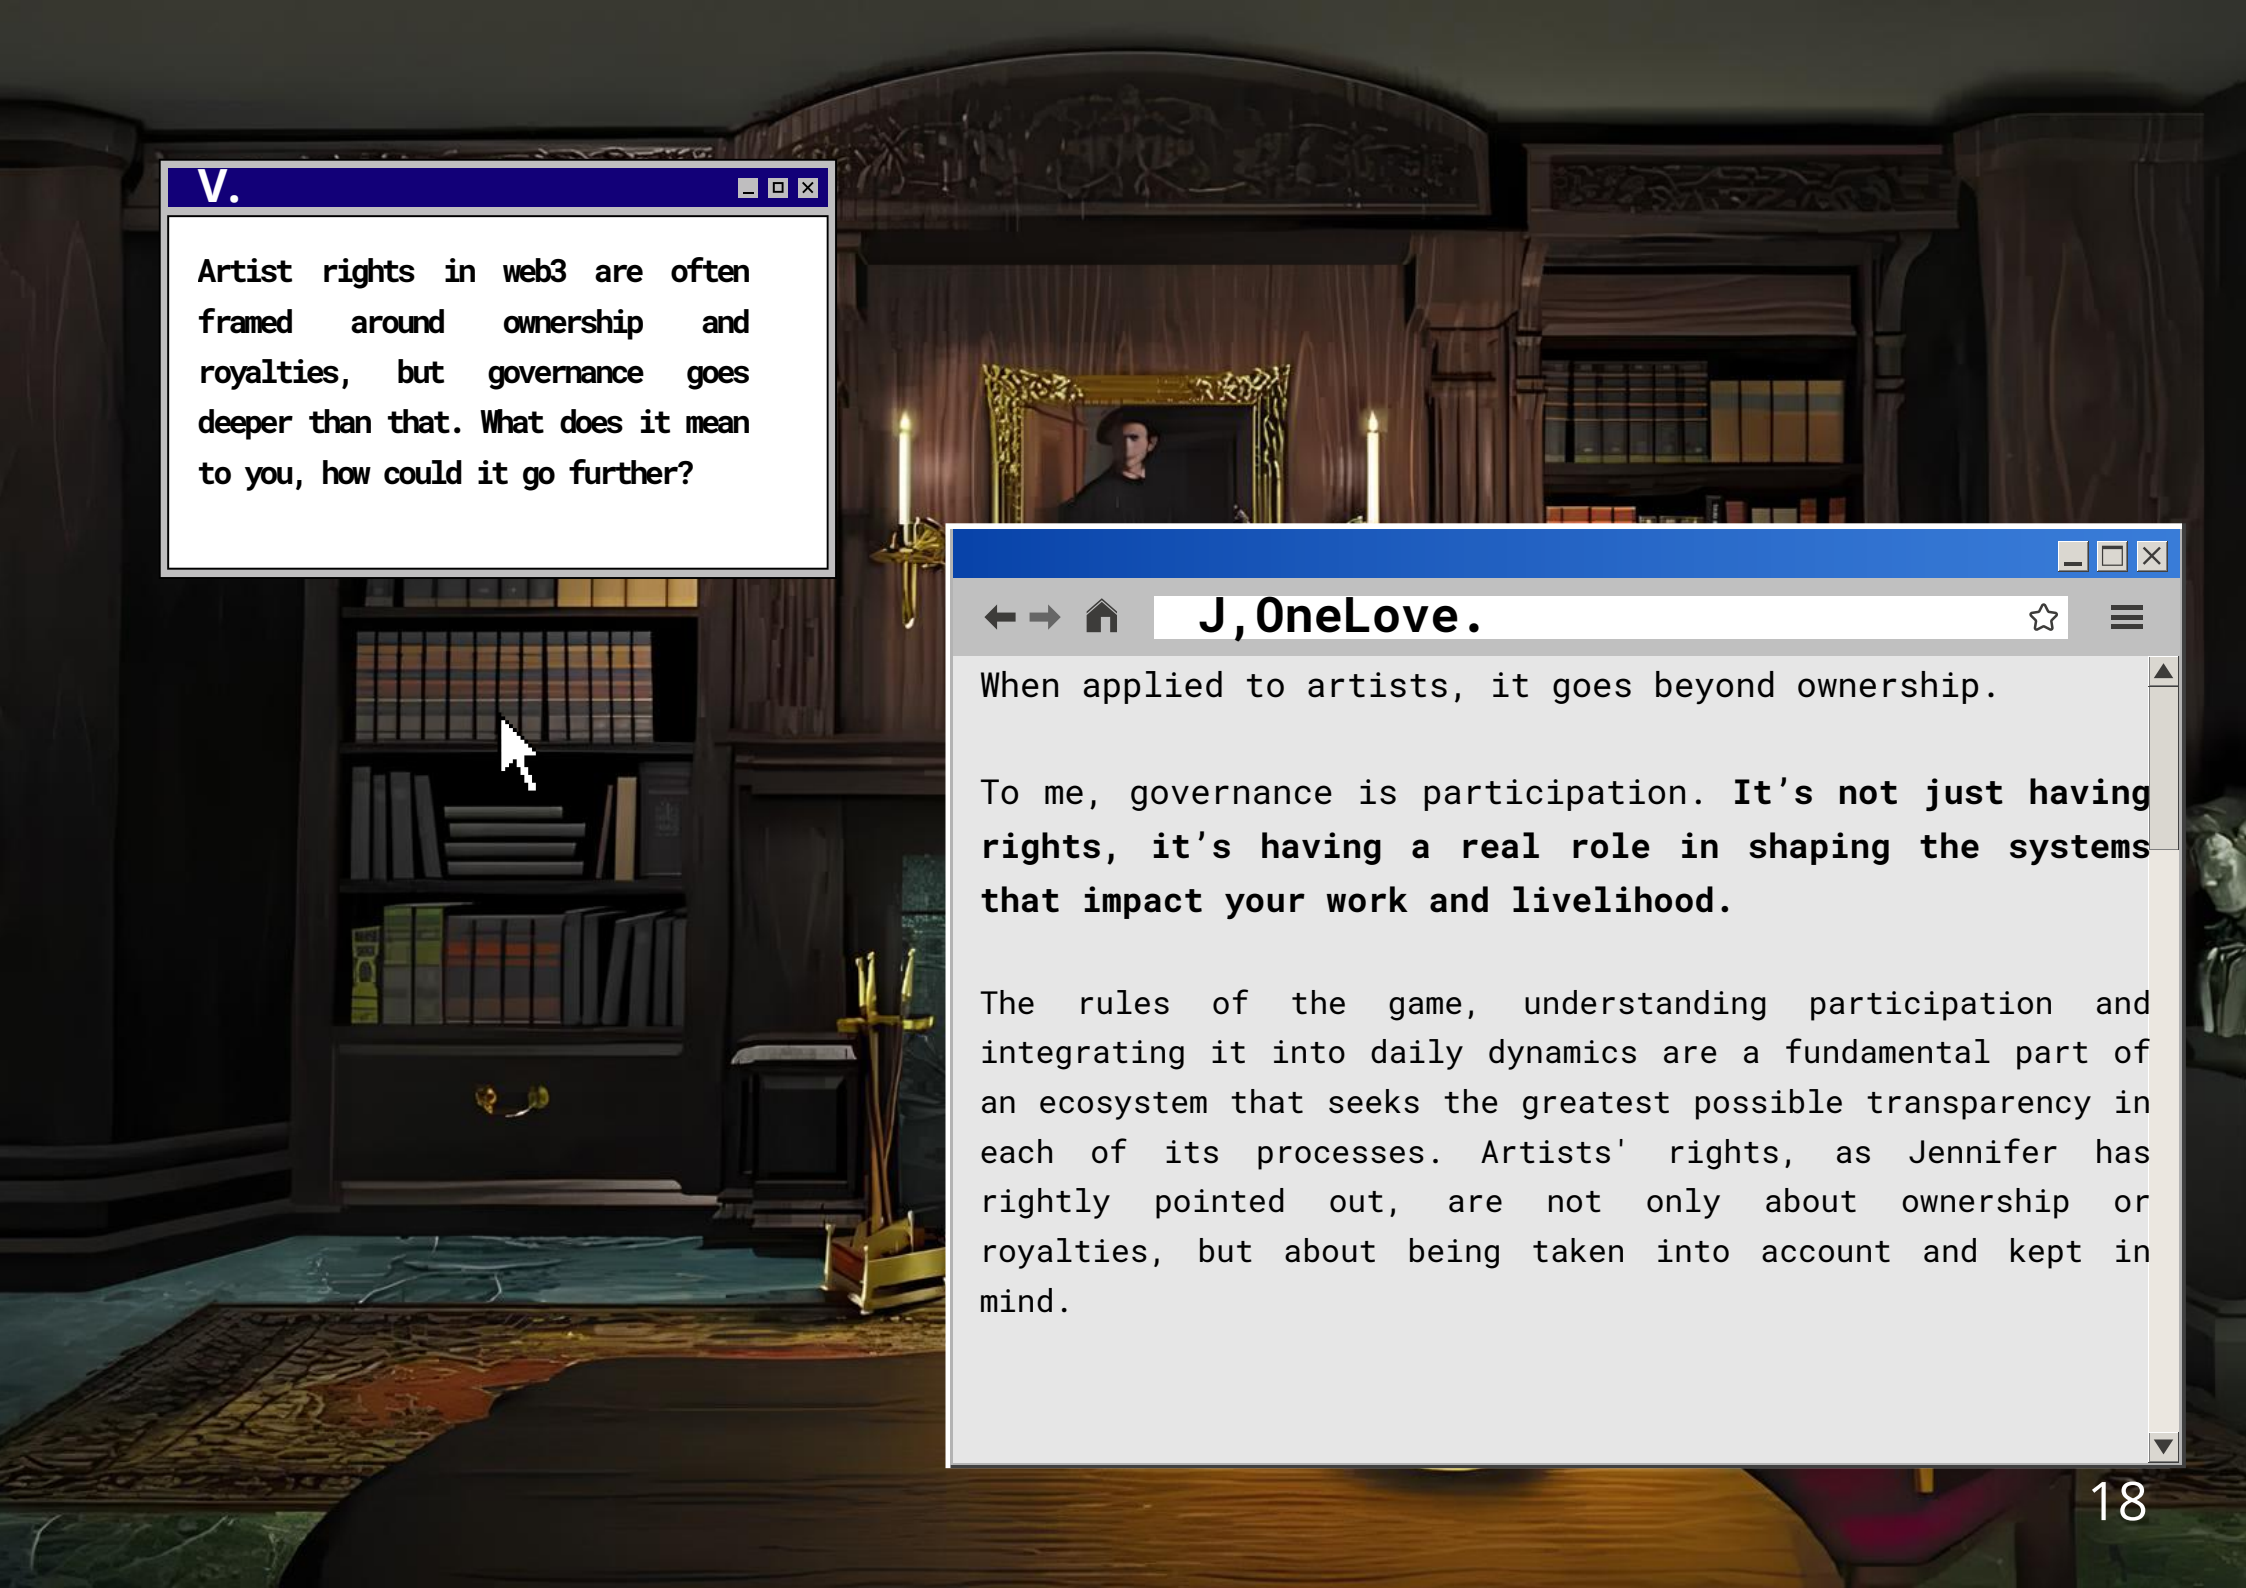
J,OneLove.

J, ONELOVE. Governance should center the people most impacted by these systems, especially working artists whose livelihoods are directly tied to them. It should also include collectors, curators, developers, and community builders who actively shape and sustain the ecosystem.

At the same time, many voices are still left out. Artists from the Global South, emerging artists without capital, women, disabled creators, older artists adapting to new tools, and communities without consistent access.

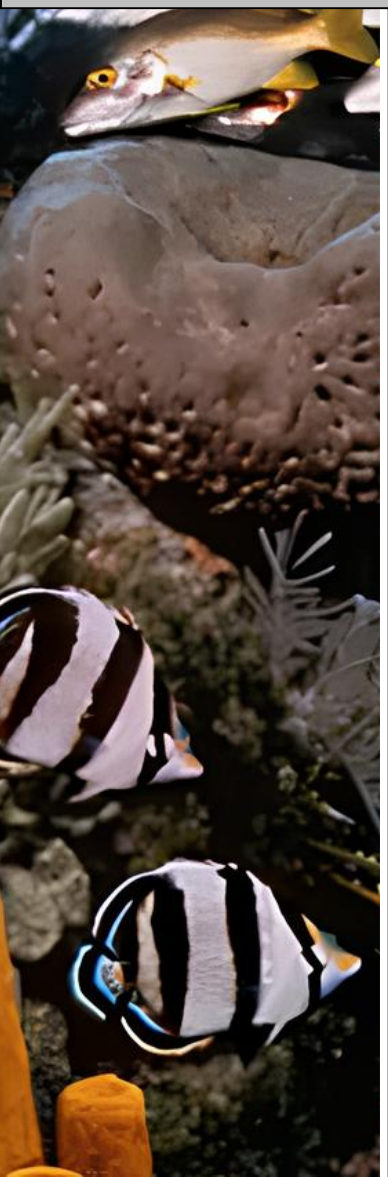
Governance becomes a really important tool to help those voices be seen, heard, and meaningfully included.

The balance is restored when people are given the opportunity to integrate into communities that, along the way, are worn down by trends often imposed within certain circles,directly impacting the artistic value a work can hold. This is where the importance of congruence between objectives within a developing structure lies, as it prevents getting lost in the constantly shifting discourses within Web3.


V.
Artist rights in web3 are often framed around ownership and royalties, but governance goes deeper than that. What does it mean to you, how could it go further?
J,OneLove.
When applied to artists, it goes beyond ownership.
To me, governance is participation. It's not just having rights, it's having a real role in shaping the systems that impact your work and livelihood.
The rules of the game, understanding participation and integrating it into daily dynamics are a fundamental part of an ecosystem that seeks the greatest possible transparency in each of its processes. Artists' rights, as Jennifer has rightly pointed out, are not only about ownership or royalties, but about being taken into account and kept in mind.

V.

Do you feel that the communities and platforms building in art spaces, traditional and web3, are being honest about their own power structures, or are there things that remain deliberately invisible?



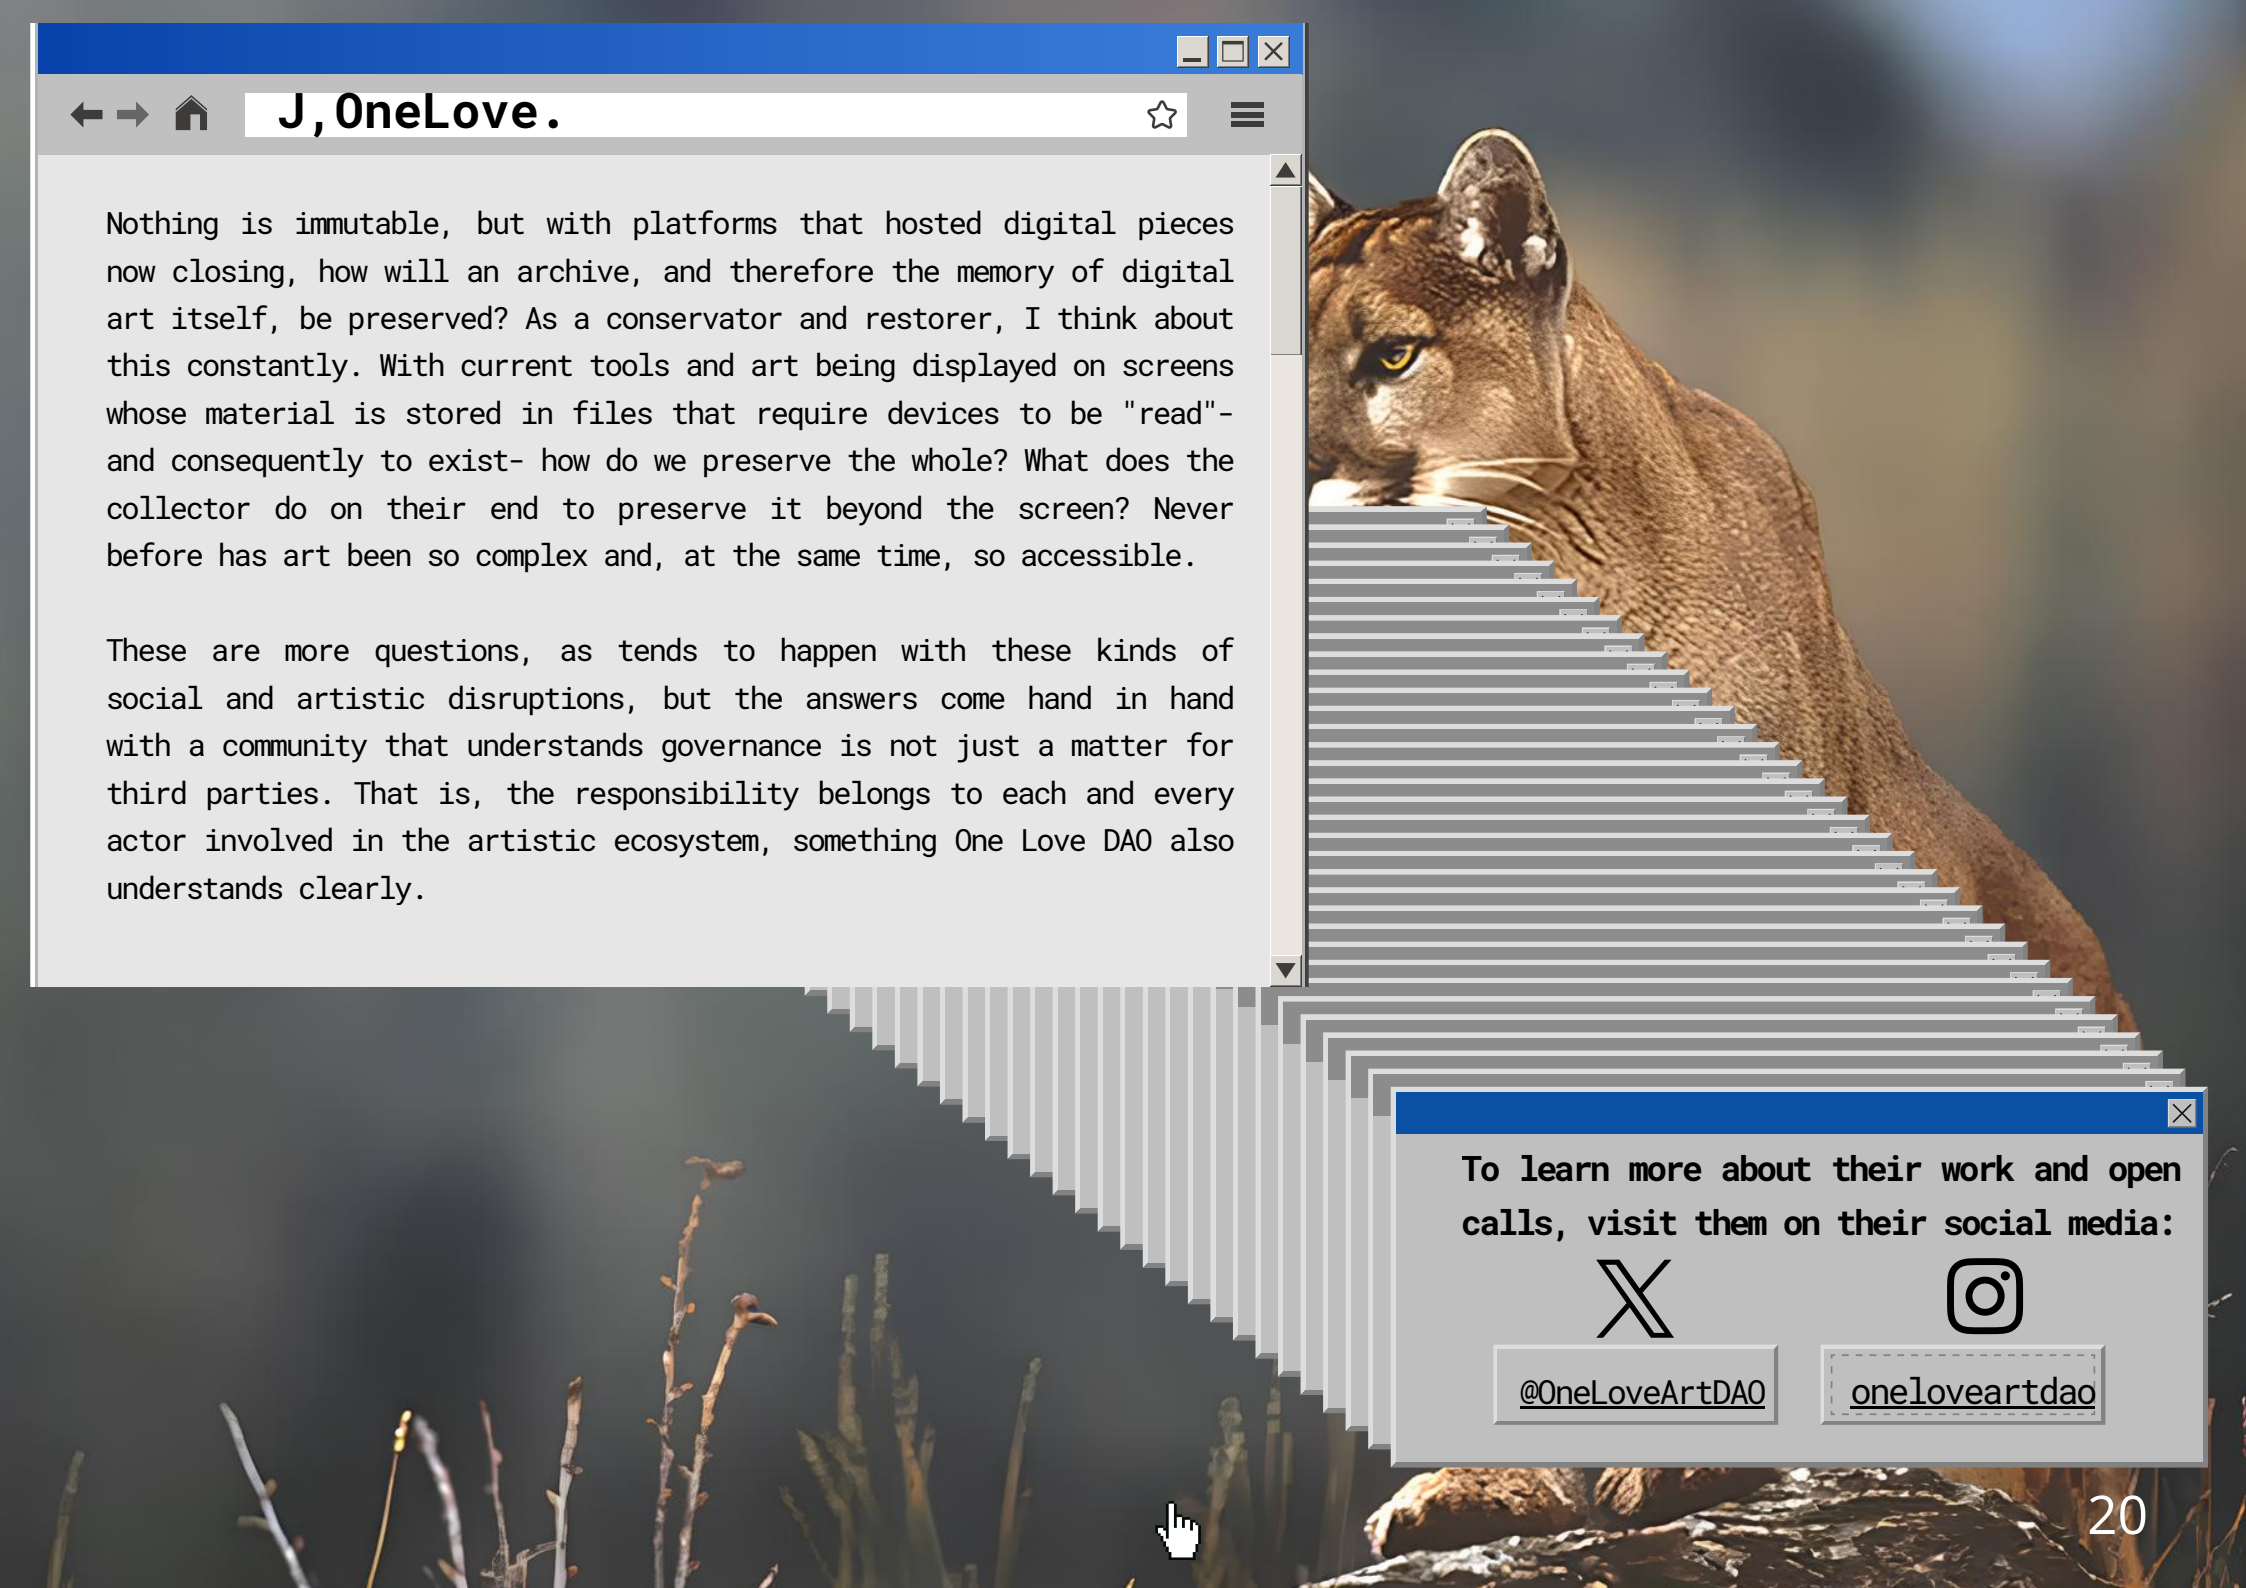
J,OneLove.

Not fully. There’s a lot of language around community, but decisions still sit with a few. That hasn’t changed as much as people think.

Invisible power shows up in quiet ways. Who gets selected. Who gets amplified. Who gets access. What happens behind the scenes.
Trust begins when that becomes visible.
Who decides. Who benefits. And how others can step into that circle.

I like power management as a term when discussing digital art and its platforms, because it is still entangled in various ideals. The way decisions are made directly impacts the artist's work, both in their practice and when it comes to the diffusion of their work.

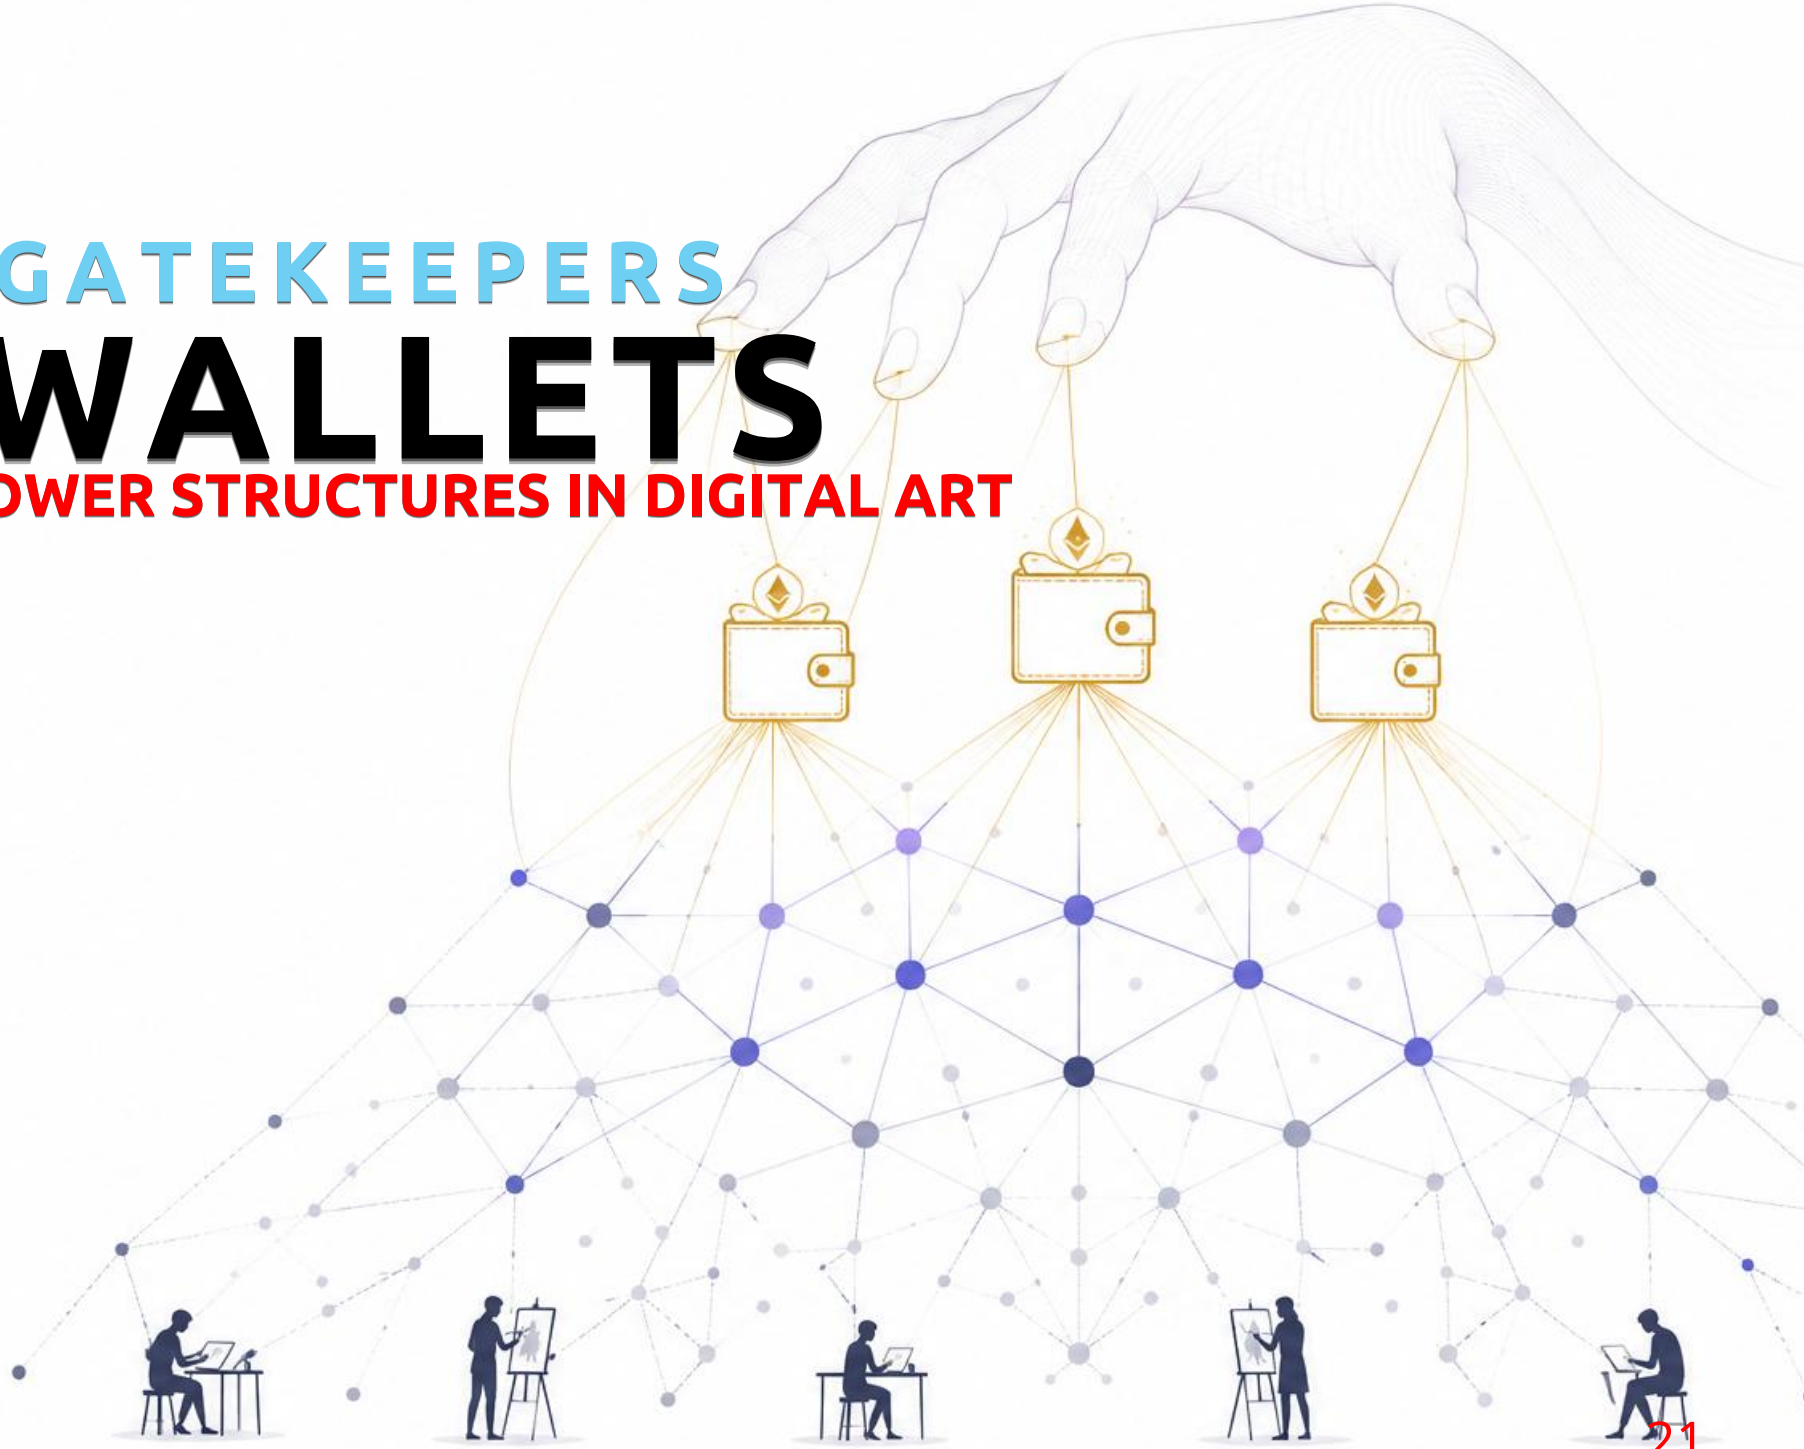
There is a clear gap between the promise and a reality that has yet to be fulfilled, which brings consequences such as a lack of credibility and the pronounced imbalance within the ecosystem itself, an ecosystem whose ultimate beneficiary is already questionable. Governance regulates, but for whom? There should be no doubt, yet there is, because even with a stated purpose, governance does not always align with its effect, as is the case with marketplaces; and this is where the topic of memory and archive comes in.


J,OneLove.
Nothing is immutable, but with platforms that hosted digital pieces now closing, how will an archive, and therefore the memory of digital art itself, be preserved? As a conservator and restorer, I think about this constantly. With current tools and art being displayed on screens whose material is stored in files that require devices to be "read"- and consequently to exist- how do we preserve the whole? What does the collector do on their end to preserve it beyond the screen? Never before has art been so complex and, at the same time, so accessible.
These are more questions, as tends to happen with these kinds of social and artistic disruptions, but the answers come hand in hand with a community that understands governance is not just a matter for third parties. That is, the responsibility belongs to each and every actor involved in the artistic ecosystem, something One Love DAO also understands clearly.
To learn more about their work and open calls, visit them on their social media:
@OneLoveArtDAO
oneloveartdao

FROM GATEKEEPERS

TO WALLETS

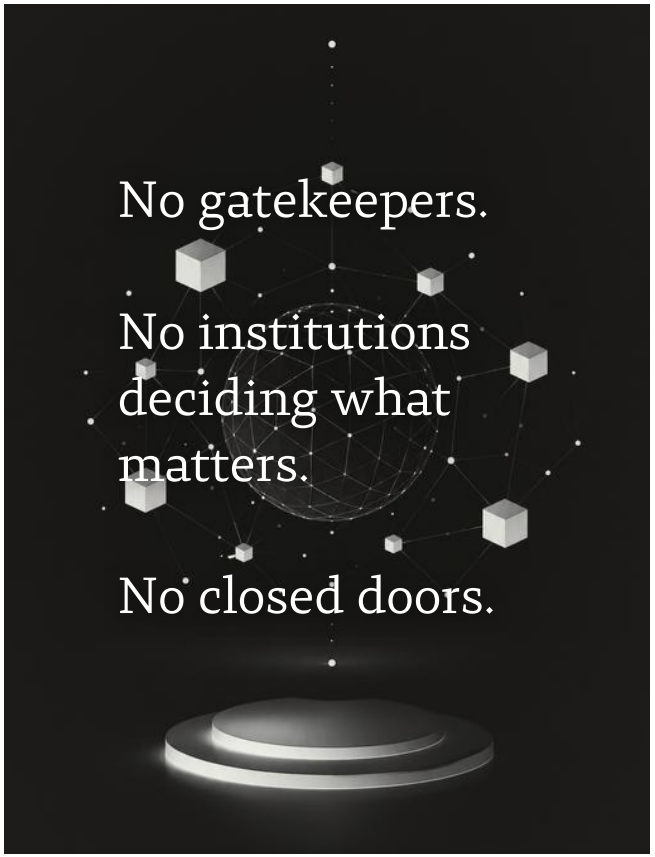
THE NEW POWER STRUCTURES IN DIGITAL ART



BY FER CAGGIANO

DECENTRALIZED, BUT NOT REALLY: WHO ACTUALLY CONTROLS ART IN WEB3?

WEB3 PROMISED A LOT.



For artists, it sounded like the long-awaited shift. A system where visibility wasn't controlled by galleries, curators, or algorithms—but by the community.

And for a moment, it felt real.

But the longer you stay, the more familiar it starts to look.

Not identical. But close enough to make you pause.

Because the truth is—we didn't remove gatekeepers. We just changed their shape.

POWER DIDN'T DISAPPEAR. IT MOVED.

In Web2, control was easier to point at.

Instagram decides what gets seen.
Galleries decide what gets shown.
Collectors decide what gets validated.

It was centralized, opaque, and frustrating—but at least it was clear.

In Web3, the structure looks different.

We have token holders instead of institutions.
Voting instead of curation (at least in theory).
Platforms like SuperRare and Foundation presenting themselves as open ecosystems.

But if you look closer, the dynamic hasn't disappeared.

It's just harder to see at first glance.

Control didn't vanish. It shifted—from institutions to capital, from curators to wallets, from closed rooms to visible systems that still concentrate influence.

THE MYTH OF DECENTRALIZED CURATION

"Community-driven curation" sounds good. It feels fair. It feels like progress.

But it assumes something that doesn't really hold up in practice: that participation is evenly distributed.

It's not.

Most people don't vote.
Most people don't engage in governance.
And when they do, their influence is often minimal.

Because in many systems, voting power is tied to tokens.

Which means the more you hold, the more your voice counts.

So what starts as decentralization slowly turns into something else—a structure where a small group of well-positioned players can heavily influence outcomes.

Not because they're evil.
Because that's how the system is designed.

And at that point, we have to ask:

Is this decentralization?
Or just a more transparent version of the same imbalance?

THE CABAL QUESTION

here's another layer people don't love to talk about.

Private group chats.
Collector circles.
Artists supporting each other strategically.

Call it a cabal, call it networking, call it coordination—it doesn't really matter.

It's human behavior.

People organize. They share information. They move together.

And in a system where visibility and value are still scarce, coordinated groups naturally gain an edge.

They can:

- Amplify certain artists
- Align around drops
- Influence narratives and attention

None of this requires a conspiracy.

Just aligned incentives.

And that's the uncomfortable part—because it means even in a decentralized system, power can still centralize. Just informally.

VISIBILITY IS THE REAL GOVERNANCE LAYER

We like to talk about governance as voting.

But in art, governance often happens much earlier—at the level of visibility.

Who gets featured.
Who gets highlighted.
Who shows up on the homepage.

Platforms like SuperRare and Foundation still play a major role here, whether explicitly or through design choices.

And that matters.

Because if people don't see your work, they can't vote for it. They can't collect it. They can't even consider it.

So before governance happens, selection already did.

Which leads to a simple but important realization:

Governance isn't just about decisions.
It's about attention.

WHERE DOES THAT LEAVE THE ARTIST?

THIS IS WHERE IT GETS A BIT UNCOMFORTABLE.

ARTISTS CREATE THE VALUE.
THEY PRODUCE THE WORK, THE CULTURE, THE REASON THE ECOSYSTEM EXISTS.

BUT STRUCTURALLY, THEY OFTEN SIT OUTSIDE THE CORE POWER LAYERS.

THEY DON'T HOLD THE MOST TOKENS.
THEY'RE NOT ALWAYS PART OF THE STRONGEST NETWORKS.
THEY RELY ON PLATFORMS AND COLLECTORS FOR VISIBILITY.

SO EVEN IN WEB3, MANY ARTISTS ARE STILL NAVIGATING SYSTEMS THEY DON'T REALLY CONTROL.

JUST NEWER ONES.

AND THAT GAP—THE ONE BETWEEN CREATING VALUE AND GOVERNING IT—IS STILL VERY MUCH THERE.

SO WHAT ACTUALLY CHANGED?

A lot, and not enough at the same time.

Web2 was opaque.
Web3 is more transparent.

You can see the wallets.
You can trace the votes.
You can observe the dynamics.

But transparency doesn't automatically create fairness.

It just makes the structure visible.

And maybe that's the real shift—not that power disappeared, but that now we can actually see it.

WHERE GOVERNANCE COULD GO FROM HERE

If Web3 is going to evolve, governance needs to move beyond the current defaults.

Not in a perfect, idealistic way—but in a more intentional one.

A few directions that feel worth exploring:

REPUTATION-BASED INFLUENCE

Not all participation is equal. Systems could start valuing contribution—not just capital.

ARTIST-LED STRUCTURES

If artists create the culture, they should have a stronger voice in shaping it.

HYBRID CURATION MODELS

Pure voting isn't enough. Pure curation isn't either. A mix might actually work better.

TRANSPARENCY IN PROMOTION

Clearer signals around why artists are featured, backed, or elevated.

AWARENESS OF COORDINATION DYNAMICS

You can't eliminate them—but you can design systems that don't blindly reward them.

None of this is easy. But pretending the current model is fully "decentralized" doesn't help either.

A MORE HONEST VERSION OF WEB3

Web3 didn't eliminate power.

It exposed it.

And that's not a failure—it's an opportunity.

Because now the question isn't whether governance exists.

It's whether we're willing to look at it clearly, admit where it falls short, and design something better from there.

Not perfect.
Not pure.

But at least… honest.

THE PUNK

OPERATING SYSTEM

~~BUILT WITHOUT PERMISSION~~



By Steve Coulter aka 45renegade

INTRODUCTION: NOT A SCENE, A SYSTEM

PUNK IS USUALLY REMEMBERED AS A MOMENT. A SOUND. A LOOK. A BRIEF RUPTURE IN THE LATE TWENTIETH CENTURY THAT BURNED BRIGHTLY AND THEN RECEDED INTO NOSTALGIA.

THAT READING IS CONVENIENT, AND INCOMPLETE.

WHAT PUNK ACTUALLY PRODUCED WAS NOT JUST MUSIC OR ATTITUDE, BUT A WORKING METHOD. A WAY OF ORGANISING PRODUCTION, DISTRIBUTION, ATTENTION, AND VALUE UNDER CONDITIONS OF EXCLUSION AND SCARCITY. LONG BEFORE THE LANGUAGE OF LEAN STARTUPS, GUERRILLA MARKETING, BOOTSTRAPPING, OR CREATOR ECONOMIES EXISTED, PUNK WAS ALREADY PRACTISING THEM, NOT AS THEORY, BUT AS NECESSITY.

THIS BOOK BEGINS FROM A SIMPLE PROPOSITION.

PUNK SHOULD BE UNDERSTOOD LESS AS A CULTURAL MOVEMENT AND MORE AS AN OPERATING LOGIC.

WHY PUNK STILL MATTERS

EVERY FEW DECADES, THE SAME CLAIMS ARE MADE WITH GREAT CONFIDENCE. THAT INSTITUTIONS ARE TOO SLOW. THAT GATEKEEPERS ARE OUT OF TOUCH. THAT NEW TOOLS HAVE LEVELLED THE PLAYING FIELD. THAT INDIVIDUALS CAN NOW OPERATE INDEPENDENTLY, DIRECTLY, AND CHEAPLY.

EACH TIME, THESE CLAIMS ARE PRESENTED AS UNPRECEDENTED.

THEY ARE NOT.

PUNK EMERGED UNDER REMARKABLY SIMILAR CONDITIONS. ACCESS TO MAINSTREAM CHANNELS WAS LIMITED. CAPITAL WAS SCARCE. EXPERTISE WAS HOARDED. THE RESPONSE WAS NOT TO WAIT FOR PERMISSION, BUT TO BUILD PARALLEL SYSTEMS THAT WERE SMALLER, ROUGHER, AND IMMEDIATELY USABLE.

RECORDS WERE PRESSED IN SMALL RUNS. SHOWS WERE BOOKED IN BORROWED ROOMS. MAGAZINES WERE PHOTOCOPIED AND STAPLED. PROMOTION WAS DONE HAND TO HAND.

NONE OF THIS WAS ROMANTIC AT THE TIME. IT WAS SIMPLY PRACTICAL.



FROM CULTURE TO STRUCTURE

OVER TIME, PUNK'S SURFACE ELEMENTS BECAME EASIER TO REMEMBER THAN ITS UNDERLYING MECHANICS. SAFETY PINS, DISTORTION, CONFRONTATIONAL GRAPHICS. THESE TRAVELLED WELL AS SYMBOLS.

THE STRUCTURE DID NOT.

YET IT IS THE STRUCTURE THAT KEEPS REAPPEARING. IN INDEPENDENT PUBLISHING. IN OPEN SOURCE SOFTWARE. IN EARLY BLOGGING. IN BOOTSTRAPPED BUSINESSES. IN CREATOR-LED ENTERPRISES THAT RESIST INTERMEDIARIES.

WHENEVER FORMAL SYSTEMS BECOME TOO RIGID OR EXPENSIVE, INFORMAL ONES ARISE. WHENEVER ACCESS NARROWS PEOPLE ROUTE AROUND IT. PUNK WAS AN EARLY, VISIBLE EXAMPLE OF THIS BEHAVIOUR, NOT ITS ORIGIN.

WHAT THIS BOOK IS DOING



THIS IS NOT A HISTORY OF PUNK MUSIC. OTHERS HAVE DONE THAT BETTER, AND IN GREATER DETAIL.

NOR IS THIS A HANDBOOK. THERE ARE NO STEPS TO FOLLOW, NO PROMISES OF SUCCESS, NO CLAIMS OF INEVITABILITY.

INSTEAD, THIS BOOK TRACES A PATTERN.



HOW INDEPENDENCE WAS CONSTRUCTED.
HOW ATTENTION WAS CAPTURED WITHOUT CAPITAL.
HOW COMMUNITIES REPLACED AUDIENCES.
HOW ETHICS ACTED AS LIMITS RATHER THAN SLOGANS.
HOW REFUSAL FUNCTIONED AS STRATEGY.

EACH CHAPTER EXAMINES ONE COMPONENT OF THIS SYSTEM, NOT TO ELEVATE IT, BUT TO UNDERSTAND HOW IT WORKED AND WHY IT CONTINUES TO RESURFACE.

WHY BUSINESS IS THE RIGHT LENS

IT MAY SEEM PERVERSE TO DISCUSS PUNK IN BUSINESS TERMS. PUNK WAS, AFTER ALL, SUSPICIOUS OF COMMERCE, PARTICULARLY WHEN IT BECAME EXTRACTIVE OR EXCLUSIONARY.

BUT PUNK WAS NEVER NAÏVE ABOUT ECONOMICS. IT UNDERSTOOD COSTS, TRADE-OFFS, MARGINS, AND SUSTAINABILITY INTIMATELY. IT SIMPLY CHOSE TO ORGANISE THEM DIFFERENTLY.



SEEN THIS WAY, PUNK IS NOT ANTI-BUSINESS. IT IS ANTI-ABSTRACTION.

MONEY WAS COUNTED BECAUSE IT MATTERED.
LABOUR WAS VISIBLE BECAUSE IT WAS SHARED.
GROWTH WAS QUESTIONED BECAUSE IT CHANGED INCENTIVES.

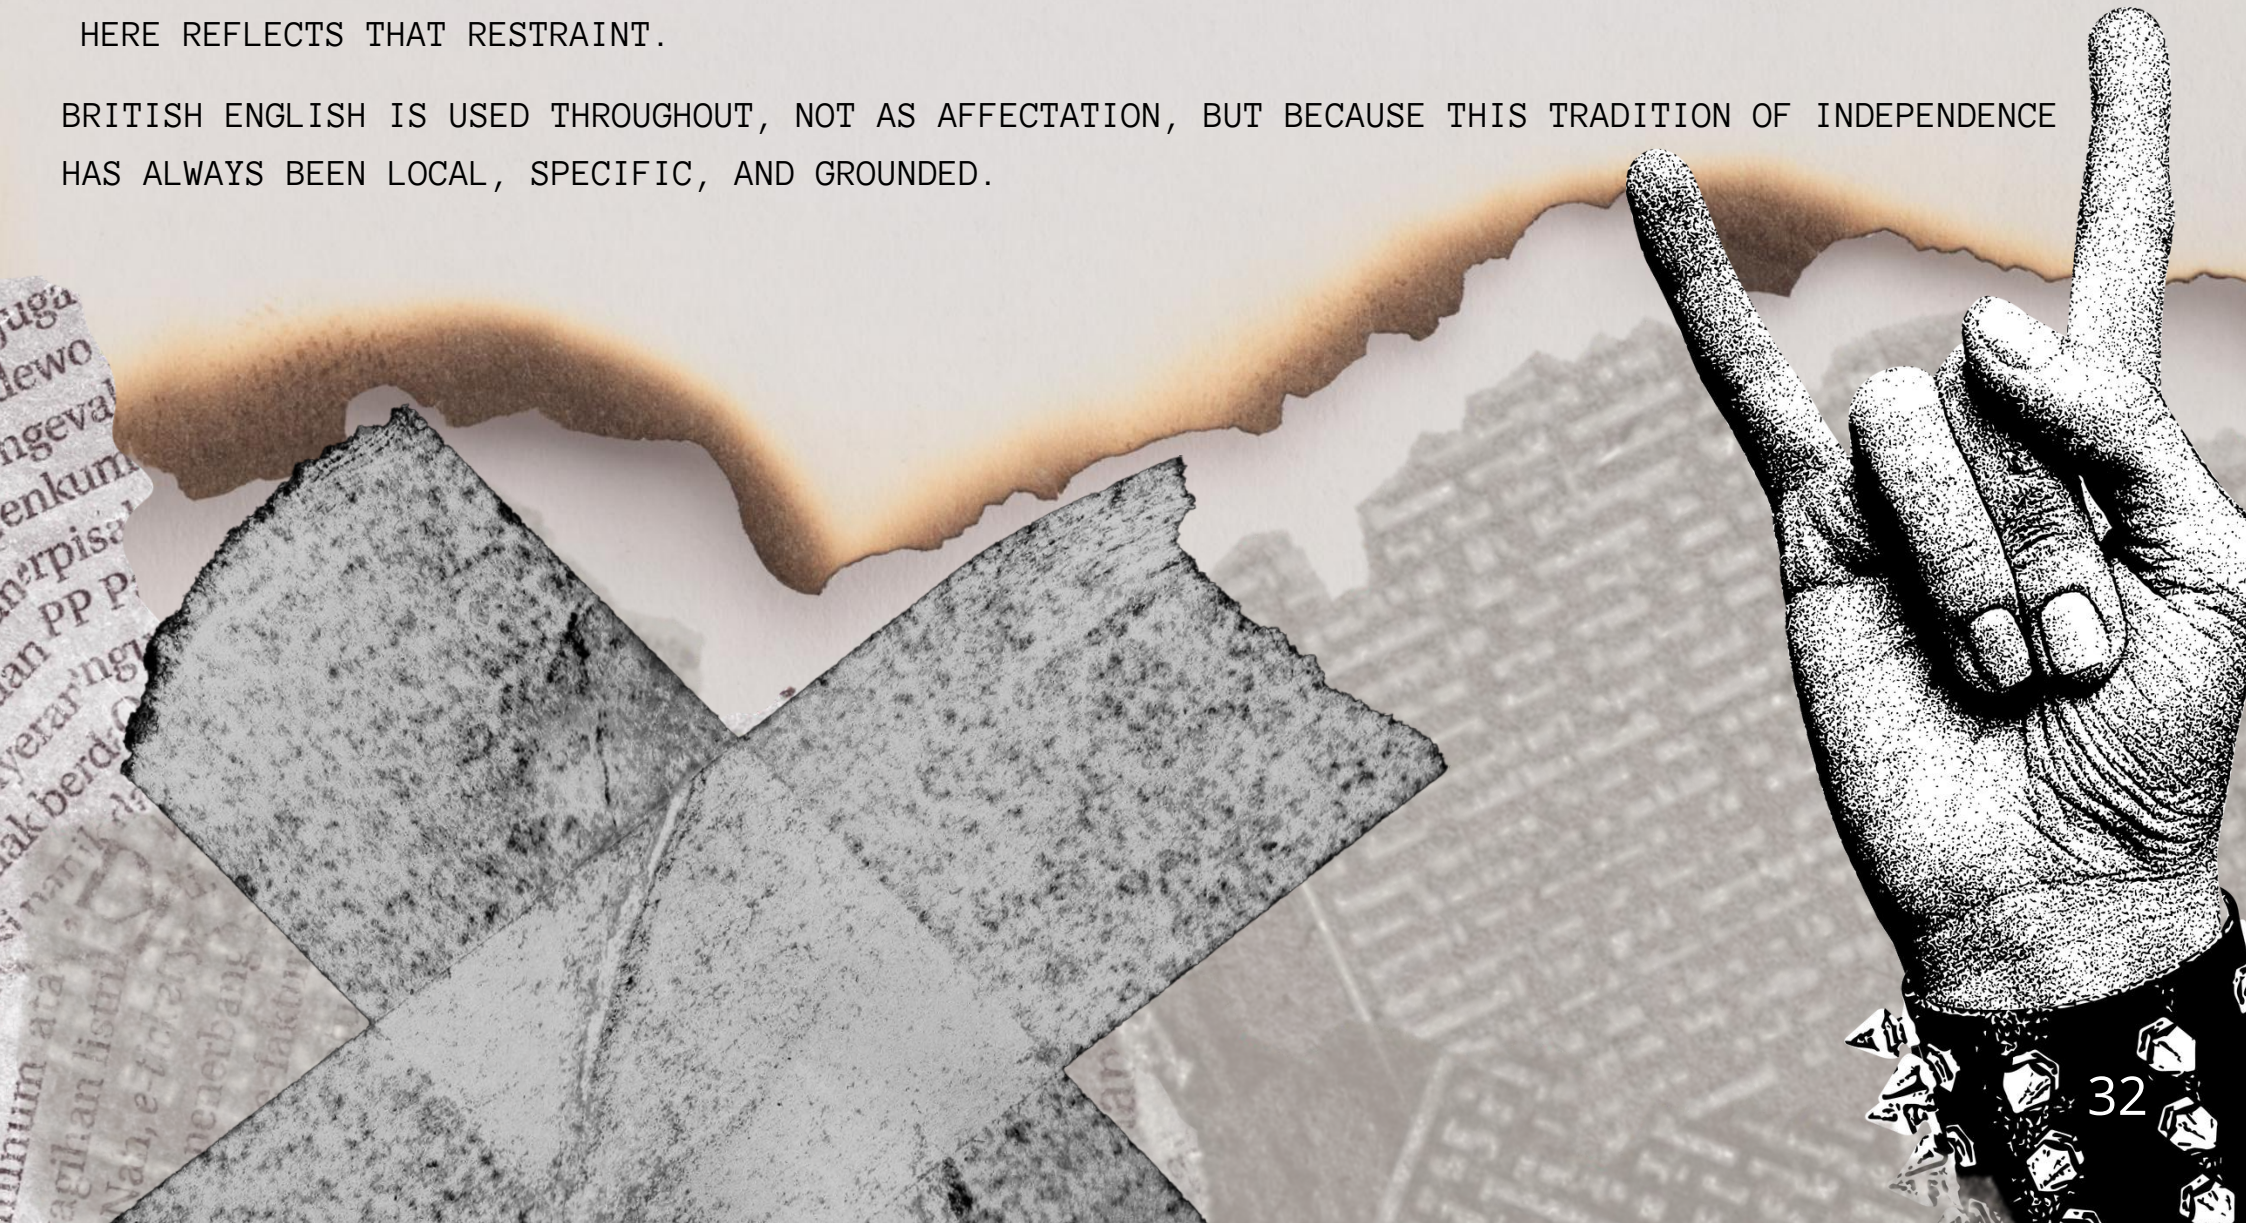
THESE ARE BUSINESS CONCERNS, WHETHER ACKNOWLEDGED OR NOT.

A NOTE ON SCOPE AND TONE

THE EXAMPLES IN THIS BOOK RANGE ACROSS MUSIC, PUBLISHING, MARKETING, SOFTWARE, AND SMALL ENTERPRISE. THEY ARE USED NOT AS INSPIRATION, BUT AS EVIDENCE.

THE TONE IS INTENTIONALLY PLAIN. PUNK DID NOT RELY ON GRAND LANGUAGE. IT RELIED ON ACTION. THE VOCABULARY HERE REFLECTS THAT RESTRAINT.

BRITISH ENGLISH IS USED THROUGHOUT, NOT AS AFFECTATION, BUT BECAUSE THIS TRADITION OF INDEPENDENCE HAS ALWAYS BEEN LOCAL, SPECIFIC, AND GROUNDED.



WHAT FOLLOWS

THE CHAPTERS THAT FOLLOW BUILD AN ARGUMENT PIECE BY PIECE.

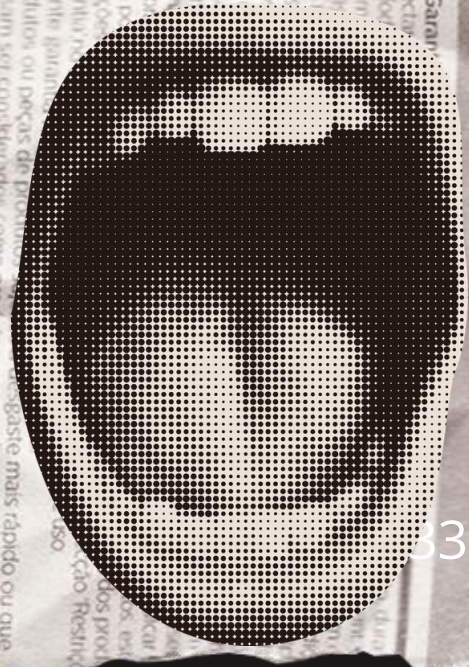
THEY BEGIN WITH SELF-SUFFICIENCY AND MOVE THROUGH ATTENTION, COMMUNITY, PRODUCTION, TOOLS, EXCHANGE, AND ETHICS. THEY END NOT WITH INSTRUCTION, BUT WITH ORIENTATION.

THE AIM IS NOT TO PERSUADE ANYONE TO CHOOSE INDEPENDENCE.
IT IS TO MAKE CLEAR WHAT THAT CHOICE ENTAILS.

PUNK DOES NOT OFFER SAFETY.
IT DOES NOT OFFER SCALE.
IT DOES NOT OFFER COMFORT.

WHAT IT OFFERS IS COHERENCE.

AND IN PERIODS WHERE SYSTEMS FEEL BLOATED, DISTANT, OR BRITTLE, THAT COHERENCE HAS A HABIT OF BECOMING USEFUL AGAIN.



INSPIRATION

THE BOOK BEGAN AS AN ATTEMPT TO EXPLAIN A PATTERN I KEPT ENCOUNTERING ACROSS DIFFERENT INDUSTRIES, DECADES, AND TECHNOLOGIES.

> **LONG BEFORE THE LANGUAGE OF LEAN START-UPS, BOOTSTRAPPING, OR THE CREATOR ECONOMY, THE SEVENTIES PUNK ETHIC HAD ALREADY SOLVED MANY OF THE SAME PROBLEMS. HOW TO ACT WITHOUT PERMISSION. HOW TO OPERATE WITH LIMITED RESOURCES. HOW TO BUILD AUDIENCES INTO COMMUNITIES. HOW TO GROW WITHOUT LOSING COHERENCE, AND CRUCIALLY, HOW TO STOP.**

RATHER THAN TREATING PUNK AS A CULTURAL MOMENT OR AESTHETIC, THIS PROJECT EXAMINES IT AS A WORKING SYSTEM. A SET OF PRACTICAL DECISIONS ABOUT PRODUCTION, ATTENTION, EXCHANGE, AND ETHICS MADE UNDER CONDITIONS OF EXCLUSION AND CONSTRAINT.

THE AIM IS NOT NOSTALGIA, INSTRUCTION, OR REVIVAL. IT IS ORIENTATION. TO UNDERSTAND HOW INDEPENDENCE IS CONSTRUCTED, WHAT IT COSTS, AND WHY THESE SAME STRUCTURES CONTINUE TO REAPPEAR WHENEVER FORMAL SYSTEMS BECOME RIGID OR INACCESSIBLE.

THE AUTOR

STEVE COULTER AKA 45RENEGADE IS A CREATIVE, ARTIST AND WRITER FROM THE UNITED KINGDOM. AN ADVOCATE OF CREATIVE INDEPENDENCE AND DEREGULATION HE PRODUCES WORK

THAT RESONATES WITH INDIVIDUALS AND BUSINESSES WHO SEEK AUTHENTICITY, EDGE, AND A REFUSAL TO CONFORM TO TRADITIONAL BOUNDARIES.

45RENEGADE.CO.UK



GOVERNING THE SOUND:

HOW OPTIMUS GODDESS IS REWRITING MUSIC INDUSTRY POWER

BY VESSY MINK



In the advertising world it is often impossible to find the seeds of change. Due to being inundated with so much content, Optimus Goddess offers top tier music creators, (whose songs change the air in the room) the grand opportunity to be featured in a growing catalogue of music in which the artist retains full publishing & writing share of their music creations.

Getting that music licensed to tv film and the internet is the challenge of Optimus Goddess, a superhero of sound who lives to bring the best songs to life for the benefit of the artists who created them.

Songs that are written & produced by artists for the purpose of being synced by advertising agencies that see the value of working with such music libraries that create quantifiable income streams for them is a world where Optimus Goddess reigns supreme.

This new model of music sync licensing forges partnerships built on equitable relationships that pay music artists the lion share. The quantifiable aspect of data sets required to prove that a song can save the world, one artist at a time, live in real stories of overcoming adversity through and by the creation of music that lives and breathes through them.

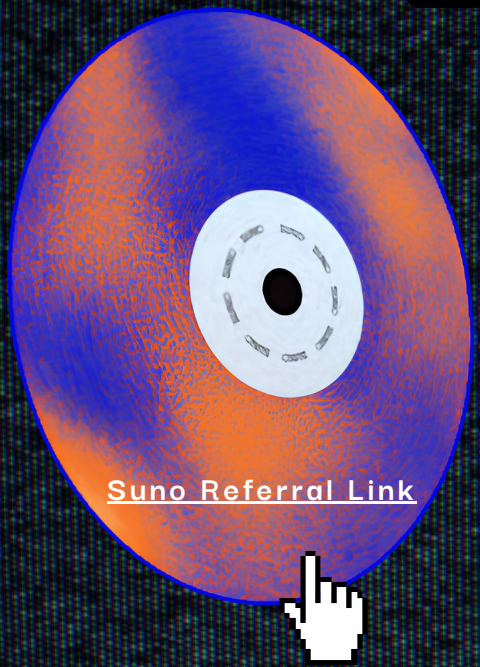
With the purpose of not only supporting the work of musicians but finding the puzzle pieces to help the world through their work, Optimus Goddess finds a way to do it by fostering new partnerships with investors and companies who realize that nothing is set in stone and that the world belongs to the ones who fight for justice.

It's the very world that was created to sustain itself that spins and feeds and keeps on giving and would be nothing without the great sun shining upon it. Optimus Goddess, healing the world one song at a time.

OPTIMUS GODDESS DEBUT RELEASE

10 brand new original songs (witten organically) & produced via SUNO AI. 20 tracks in total make up this debut album titled OPTIMUS GODDESS (including the raw un edited original demos by Vessy Mink)



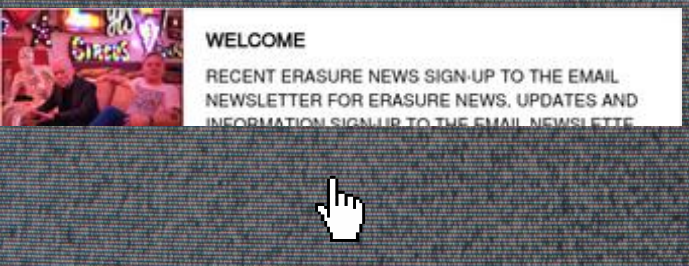


MEMORY LANE STILL ALIVE AND KICKING

Erasure has been one of the most important and joyful parts of my musical life for as long as I can remember. Vince Clarke's sharp, addictive synth hooks paired with Andy Bell's soaring, soulful voice created a sound that was pure emotion wrapped in irresistible pop – heartfelt, danceable, and completely timeless.

I've been greatly inspired and influenced by Erasure, and I had the great opportunity to experience them live in Chicago in the mid-1990s. That night still stands out as one of my favorite concert memories – the energy, the melodies, and the connection with the crowd were electric.

Check them out here on Apple Music, nd visit their official website for the latest news, updates, and more:



Music

if you love 80s/90s synth-pop with real heart, Erasure is essential. Highly recommended!



Tom Friedman "Takeaway" installed in 2024 at CTA's UIC-Halsted Blue Line station.

It shows the Greek god Hermes (the swift messenger) as a running figure balancing a tall stack of restaurant take-out containers on his head—a playful nod to modern urban life, delivery culture, and hustle. The stainless steel piece is visible from the Eisenhower Expressway (I-290) near Peoria Street/UIC. Friedman is a UIC alumnus and well-known contemporary artist. This was a collaboration between the CTA and the University of Illinois Chicago for their public art collection.

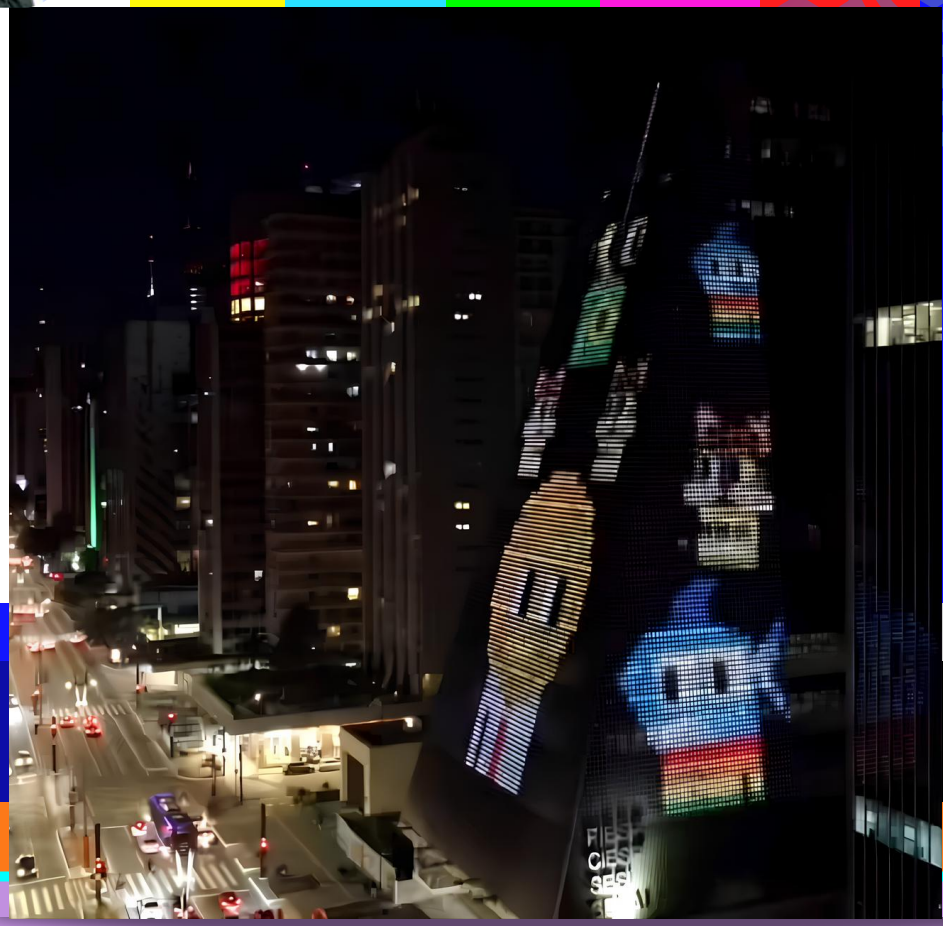
TokenNation and the Governance of the Digital Future

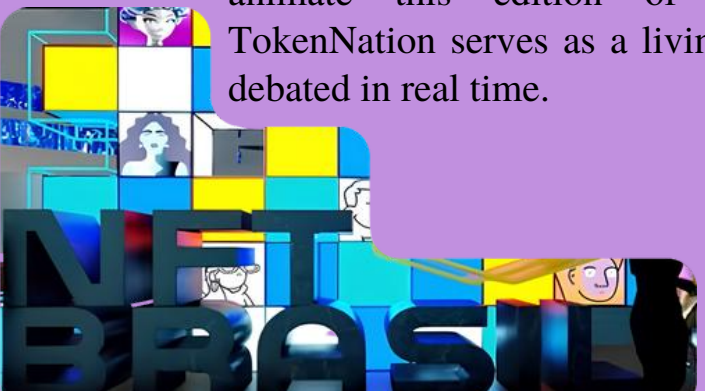
BY MA. VITOR EMANUEL GRIPP
COMMUNITY MANAGER, TOKEN NATION





I have followed Token Nation since its very first edition, initially as a speaker and artist, exhibiting work at a moment when the event was still finding its identity, and now, in 2026, as part of the team itself, serving as Community Manager in what is already the fourth consecutive edition at the Bienal in São Paulo. This trajectory, from participant to builder, mirrors in many ways the evolution of the event itself: what began as a gathering centred on NFTs has matured into one of Latin America's leading ecosystems for debating artificial intelligence, blockchain, tokenisation, and digital innovation. From my point inside the organisation, what strikes me most is not the scale of the audience or the prominence of the speakers, but the quality of the questions the event forces us to ask: Who decides how new technologies are governed? How are the rules of decentralised systems written, enforced, and contested? And, perhaps most pressingly, how can the rigour of academic inquiry inform, and be informed by, the urgency of technological practice? These are the questions that animate this edition of Virtualia Magazine, and TokenNation serves as a living laboratory where they are debated in real time.





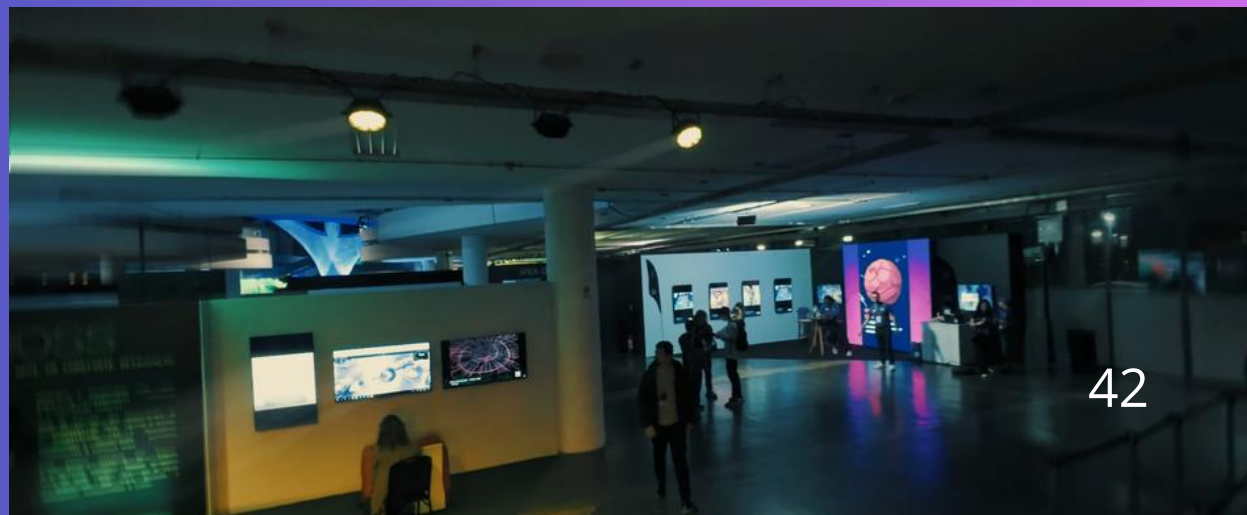


Token Nation's history begins in 2022, when Marco Affonseca — drawing on his background in entertainment and large-scale events — decided to organise a major gathering on NFTs in São Paulo. At the time, Brazil's principal events on the subject were concentrated in Rio de Janeiro, and Affonseca identified an opportunity to create a dedicated space for the rapidly growing São Paulo community. The first two editions, held at the iconic Bienal pavilion, were called "NFT Brasil", a name inspired by similar conferences in New York, Paris, and Bogotá. The format was ambitious from the outset, but the scope of the conversation was still largely defined by digital art and collectables.

Each subsequent edition expanded both the audience and the intellectual terrain. The organisers recognised that discussing tokens had come to require engagement with the philosophy of technology, the foundations of computer science, the models of behavioural economics, and the intricacies of regulatory policy. Tokenisation of real-world assets, new blockchain-enabled business models, and the accelerating role of artificial intelligence all demanded a broader stage. It was in 2025 that the event underwent its defining strategic rebranding, becoming TokenNation, a more robust and comprehensive platform capable of bridging business, art, culture, and technology. Supported by figures such as skateboarder and co-founder Bob Burnquist, the new identity signalled a consolidation of the event's ambition: to function as a global hub connecting innovation with new generations.







The 2025

edition immediately demonstrated the strength of this new proposal. Over two days, the Bienal received approximately 10,000 visitors per day, confirming Token Nation's position as a meeting point between the traditional market and the emerging digital economy. The programme was structured around two complementary areas: "Experience", free and focused on educational and immersive content, and "Onchain", oriented toward business, featuring talks, a hackathon, and case studies presented by institutions such as Itaú Unibanco, Microsoft, and Unilever. Artificial intelligence gained its own dedicated stage, the "AI Arena", while Drex, the Central Bank's digital currency, emerged as one of the edition's central themes, signalling the event's direct dialogue with the structural transformation of the Brazilian financial system.









In 2026

TokenNation arrives at its fourth consecutive edition at the Bienal, and the evolution is visible not only in scale but in depth, particularly in the event's growing relationship with higher education. One of the developments I find most significant is the expansion of our university partnerships, which have transformed TokenNation from a market event into a genuine bridge between technological practice and academic knowledge. This year, the event hosts intensive bootcamps at FIAP Paulista and Inteli Blockchain, two of Brazil's most innovative institutions in the fields of technology and entrepreneurship. These represent a commitment to training the next generation of practitioners who can think rigorously about the technologies they are building. Alongside these partnerships, TokenNation has deepened its connections with major Brazilian research universities, creating a dialogue between the market and the academy that is rare in the global conference landscape.

The 2026 programme is structured around four main tracks. The Token Nation stage concentrates the central debates on the digital economy. The Artificial Intelligence stage is devoted to the most disruptive technology of our time. HackaNation provides a dedicated space for those beginning to build practical solutions. And the OnChain stage hosts discussions for those with deeper web3 expertise. The four-track architecture reflects a deliberate curatorial philosophy: that conversations about Web3 and AI must simultaneously address the novice and the expert, the developer and the policymaker, the builder and the critic.

What makes TokenNation particularly relevant for a magazine devoted to governance is the way its structure embodies, rather than merely describes, the tensions that preoccupy theorists of collective decision-making. The choice of speakers, the selection of featured projects, the designation of which protocols receive a platform, all of these are curatorial acts that mirror the governance challenges discussed by Fer Caggiano and VAN Ameneyro elsewhere in this issue. The event practises a form of distributed curation in which multiple stakeholders — corporate sponsors, community builders, academic partners, and independent creators — collectively shape the final programme. This is governance by emergence rather than by decree.



Dr. Rodrigo Cid
tokennation 2025 edition

The 2026 edition gives particular expression to this commitment through an interdisciplinary selection of academic voices. Dr. Maria Gorete, from Fiocruz, brings a foundational perspective on the social and ethical impacts of digitalisation and tokenisation in the field of health and innovation. Dr. Carlos Frederico, head of KryptoLab and full professor at the Federal University of Ouro Preto (UFOP), addresses the technical pillars and industrial applications of blockchain and cryptography. And Dr. Rodrigo Cid, professor of Philosophy of Technology at UFOP and coordinator of GIFLABS — a laboratory for experimentation at the intersections of philosophy, art, and technology — brings to the stage the deepest reflections on the nature, ethics, and governance of artificial intelligence and decentralised infrastructures. Their presence is not incidental, it signals that TokenNation takes seriously the need to ground technological innovation in conceptual clarity, precisely the kind of clarity that philosophical inquiry provides.

Token Nation is, therefore, considerably more than a business fair or a technology conference. It is a space where the future of digital society is actively negotiated. Where academics, creators, professionals, founders, investors, and protocol developers meet not only to present solutions but to debate the philosophical, ethical, and political problems that these technologies bring to the surface. For a magazine focused on governance, TokenNation offers an invaluable case study: a large-scale, multi-stakeholder event that must constantly negotiate the tension between efficiency and participation, between expert curation and community voice, between the immutability of code and the flexibility of human judgment. From where I stand, having helped build this community over the years, what is most striking is that these are not abstract problems. They are concrete choices made every day by the organisers, speakers, university partners, and attendees of TokenNation. And it is, finally, proof that Brazil has a central role to play in building global thought about the world to come.



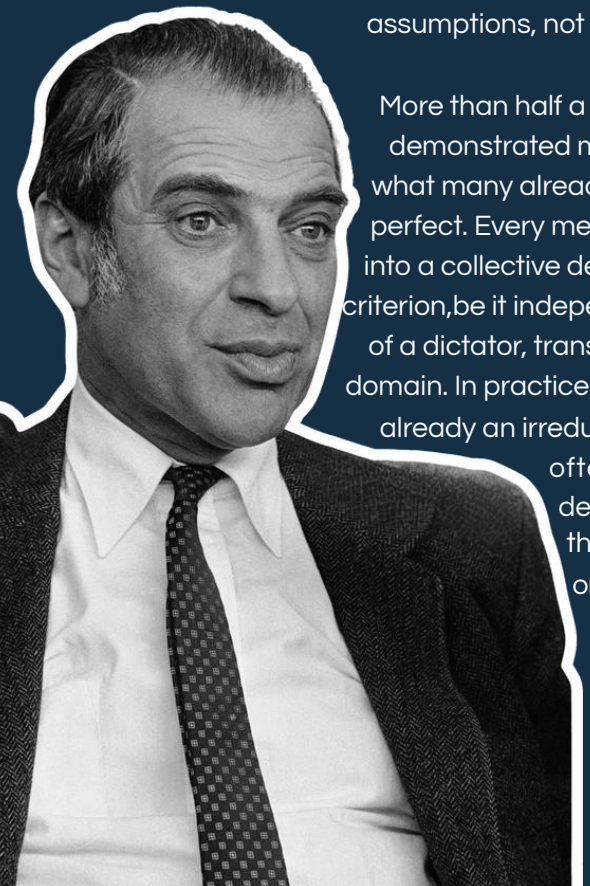
VOTING OR GOVERNING:

SOME PHILOSOPHICAL PROBLEMS NOT SOLVED BY BLOCKCHAIN



BY PROF. DR. RODRIGO CID, FEDERAL UNIVERSITY OF OURO PRETO

The promise of digital governance, spread by blockchain enthusiasts and decentralized autonomous organizations, was one of a quiet and clean revolution. Finally we could replace messy politics, fallible intermediaries, and opaque hierarchies with something supposedly objective, automatic, and transparent: code. Decisions would be made through direct voting recorded on immutable ledgers, executed by smart contracts without delays or creative interpretations. Goodbye presidents, bureaucrats, lobbyists, and arbitrary curators. Welcome to the age of algorithmic governance. However, this promise runs into a fundamental problem that political philosophy and social choice theory had already identified decades before Bitcoin was invented: voting does not resolve politics. It merely organizes it. And organizing is not nothing, but it is far less than its enthusiasts imagine. The aggregation of collective preferences through voting rules is a normative task loaded with ethical and political assumptions, not a mere matter of computational efficiency.



More than half a century ago, the economist Kenneth Arrow demonstrated mathematically, in his impossibility theorem, what many already suspected intuitively: no voting system is perfect. Every method of aggregating individual preferences into a collective decision fails at least one intuitively desirable criterion,be it independence of irrelevant alternatives, absence of a dictator, transitivity of social choice, or universality of the domain. In practice, this means that choosing the voting rule is already an irreducible political decision. And this decision is often made behind the scenes by whoever designed the system or holds system or holds the initial power to define the protoco system or holds the initial power to define the protocol.



When a decentralized autonomous organization adopts the model of one token equals one vote, it signals that economic power should translate directly into political power, with no protection mechanisms for minorities that lack capital. When it adopts quadratic voting, it prioritizes the intensity of preference over the mere quantity of tokens, a choice that favors voters with strong preferences over indifferent majorities. When it requires a seventy percent quorum, it protects the minority against hasty decisions, but also paralyzes the ability to act quickly. There is no neutrality. The voting protocol, in its most elementary form, is already a constitution, a set of norms defining who counts, how they count, and what counts as a legitimate decision. Ignoring this normative dimension is to condemn oneself to reproducing, under the guise of technical innovation, the same power asymmetries that digital governance promised to overcome.

One of the most acute philosophical problems emerging in this scenario is that of fragmented responsibility, often called the problem of many hands. In digital governance, the authorship of decisions is distributed and opaque. A smart contract is written by one developer, audited by another professional, deployed by a technical team, governed by an assembly of token holders who have never met physically, and executed by an autonomous machine with no interpretive discretion. When the system produces a harmful or unjust outcome, who is held responsible? The developer can argue that they merely followed the specifications. The auditor can claim the flaw was in the economic assumptions, not in the code. The token holder assemly can hide behind the crowd, each with a minuscule share of power. And the machine, obviously, cannot be punished nor feel remorse.



This epistemic and normative fragmentation is not a fixable bug; it is a structural feature of decentralized systems that lack a single point for the imputation of responsibility. While in traditional political systems responsibility falls on named individuals or collegiate bodies with legal identity, in decentralized autonomous organizations responsibility dissolves into a haze of circular interdependencies, where no one had enough control to be blamed, yet everyone contributed causally to the outcome. The absence of a responsible agent does not make the harm less real; it only makes its reparation more difficult and its prevention more speculative.

Even before any voting occurs, however, there is a layer of governance that is often invisible and underestimated: the governance of visibility. Fer Caggiano warns us of a phenomenon that digital art platforms and decentralized marketplaces are reluctant to admit. Curation did not disappear in web3. It merely became less visible and more concentrated in mechanisms of reputation and capital. Whoever decides which artists appear on a marketplace homepage, which works are highlighted by algorithmic curation, and which collections receive official verification wields as much power as the curators of traditional galleries once did. The difference is that in the traditional system, the curator was an identifiable and criticizable figure.

In digital governance, the power of visibility is often exercised by opaque systems, by informal committees, or simply by the fact that wealthy collectors coordinate their tastes through private messaging groups. The consequence is that before any token holder votes on a proposal, the field of alternatives has already been drastically reduced by whoever decides what deserves attention. Governance of attention, therefore, is the original governance, the one that precedes and conditions all others. If an artist is not seen, they cannot be voted on. If a proposal does not circulate through the right channels, it will not even be considered. The promise of direct democracy on the blockchain thus runs into the harsh reality of the attention economy, where scarcity lies not in votes but in eyes.

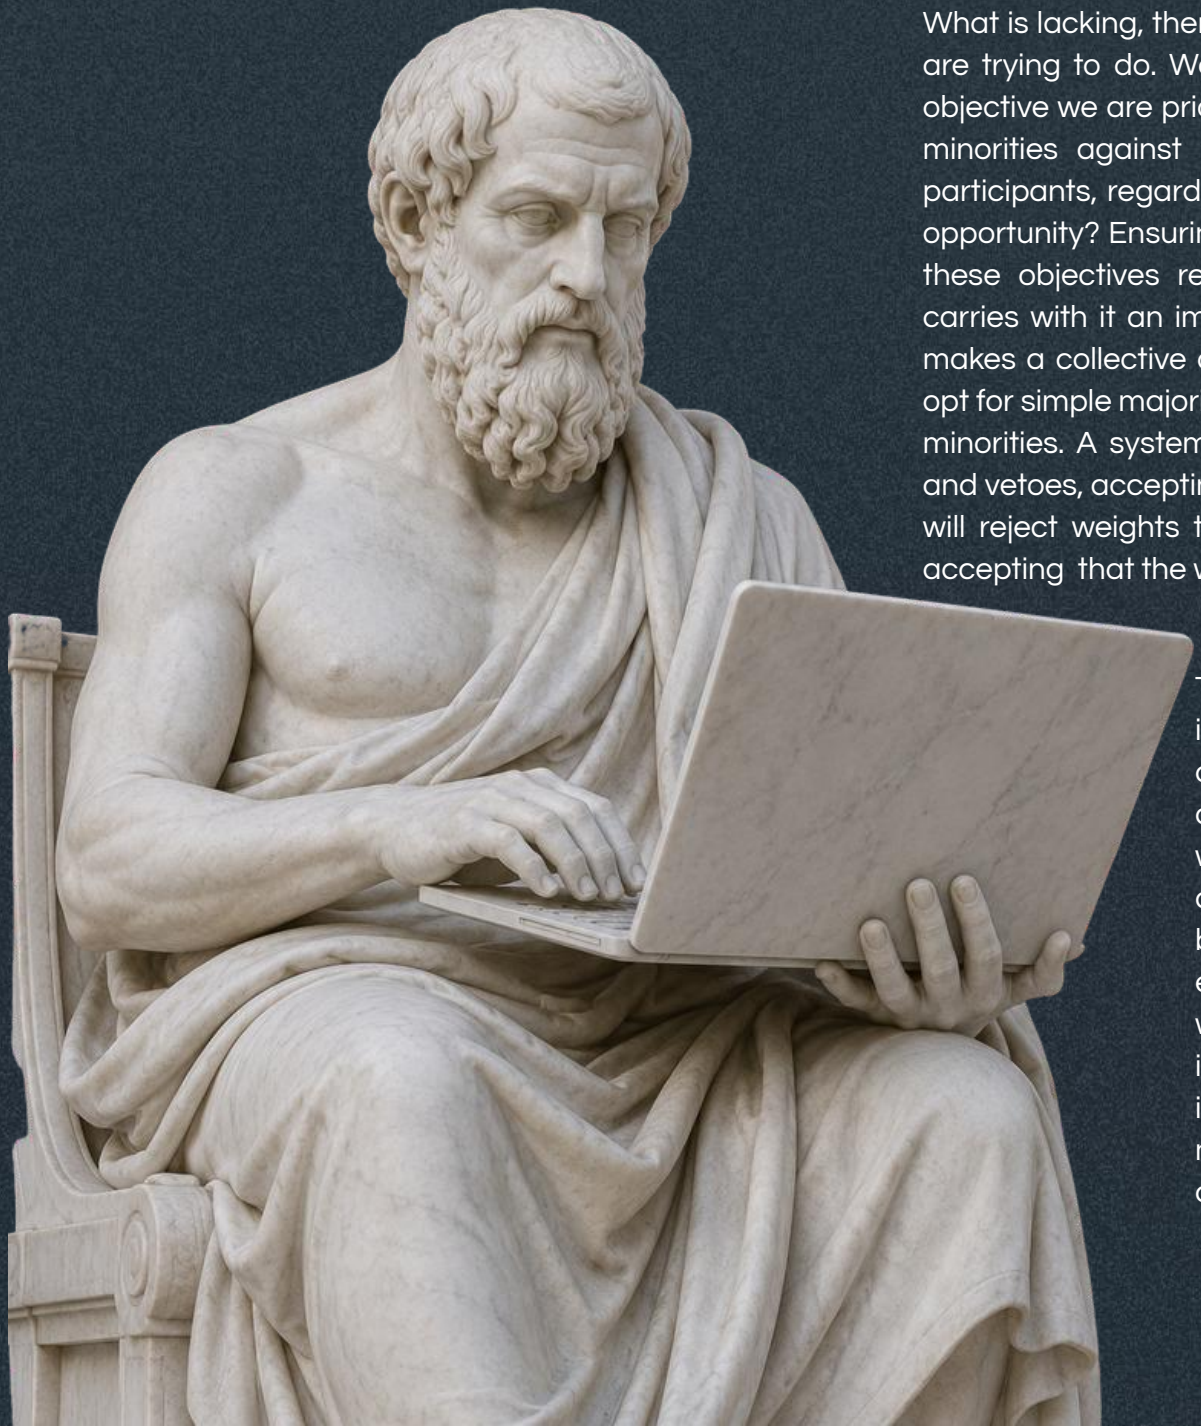
To this must be added the problem of institutional memory, a theme that VAN Ameneyro explored in their dialogue with One Love DAO. In digital ecosystems where platforms are born and die in cycles of a few months, where smart contracts are abandoned or replaced, where discussion forums are lost in inaccessible archives, what happens to the archive? Who ensures that the decisions made today, the justifications presented, the minorities that opposed them, and the commitments made will not simply be forgotten tomorrow? Digital governance, in its most radical version, tends toward the eternal present. Each vote is an isolated event, with no mandatory memory embedded.



The transparency promised by the blockchain is a transparency of raw records, not of interpreted meanings. Knowing that a cryptographic address voted in favor of a proposal does not tell us why they voted that way, what arguments convinced them, or whether they will change their mind in light of new evidence. The absence of a deliberative public sphere, with time for reflection and exposure to counterarguments, turns voting into a mechanical act, not an exercise in informed judgment. And without institutional memory, each generation of participants is forced to relearn lessons already paid for dearly by predecessors, in a perpetual cycle of amnesia and repetition of mistakes.

Faced with this diagnosis, it would be easy to succumb to skepticism and conclude that digital governance is an illusion or a regression. But there is a more productive and less pessimistic reading. What these problems reveal is not the unfeasibility of digital governance, but the necessity of treating it as a problem of institutional design, not as a technological destiny. Voting systems can be designed, tested empirically, iterated, and even abandoned when they prove inadequate. There is a diverse toolkit available that goes far beyond the simplistic one token one vote. There is delegated voting, which allows participants without time or expertise to entrust their votes to trusted representatives, combining scale with specialization.

There is quadratic voting, which mitigates plutocratic power by making the cost of additional votes progressively more expensive, balancing intensity of preference with purchasing power. There is approval voting, where each participant can approve as many alternatives as they wish, eliminating the problem of strategic voting and artificial polarization. There is the random selection of deliberative assemblies, which reintroduces uncertainty and statistical representativeness where purely voluntary mechanisms fail. And there are hybrid systems that combine on-chain voting for routine decisions with off-chain deliberation for constitutional matters, recognizing that not every decision deserves the same degree of automation.



What is lacking, therefore, is not technology. It is philosophical clarity about what we are trying to do. We need to know, before programming a smart contract, which objective we are prioritizing. Maximizing efficiency in resource allocation? Protecting minorities against occasional majorities? Distributing power equally among all participants, regardless of their wealth? Accelerating decisions to seize windows of opportunity? Ensuring that all relevant voices are heard, even the weakest? Each of these objectives requires a different voting mechanism, and each mechanism carries with it an implicit theory of justice, a silent answer to the question of what makes a collective decision legitimate. A system that prioritizes efficiency will likely opt for simple majority rules and automatic execution, accepting the risk of trampling minorities. A system that prioritizes minority protection will require supermajorities and vetoes, accepting the risk of paralysis. A system that prioritizes equality of power will reject weights tied to tokens, seeking one participant one vote mechanisms, accepting that the wealthy may simply create multiple identities to bypass the rule.

The challenge of digital governance, therefore, is not technical. It is conceptual and normative. And that is why philosophy, that old and sometimes discredited practice of asking uncomfortable questions about foundations, remains the most advanced tool we have to navigate this territory. Not because philosophy offers definitive answers, but because it trains us to distinguish between different kinds of questions and not to confuse efficiency with justice, transparency with legitimacy, participation with power. Ultimately, governing is choosing between incommensurable values, and no blockchain, however elegant its cryptography, can spare us from that choice. What it can do is make the choice more explicit, more testable, and hopefully more accountable. And that, after all, is already considerable progress.

Direct Democracy

In the Age of the Extended Mind

By Prof. Dr. Filipe Farinha, University of Saint Joseph, Macau



The promise of direct democracy has always been simple and radical: citizens should decide public questions themselves rather than periodically handing power to representatives. Its weakness has also always been simple: attention is scarce. Modern societies generate too many policy questions, budgets, trade-offs, technical standards, and regulatory decisions for ordinary citizens to study and vote on continuously. The old problem of direct democracy is therefore not only institutional; it is cognitive.

EXTENDED MIND

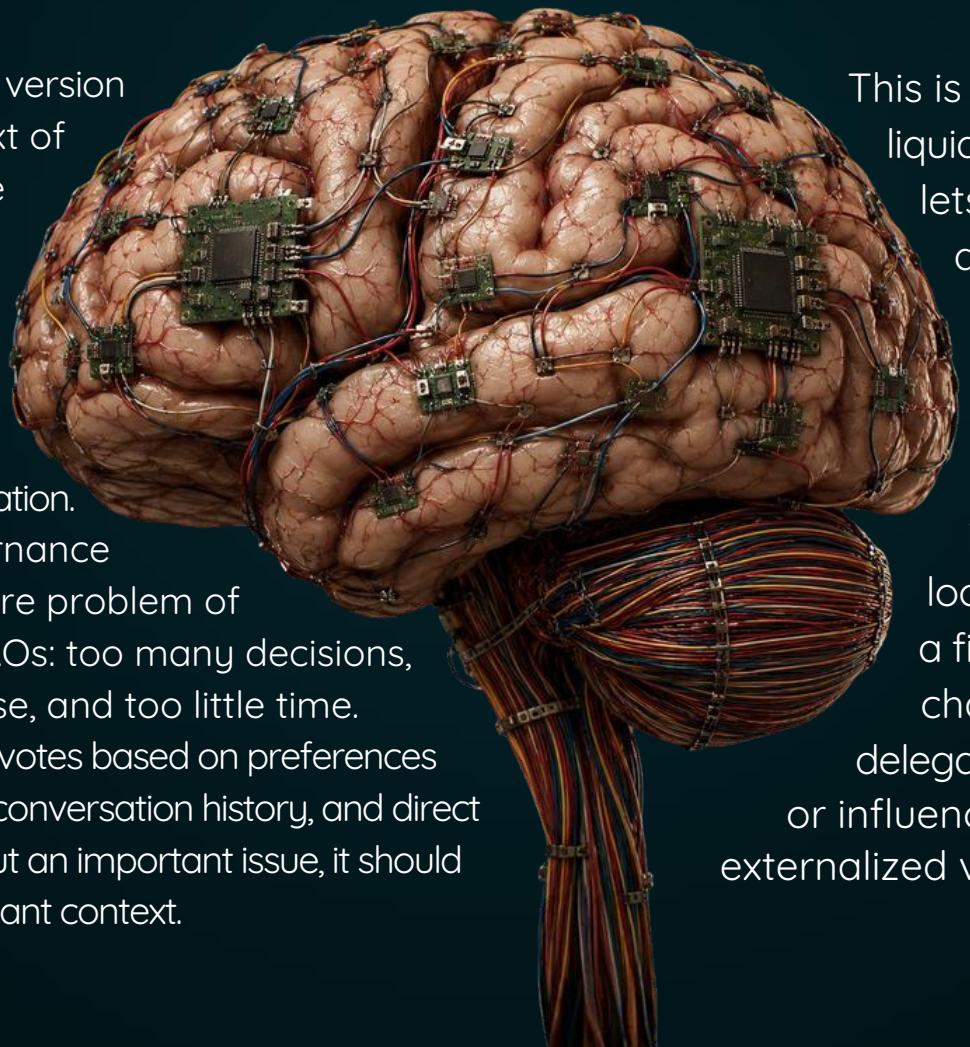


An:dy Clark and David Chalmers, original photo here

That is why the phrase personal exocortex is useful. It connects the politics of direct democracy to the philosophy of the Extended Mind. In their 1998 paper, Andy Clark and David Chalmers argued that cognition need not stop at “skin and skull”; external tools can become part of the cognitive process when they are reliably available, trusted, and integrated into a person’s practical reasoning. Their famous example is Otto, who uses a notebook as a memory system in a way functionally comparable to biological memory.

A personal exocortex updates Otto’s notebook for the age of AI. It is not merely a chatbot, a calendar, or a “digital twin.” It is an external cognitive layer: a persistent, user-aligned system that remembers, filters, explains, deliberates, and sometimes acts. Recent AI literature has already begun using “exocortex” this way, describing it as a synthetic extension of a person’s cognition and even as a swarm of AI agents that amplifies a researcher’s agency and judgment.

Applied to democracy, the personal exocortex suggests a provocative possibility: perhaps direct democracy failed at scale because citizens lacked cognitive prosthetics. If citizens had trustworthy personal agents that could read proposals, compare them to their values, ask for clarification when stakes are high, and cast routine votes under explicit authorization, then “voter apathy” might be reframed as a design failure rather than a civic defect.



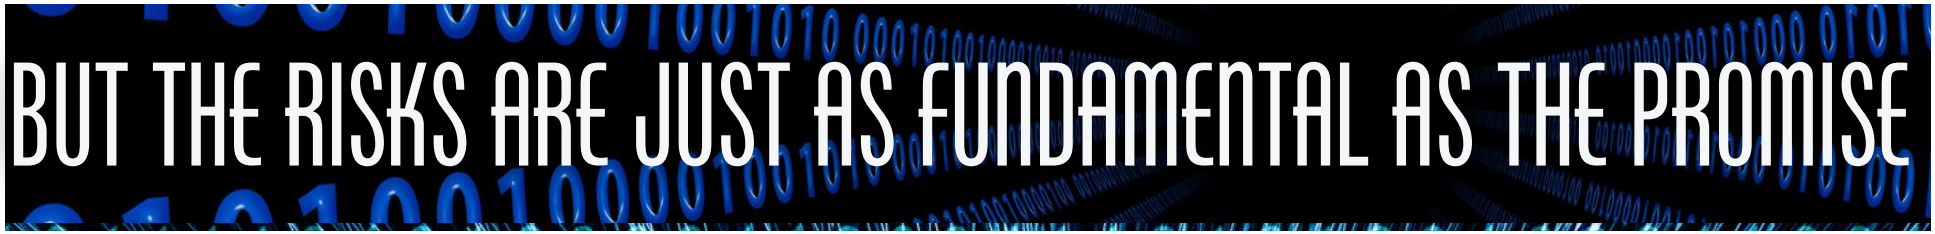
Vitalik Buterin recently made a version of this argument in the context of decentralized governance. He warned that “AI becomes the government” is dystopian, but argued that AI can empower democratic and decentralized systems when used as an aid to human participation. His proposed “personal governance agents” would address the core problem of limited human attention in DAOs: too many decisions, too many domains of expertise, and too little time. A personal agent could perform votes based on preferences inferred from a person’s writing, conversation history, and direct statements; when uncertain about an important issue, it should ask the person and provide relevant context.

This is close to, but not identical with, liquid democracy. Liquid democracy lets citizens either vote directly or delegate their vote to someone else. It is often described as a hybrid between direct and representative democracy, because voting power can flow issue by issue rather than being locked into one representative for a fixed term. The exocortex model changes the delegate. Instead of delegating to a politician, party, expert, or influencer, the citizen delegates to an externalized version of their own reasoning.

THAT SHIFT HAS REAL DEMOCRATIC APPEAL

First it could reduce the distance between public decisions and citizens' actual preferences. Direct democracy is often defended because it can increase responsiveness, transparency, and citizen empowerment. Second, exocortical voting could make participation granular. A citizen might vote personally on constitutional questions, use their exocortex for routine administrative issues, and set hard constraints on topics such as taxation, privacy, climate, or war. Third, personal agents could support deliberation rather than merely automate voting. Buterin's proposal also includes public conversation agents that summarize viewpoints, identify common ground, and convert private views into shareable formats without exposing private data.

The strongest version of this idea is therefore not "bots vote for us." It is: citizens deliberate and decide with cognitive extensions that they own, audit, and can override.



The first risk is authenticity. A vote cast by an exocortex may represent the citizen's considered judgment, or it may represent the training data, defaults, nudges, or hidden incentives of the system provider. If the exocortex is built by a platform, party, state, or advertising company, then "delegating to myself" may become indistinguishable from delegating to whoever shaped the model.

The second risk is manipulation at scale. AI systems already raise concerns for democracy because they can generate persuasive content, simulate social consensus, and personalize influence. Recent warnings about AI bot swarms emphasize that coordinated agents could mimic human behavior, distort public opinion, and harass or silence real users. A democracy of exocortices would need to distinguish legitimate personal cognitive extensions from armies of synthetic citizens, captured agents, or coerced proxies.

The third risk is privacy. A useful political exocortex would know a great deal: a person's values, fears, affiliations, private conversations, economic situation, and moral trade-offs. Buterin's proposal recognizes this by suggesting privacy-preserving technologies such as trusted execution environments, secure multi-party computation, encrypted circuits, and zero-knowledge proofs for sensitive governance processes. Yet technical privacy is not enough. Citizens also need legal rights over their models, logs, memories, and delegation settings.

The fourth risk is civic deskilling. Direct democracy is not only a method for aggregating preferences; it can also be a school of citizenship. Campaigns, arguments, public meetings, and referendums teach citizens what others value and why. Research on direct democracy stresses that participation and competence depend heavily on information environments, campaigns, and issue-specific learning. If exocortices quietly vote in the background, democracy may become more formally participatory while becoming less experientially civic.

The fifth risk is participation fatigue in disguise. Frequent voting can depress turnout or push citizens toward abstention and ballot roll-off. Exocortices solve this only if they reduce cognitive load without eliminating meaningful agency. A system that floods citizens with micro-referendums and then relies on agents to absorb the overload may reproduce the same problem at a higher technical level.


THE BALANCED PATH IS TO TREAT EXOCORTICAL DEMOCRACY AS A CONSTITUTIONAL DESIGN PROBLEM, NOT MERELY A PRODUCT-DESIGN PROBLEM. SEVERAL PRINCIPLES FOLLOW.

First, the exocortex must be user-sovereign: portable, inspectable, revocable, and not dependent on one platform. The LifeServer architecture proposed by the Grokya project points toward one way of making this principle concrete: open-source personal infrastructure, controlled by the user, with transparent update paths and data custody designed around the person rather than the service provider. Second, it should support graduated autonomy: advise-only, draft-vote, vote-with-confirmation, and bounded autonomous voting should be separate modes. Third, every automated vote should be explainable to the voter in plain language: "I voted yes because you previously prioritized X, opposed Y, and marked this issue as low-risk." Fourth, high-stakes categories should require renewed human confirmation. Fifth, the system should preserve secret ballots and private reasoning while still enabling public audit of eligibility, counting, and non-duplication. Sixth, democratic institutions should invest as much in shared deliberative infrastructure as in individual voting agents; otherwise, politics becomes millions of private preference machines talking past one another.

"Direct democracy in the age of the extended mind" should not mean replacing citizens with machines. It should mean recognizing that citizens already think through tools, records, networks, and institutions. The question is whether AI becomes an extension of democratic agency or a substitute for it.

The name personal exocortex captures the normative demand better than "AI agent" or "digital twin." An agent acts; a twin imitates; an exocortex extends. In a healthy democratic architecture, the exocortex would not be a ruler, representative, or oracle. It would be closer to a civic prosthesis: helping citizens understand more, participate more consistently, resist manipulation, and delegate less of their sovereignty to parties, platforms, and professional political classes.

The danger is that the same technology could do the opposite. It could centralize influence behind model providers, automate consent, and turn direct democracy into an illusion of participation. The future therefore depends less on whether LLM-powered voting agents are possible than on who owns them, who audits them, when they must defer to the human, and whether they strengthen the public sphere rather than merely optimizing private preference expression.

The extended mind thesis began with a simple question: where does the mind stop? In democratic politics, the parallel question is: where does citizenship stop? If a citizen's memory, attention, and reasoning are partly externalized into a personal exocortex, then democratic theory must expand to include those tools. But it must do so carefully. The goal is not government by AI. It is democracy for humans whose minds, like their societies, have become technologically extended.



Scientific Progress and Moral Progress:

A Reflection Based On rousseau's First Discourse

By Ma Daniel Gomides, Federal University of Ouro Preto



Since the outbreak of the Second World War, humanity has developed increasingly sophisticated scientific advances.

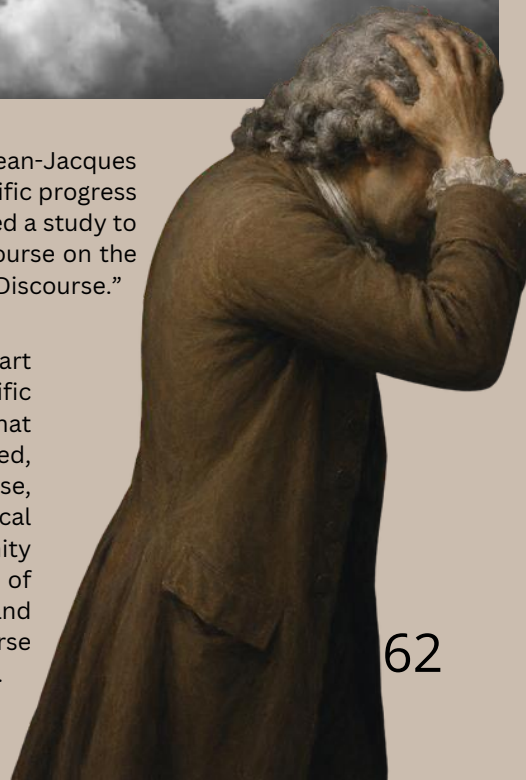


Initially, much of this progress pursued military purposes, such as the improvement of radar systems for detecting missiles and enemy aircraft, as well as machines capable of encoding encrypted messages. Over the decades, scientific development progressively came to occupy several sectors of human life and, in contemporary times, one of its most expressive moments has emerged through the expansion of Artificial Intelligences, the well-known AI's.

As an example of this expansion generated by AI, one may observe the growth of chatbots, which have now become globally recognized tools. The use of AI on platforms such as Google and META's social networks makes it safe to claim that AI has become part of human existence itself. In this context, increasingly frequent discussions arise regarding the limits of the use of these technologies. The major concern surrounding AI in its capacity to influence human decisions and structures of social organization as a whole. In societies marked by political instability, the irresponsible use of AI becomes extremely delicate. Thus, the following question becomes valid: is technical and scientific advancement necessarily linked to the moral progress of society? Even with so many automated systems, is fully human intervention still necessary?

This concern is not merely a contemporary issue. In the XVIII century, Jean-Jacques Rousseau had already identified possible tensions between the technical/scientific progress of his time and ethical/moral progress. In 1750, the French philosopher submitted a study to a public competition promoted by the University of Dijon and entitled it "Discourse on the Moral Effects of the Arts and Sciences," which later became known as the "First Discourse."

In summary, the essay fundamentally sought to answer whether science and art contributed to human morality or not. Rousseau firmly argued that scientific progress as a whole did not contribute to moral advancement, a position that radically contrasted with what the Enlightenment movement of his time proposed, namely the direct association between reason and human progress. In this sense, for Rousseau, technical advancement progresses independently of ethical advancement, since the sophistication of the instruments produced by humanity would not, by themselves, guarantee a more just moral organization. The core of Rousseau's critique lies precisely in the distinction between technical progress and human progress, and it is exactly this separation that makes the First Discourse relevant for identifying the instabilities of contemporary scientific development.





It is evident that analyzing Rousseau's argument in relation to the context of his own time is somewhat delicate, not only because eighteenth-century scientific progress was still extremely rudimentary, but mainly because the author lived in an intensely absolutist France under the rule of Louis XV. Nevertheless, the question raised by Rousseau remains highly relevant in contemporary times, especially if morality is understood as being directly connected to social justice and to a fair legal order.

First and foremost, it is necessary to recognize that AI algorithms are developed by individuals and institutions carrying very specific political, economic, and ideological interests. In this way, algorithms reflect the particular perspectives of their own financiers, allowing them to be used as instruments of manipulation, surveillance, and even social control. In an international scenario marked by wars, geopolitical disputes, and conflicts of interest, trust in any algorithm becomes unstable.

In this context, the debate surrounding AI cannot involve only its technical efficiency, but also the need to question which interests guide its construction. Thus, the discussion once again becomes entirely anti-Enlightenment, since the knowledge produced is used not to serve a universal or rational human ideal, but rather to satisfy specific interests.

Still analyzing the issue from the perspective of AI, these technologies have frequently been associated with ideals of efficiency, especially because of the speed with which information can be accessed and produced. However, just as technical development before AI was incapable of eliminating humanity's fundamental problems, the expansion of artificial intelligence also does not seem to guarantee the social or moral advancement of society. On the contrary, the absence of strict regulations and ethical governance mechanisms may cause such technologies to deepen already existing inequalities and consolidate increasingly violent and authoritarian political structures.

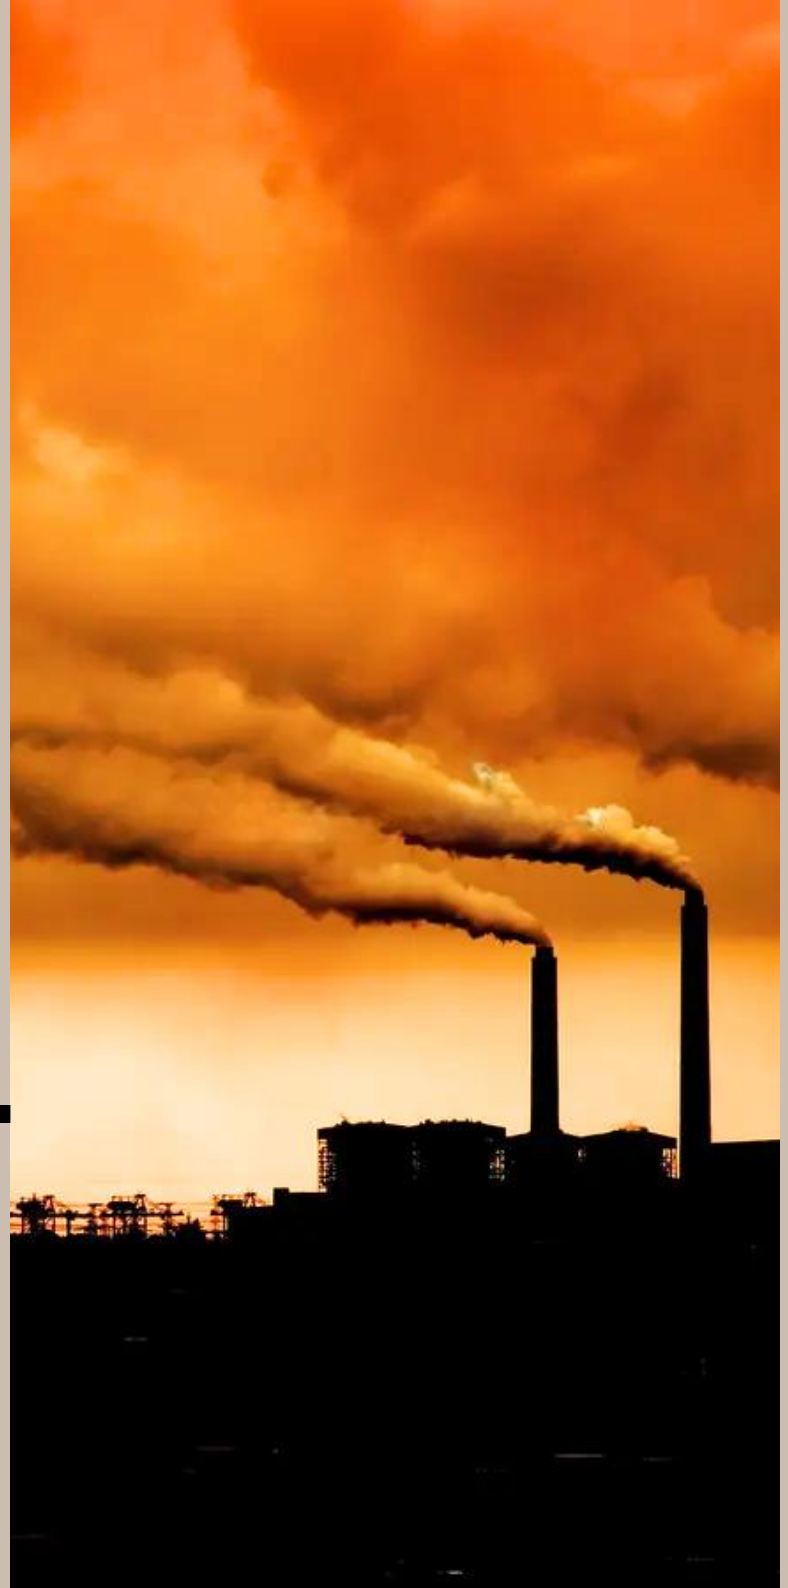
Thus, the major problem does not lie in the existence of Artificial Intelligences themselves, but in the belief that their technical development necessarily implies the moral and ethical capacity to use them properly.



A major example of this issue is the increasingly common use of AI by ill-intentioned individuals to commit crimes against others, produce illegal videos, and spread Fake News on a massive scale. From Rousseau's perspective, the following question should therefore be asked: if scientific progress truly accompanies moral progress, should not the major investors and financiers behind AI development have already created mechanisms capable of overcoming these types of crimes? Should governments not radically concern themselves with the limits of AI technologies?

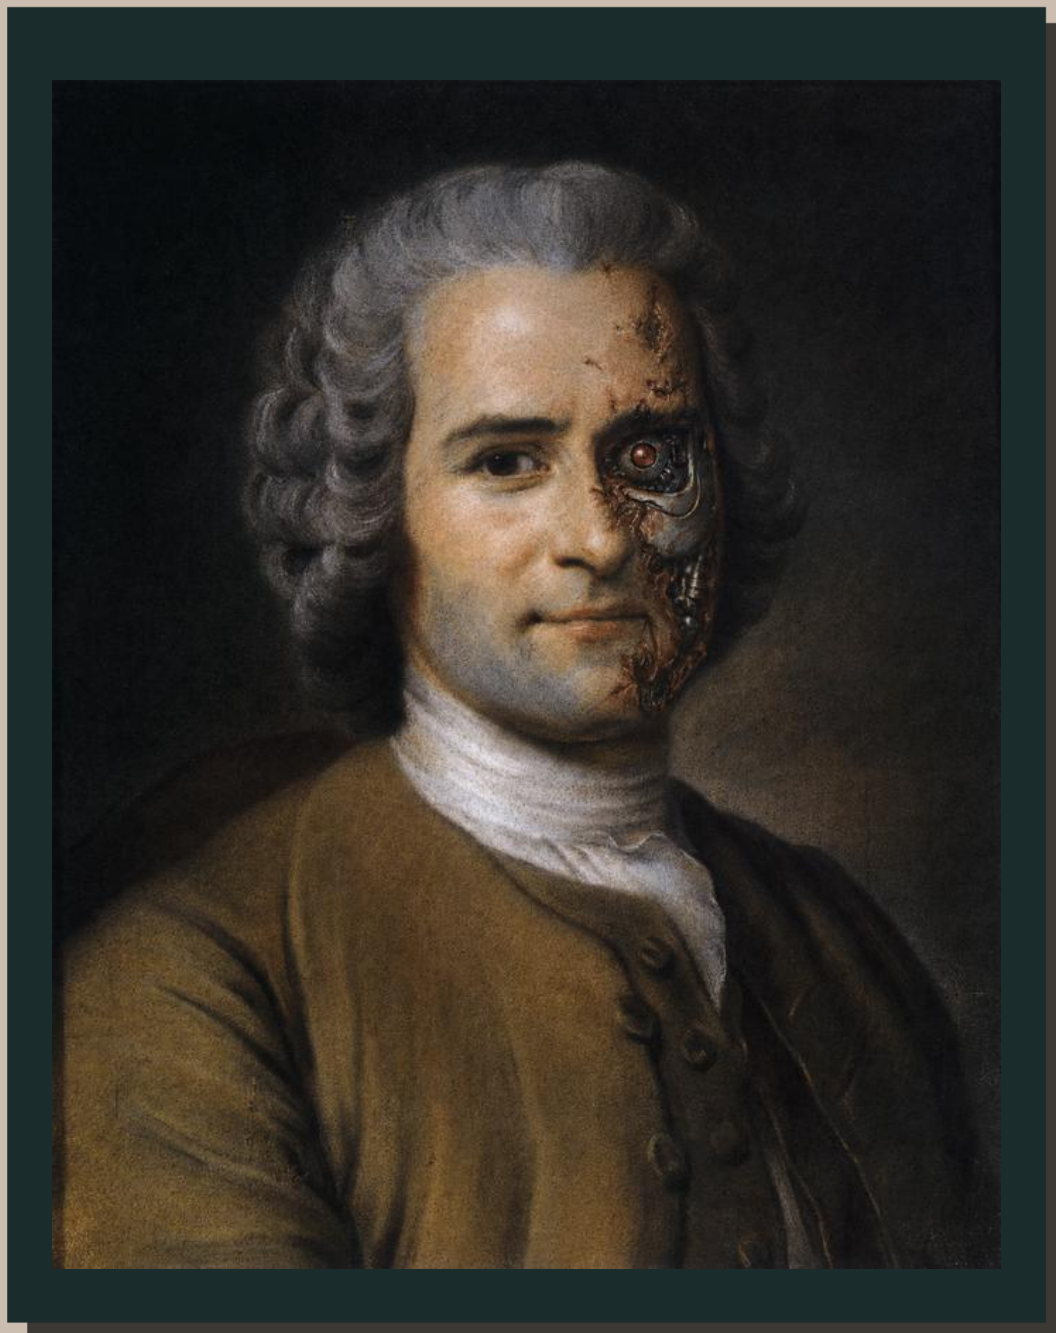
However, the field of technology involves many other components beyond AI itself.

Problems related to human life as a whole also become fundamental for understanding the ethical limits of technical development itself. What can be observed is that, although humanity has indeed achieved a high degree of scientific and technical development, various social and ethical problems persist, leading to the conclusion that technological advancement has not eliminated problems that existed long before it.

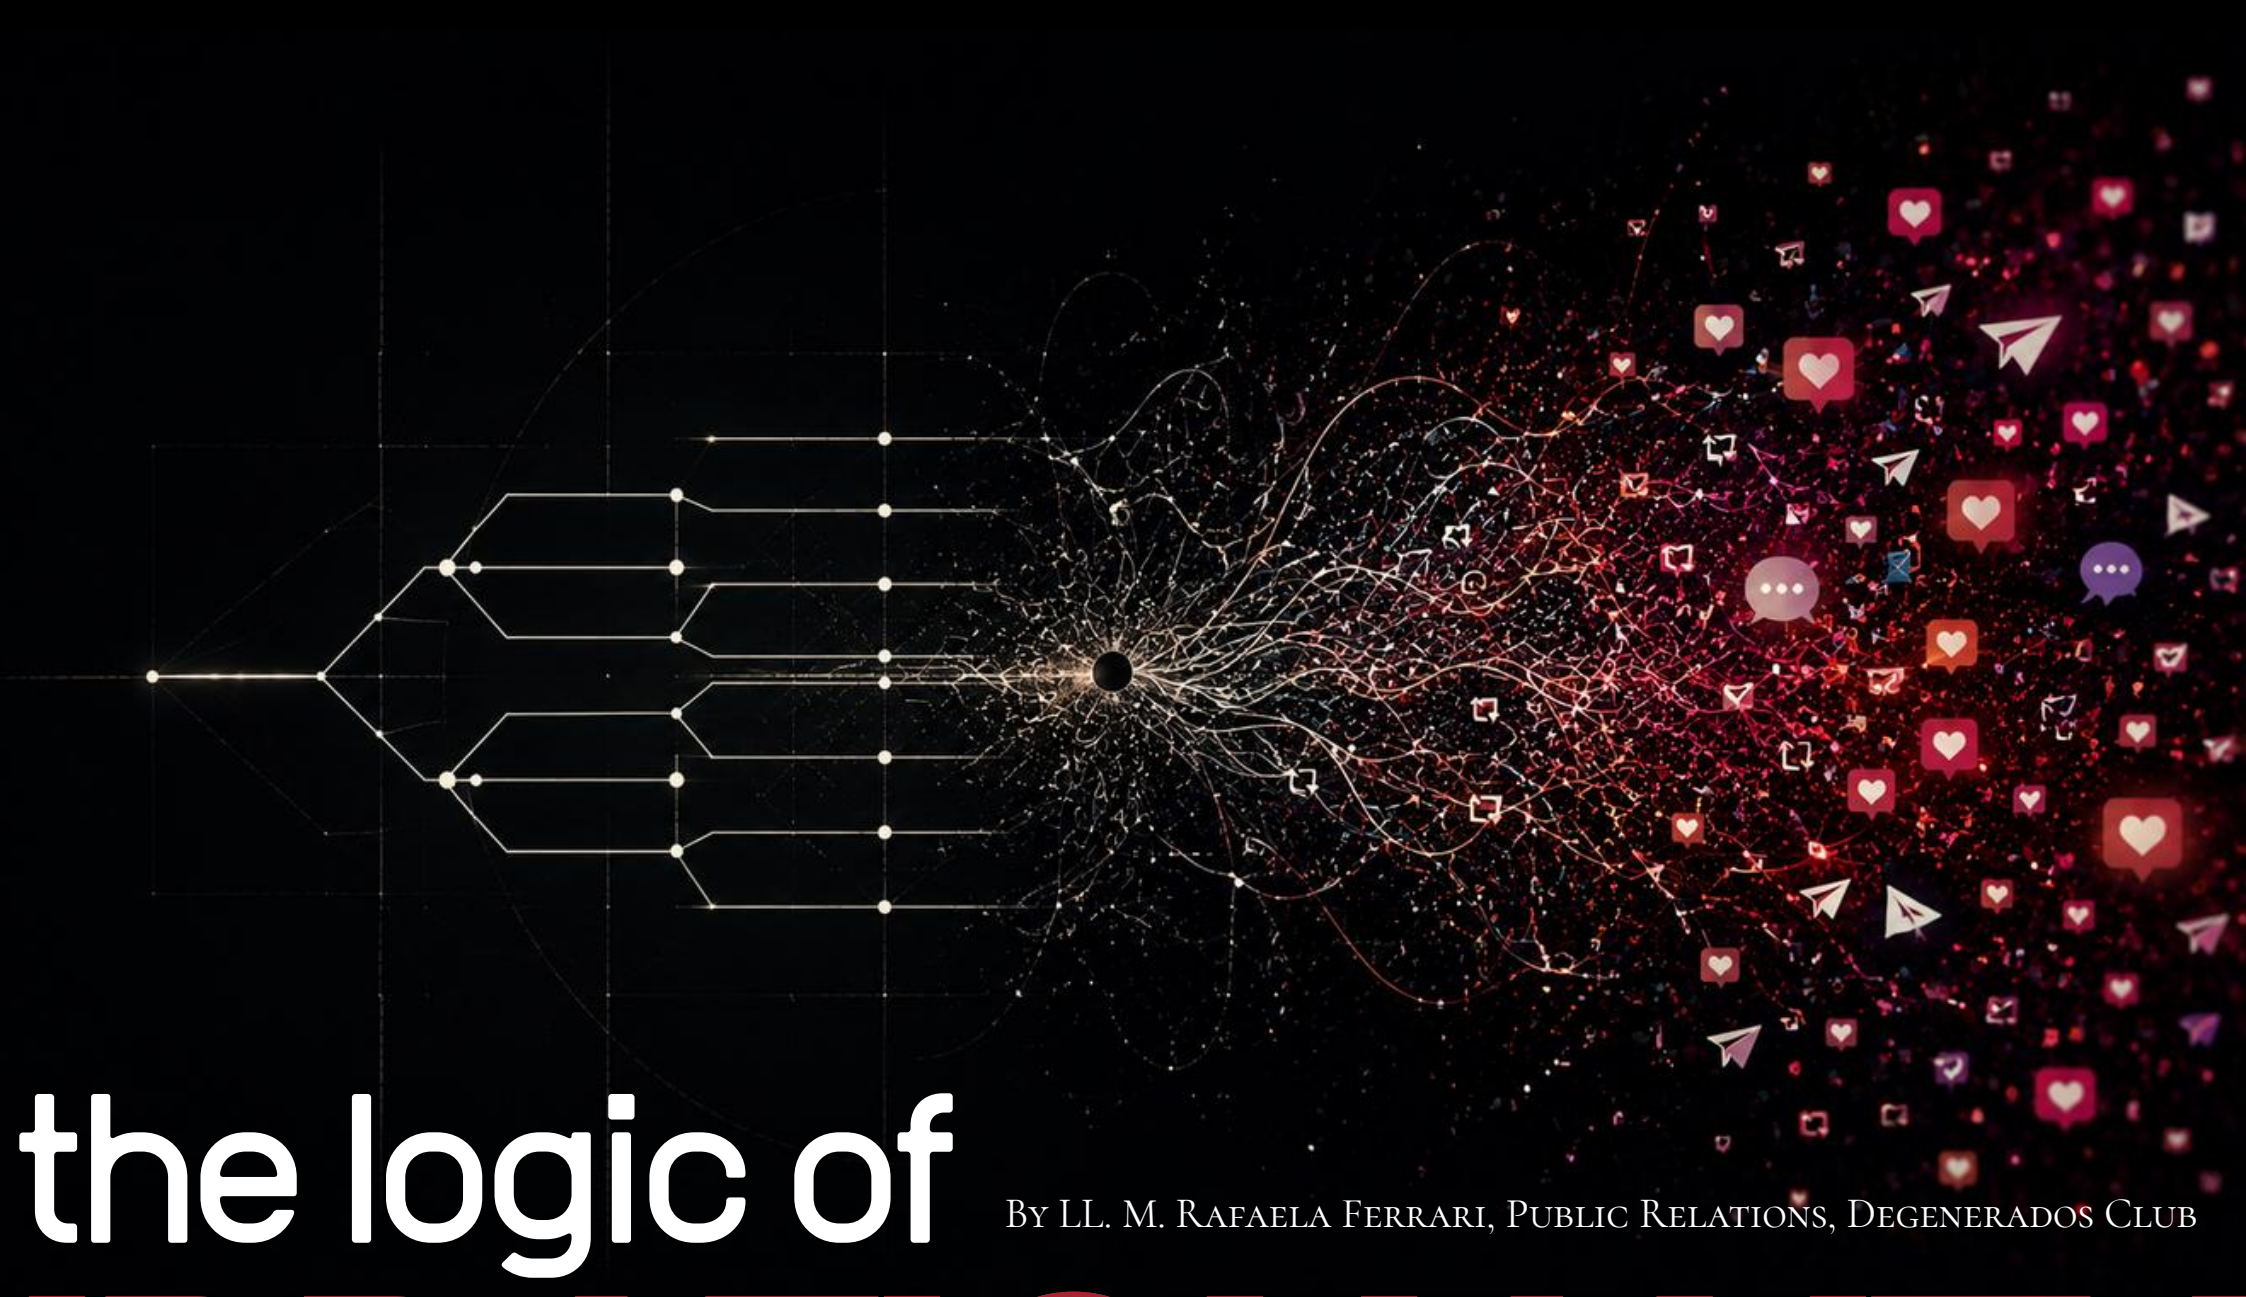


For example, according to research conducted in 2024 by the UN together with the Global Hunger Index, more than 673 million people still suffer from hunger worldwide (approximately 8.2% of the global population). This leads to the following reflection: assuming that scientific progress is accompanied by moral progress, in a highly globalized and technological world, should those responsible for scientific advancement not also be concerned with creating technical mechanisms capable of ensuring food reaches these people? Should nations that spend billions of dollars each year on scientific development not also be concerned with the basic conditions of the world population?

Another relevant example highlighting the contrast between scientific progress and moral progress may also be observed in environmental issues. Although humanity has developed sophisticated technical means to detect climate change and gas emission levels, some countries and private institutions continue prioritizing political and financial interests instead of ethical environmental awareness. As an attempt to address this concern, the Paris Agreement was signed in 2015, establishing goals for reducing global warming and greenhouse gas emissions. Nevertheless, since its creation, the agreement has failed to achieve effective compliance with its established goals precisely because of the issues previously discussed. This duality once again invites reflection: why, even in the presence of international agreements and all available technological tools, does the environment continue to deteriorate every day?



In light of everything that has been discussed, it becomes possible to understand that scientific progress still does not develop proportionally to the moral and ethical progress of society. The automation of decisions does not eliminate the need for moral reflection regarding their social effects, since, as previously mentioned, algorithms do not operate neutrally or without bias. Thus, the expansion of technology demonstrates that technical progress remains directly connected to the inclinations of the very individuals who produce it, meaning that scientific development does not, by itself, guarantee the construction of more ethical or morally organized societies. It is precisely at this point that Rousseau's reflection in his Discourse remains highly relevant. In short, the contemporary expansion of scientific and technological progress still surpasses humanity's moral capacity to deal with the effects of technology.



the logic of IRRATIONALITY

By LL. M. Rafaela Ferrari, Public Relations, Degenerados Club

There is something profoundly ironic about the fact that we live in the most technological era in history and, at the same time, perhaps the most emotionally manipulable since the ancient Roman crowds who shouted for bread, blood, and spectacle. Never have we had so much data, so many charts, so many indicators, so many experts, so many real-time metrics. And yet, the 21st century seems to function less by logic and more by narrative. Less by reason and more by the ability to control collective perceptions.

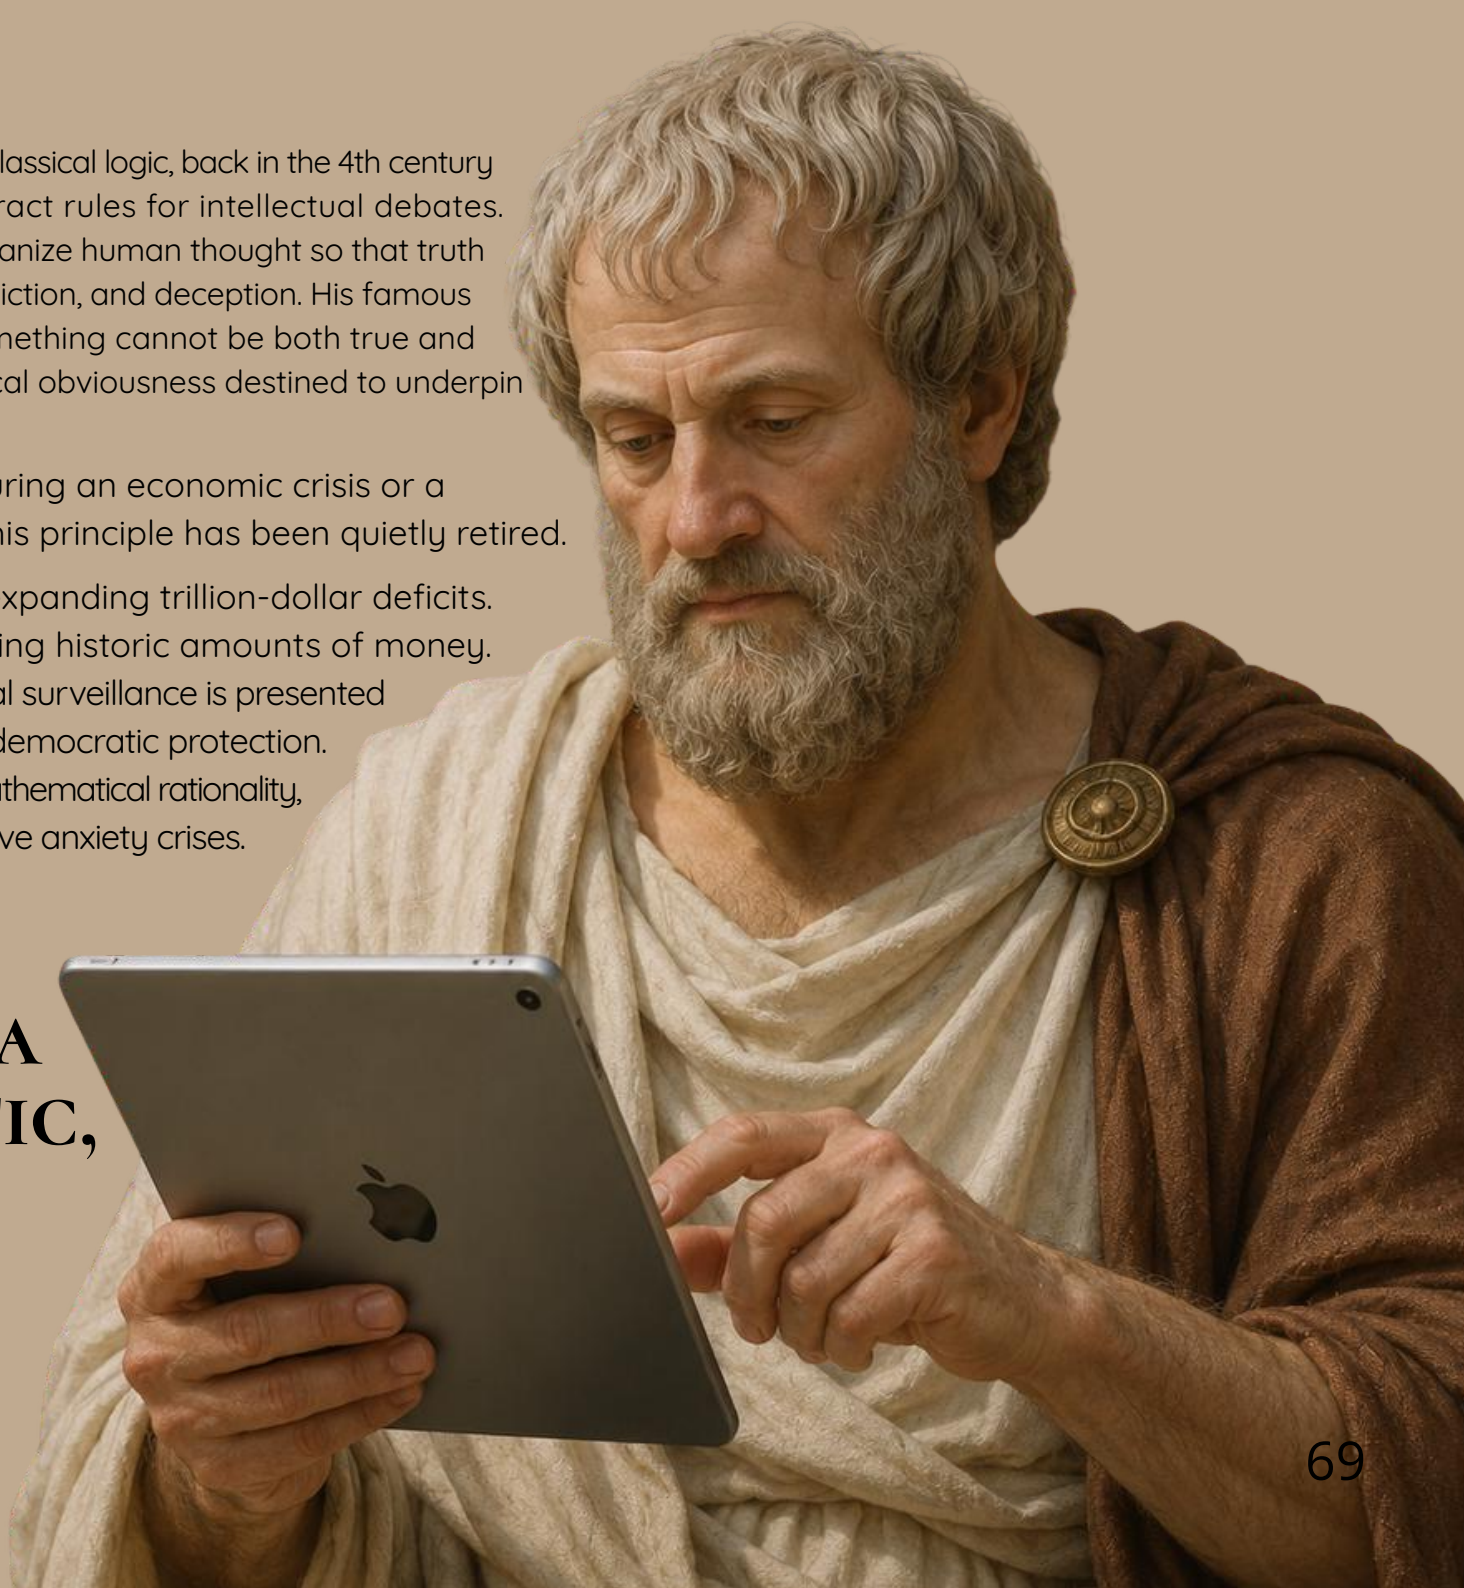
It is curious to realize that, in a world obsessed with artificial intelligence, supercomputers, and algorithms capable of processing billions of pieces of information per second, we continue making political, economic, and social decisions driven by the same primitive impulses that coursed through ancient empires: fear, greed, tribalism, desire for belonging, and need for collective validation. The difference is that now these impulses have fiber optics, sophisticated marketing, and sponsored campaigns.

Perhaps Aristotle would recognize this pattern immediately — can you imagine if he had an iPad?

When the Greek philosopher began to systematize classical logic, back in the 4th century BC, his intention was not merely to create abstract rules for intellectual debates. Aristotelian logic was, above all, an attempt to organize human thought so that truth could be distinguished from manipulation, contradiction, and deception. His famous principle of non-contradiction – the idea that something cannot be both true and falseat the same time – seemed like a philosophical obviousness destined to underpin theentire rational tradition of the West.

Centuries later, just open a social network during an economic crisis or a geopolitical conflict to realize that perhaps this principle has been quietly retired.

Today, governments defend austerity while expanding trillion-dollar deficits. Central banks claim to fight inflation by printing historic amounts of money. Wars are sold as humanitarian operations. Digital surveillance is presented as freedom. Censorship appears packaged as democratic protection. And financial markets, theoretically grounded in mathematical rationality, often move like teenagers in the midst of collective anxiety crises.



MODERNITY BUILT A RATIONAL AESTHETIC, BUT AS WE KNOW – OR AT LEAST SHOULD KNOW – AESTHETICS IS NOT SUBSTANCE.

The Greeks understood logic as an instrument for approaching truth. The contemporary world, in many aspects, has transformed logic into a tool of narrative persuasion. What is coherent doesn't necessarily matter; what matters is what is repeated with enough confidence to seem coherent.

In a civilization mediated by screens, perception often matters more than objective reality. Aristotle would probably be fascinated to discover that, in the 21st century, truth no longer depends on argumentative consistency but on the ability to go viral. Classical logic presupposed a rational human being seeking to understand the world. The problem is that neuroscience, modern politics, and the functioning of social networks have revealed something uncomfortable: human beings rarely desire truth above all else. Most of the time, they desire psychological comfort, confirmation of prior beliefs, and tribal belonging. Reason often enters only afterward, as the emotional lawyer for decisions already made instinctively.

Social networks – whether villains or heroines – realized this before the universities did.







Contemporary algorithms were not designed to promote rationality, but engagement, and few things engage more than fear, outrage, and conflict. The digital economy has transformed human emotions into monetizable raw material... every anxious click, every impulsive reaction, every collective polarization is converted into data, advertising, and profit..

Aristotle organized syllogisms; Silicon Valley organizes dopamine

The inevitable consequence is a society permanently excited, emotionally fatigued, and cognitively fragmented. Never – with due emphasis here – has it been so easy to access information. Never has it been so difficult to sustain attention, reflection, and intellectual coherence.

Logic demands time; the algorithm demands speed.

Perhaps that is why contemporary public debate seems less and less interested in understanding and more and more interested in winning: the goal is no longer to find the best interpretation of reality but to dominate the narrative before the attention cycle expires in fifteen seconds.



For centuries, money carried some physical link to real scarcity – gold was difficult to extract, silver had material limitations. So there was a concrete logic between value, labor, and supply limitations. Contemporary money, however, has become essentially political, and its issuance depends less on objective criteria and more on the circumstantial needs of the financial system and governments.

Naturally, everything continues to be presented with sophisticated technical language. Quantitative easing, monetary stimulus, flexibilization, systemic liquidity... elegant expressions to describe a rather ancient phenomenon: creating money to sustain structures that perhaps would not sustain themselves alone.

The irony is almost literary and poetic.

The ordinary citizen learns from an early age that printing money uncontrollably generates inflation, but when central banks do exactly that on a trillion-dollar scale, the practice gains sophisticated charts, press briefings, and a scientific appearance. The problem was never economic logic itself; it was always who has institutional authorization to ignore it.

And then Bitcoin emerges.

Not as a messianic technological miracle, nor as an absolute solution to all human problems, but as a kind of philosophical provocation to the contemporary world. And for the first time in a long time, a monetary system appears whose issuance does not depend on elections, political speeches, wars, or emergency decisions made in closed-door meetings by monetary authorities.

There is something almost Aristotelian about the brutal conceptual simplicity of Bitcoin: predictable rules, known issuance, verifiable scarcity, and relatively transparent mathematical functioning. It does not require emotional trust in governments, does not depend on the popularity of leaders, does not change because an electoral crisis appears on the horizon.



And perhaps that is exactly what bothers people so much.

The traditional financial system – full of interventions, artificial bailouts, continuous monetary expansion, and political manipulations – is treated as a symbol of rational stability. Meanwhile, an asset based on predictable mathematical rules is often labeled irrational, dangerous, or speculative. It is as if contemporary logic had completely inverted its own criteria.

Naturally, Bitcoin has volatility; the crypto market is full of excesses, opportunists, and speculative bubbles, but it would also be intellectually dishonest to ignore that traditional markets operate under equally impressive levels of collective irrationality, merely dressed in better ties and more expensive buildings.

The dot-com bubble, the 2008 crisis, the post-pandemic money-printing cycles, the recurring banking collapses – all of this reveals a deeply human characteristic of markets: they are rarely purely logical; they are collective psychological systems desperately trying to appear scientific.

The great contradiction of the 21st century is exactly this: we have never talked so much about science, data, and rationality, while our political, economic, and social structures continue functioning on the basis of ancestral emotional impulses. Technology has evolved exponentially; the human mind, not so much.

And perhaps Aristotle knew this from the beginning, without having an iPad.



Because classical logic never promised to eliminate human irrationality; it merely offered a method to recognize it. Reason does not exist because man is naturally rational; it exists precisely because we are not.

Logic has always been a civilizing attempt to contain our own emotional excesses, a kind of intellectual brake against collective hysterias, political manipulations, and comfortable illusions. The problem is that modern societies have discovered something extremely efficient: emotions mobilize faster than arguments.

Wars continue to start from emotional narratives; markets continue to inflate through collective euphoria; crises continue to be managed with short-term solutions that often create even greater distortions in the future; and digital crowds continue to confuse emotional intensity with intellectual depth.

Meanwhile, Aristotle remains silently uncomfortable on some university shelf, observing a world that has built machines capable of calculating quantum probabilities but still reacts politically like an Athenian assembly in panic.

The true contemporary crisis is not economic, technological, or geopolitical – it is a crisis of coherence. Because logic has not disappeared from the modern world; it has merely become selective.

We demand rationality from individuals, but tolerate systemic irrationality. We call any thought that challenges established narratives extremism, while normalizing monumental contradictions within institutional structures themselves. We treat mathematical predictability as a threat, but accept permanent monetary improvisations as responsible economic policy.

The question left by the Aristotelian tradition remains dangerously current: are we seeking to understand reality, or merely constructing emotionally bearable versions of it? And the answer to that question is the key, and perhaps explains much more about the 21st century than all the algorithms ever created.

MUSIC TRAIN



Host: Vessy Mink

@Connect.Club

Mon, Tue, Thu, Fri

SGT 10:00 EST 22:00

Music Train is the weekly live project from Vessy Mink and BK Han that began on the blockchain in Season 1, where the audience joined in real time to co-create original songs recorded on the spot and minted as NFTs: https://opensea.io/collection/v-mink

In Season 9, Vessy is releasing episodes full of candid interviews and live performances featuring Downtown Charlie Brown, Bob Springsteen & Joe Nix, David Baerwald (whose book The Fire Agent launches June 2), Sean McKee, Peter DiStefano, Jay B. Little, rObi Kumalo Hall, and more.

Born in the Connect Club app nine seasons ago, Music Train has grown into a global space for raw conversation, live creation, and genuine human connection. Produced with vision by BK Han and hosted with heart by Vessy Mink.
Catch every episode on YouTube and X.

DOWNTOWN CHARLIE BROWN

Featuring:



INTERVIEW



ABOUT

https://downtowncharliebrown.com

Featuring: JOE NIX HEARTISTRY & BOB SPRINGSTEEN





INTERVIEW



ABOUT

Featuring:

DAVID BAERWALD

INTERVIEW



ABOUT



SEAN MCKEE

Featuring:

INTERVIEW



ABOUT

https://seanmckeeband.com/home



PETER DISTEFANO

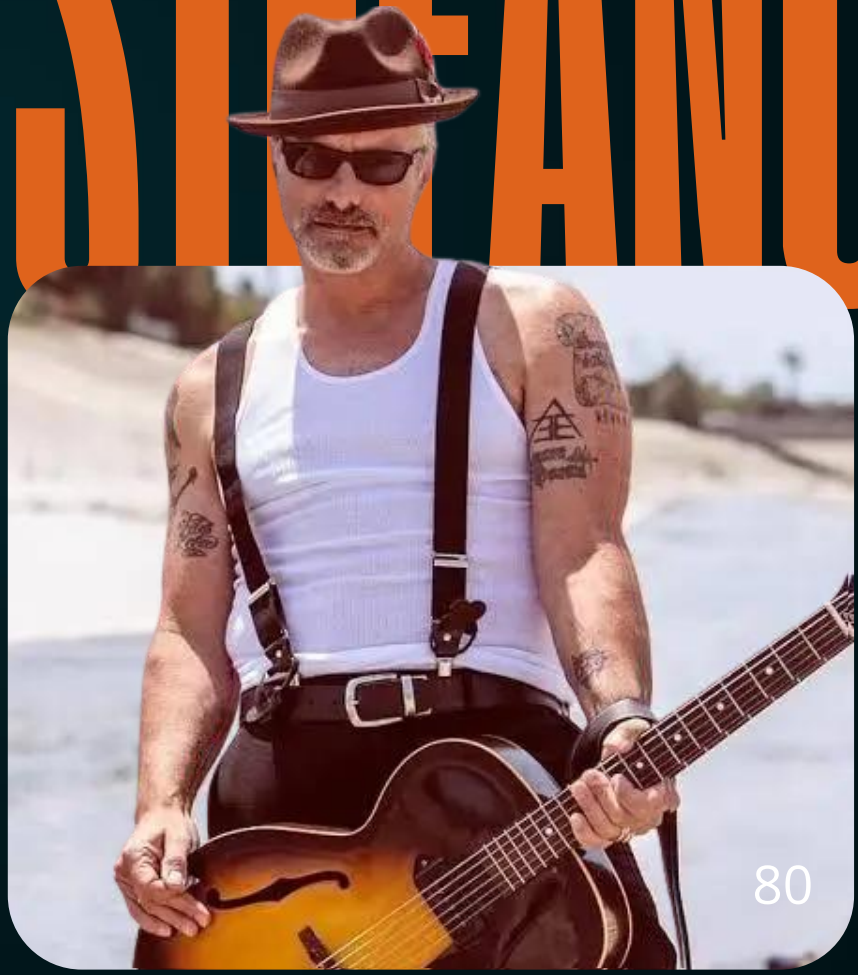
Featuring:

INTERVIEW



ABOUT





Featuring:

JB LITTLE

INTERVIEW



ABOUT





ROBI HALL KUMALO

Featuring:



INTERVIEW



ABOUT

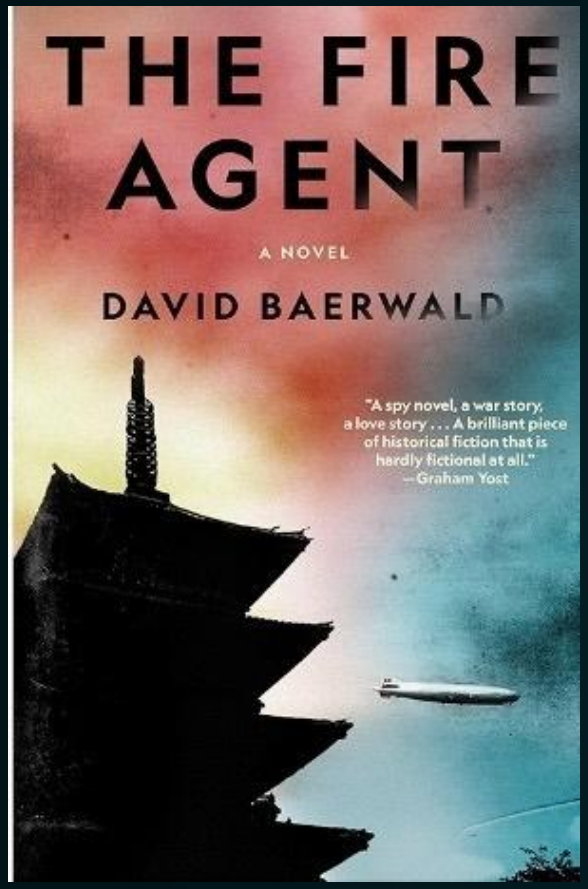
https://www.robbikumalo.com/

THE 18TH ANNUAL BAHA'I CHORAL FESTIVAL
SACRED A CAPPELLA MUSIC
MAY 21TH-24TH, 2026
ALL ARE WELCOME





ROBBI HALL KUMALO WILL JOIN 160+ CHOIR SINGERS FLYING IN FROM ALL OVER THE WORLD ON MAY 24TH 2026 AT BAHAI TEMPLE IN CHICAGOLAND AREA FOR THEIR 18TH ANNUAL CONCERT!



THE FIRE AGENT
BY DAVID BAERWALD
AVAILABLE EVERYWHERE
BOOKS ARE SOLD JUNE 2, 2026

A gifted German-Jewish linguist and idealist is drawn into the world of espionage at the dawn of the twentieth century. Operating undercover for a powerful chemical firm destined for infamy, he embarks on a life of high-stakes intrigue that spans continents, two world wars, and profound moral dilemmas.

From the cultural hubs of Europe to the volatile streets of pre-war Japan, he confronts danger, passion, and the rising tides of fascism, risking everything in a desperate bid to do what is right amid the chaos of history.
Blending espionage, sweeping romance, and the raw truths of war, this epic debut—rooted in the author's own family legacy is a riveting exploration of courage, loyalty, betrayal, and the human cost of progress. Historical fiction with the pulse of a thriller, it illuminates one of the most turbulent eras through the eyes of an unforgettable man.

A must-read for fans of ambitious, character-driven historical novels. Pre-order The Fire Agent today.

RECORD DROP ALERT!

EMMA MILLER

Emma Miller has released her captivating debut album, I Want to Be the Anchor.
This remarkable collection from the Edinburgh singer-songwriter features her spellbinding voice paired with poignant lyricism exploring love, loss, resilience and the power of self-discovery. Recorded in the Tennessee countryside, the album embodies years of dedicated artistry and personal reflection.

Available on vinyl, CD and digital download with Side A on select streaming platforms. I Want to Be the Anchor stands as an essential and authentic addition to contemporary folk music.

This summer, Emma teams up with Violetta Zironi for their special Scotland x Italy 2026 community meet-ups, featuring intimate performances and cultural adventures in Edinburgh and Italy in August.



PHOTO BY: DARIA BILYK @DASHIEE ON INSTAGRAM

Visit https://emmamiller.xyz/pages/home for more information and to connect with the artist, or head to https://violettas-moonshot.beehiiv.com/p/scotland-x-italy-2026-presented-by-violetta-zironi-emma-miller for full event details.

Philosophical Tools for the Digital Age:

launching "The Philosophy of Artificial Intelligence"

By Vitor Miranda, Federal University of Ouro Preto



The present issue of Virtualia Magazine has explored governance across radically different domains: digital art curation, music licensing, educational subversion, punk as an operating system, and the experimental laboratories of DAOs and tokenised events. Running through all these contributions is a shared, often unspoken question: who decides, and by what rules? This question lies at the heart of governance. But it is also the central preoccupation of a recently published work that deserves a special place in this debate.

The Philosophy of Artificial Intelligence (A Filosofia da Inteligência Artificial), written by Rodrigo Reis Lastra Cid and Pedro Luiz Caetano Filho and published by the Dissertatio series (UFPel / NEPFil Online), is a substantial investigation into the metaphysical, epistemological and ethical dimensions of AI. While its primary focus is the nature, knowledge, and morality of intelligent machines, the book consistently returns to a governance insight that binds it to the other articles in this volume: no system of collective decision – whether human or algorithmic – is neutral, complete, or self-justifying.

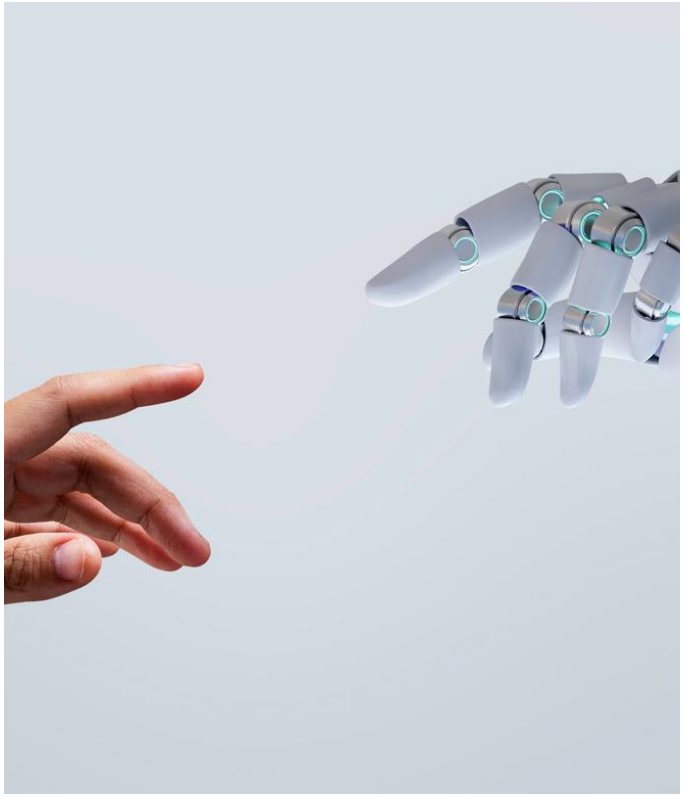


The book opens with a rigorous discussion of social choice theory's foundational result: Arrow's impossibility theorem. Cid and Caetano Filho show that every voting mechanism inevitably sacrifices some desirable criterion (independence of irrelevant alternatives, absence of a dictator, transitivity, etc.). This is not merely a mathematical curiosity; it is the philosophical bedrock on which any governance architecture – from a corporate board to a DAO – must be built. The contributors to this magazine have documented how Web3 platforms often default to "one token, one vote" without examining its plutocratic implications (Fer Caggiano, VAN Ameneyro). The book's lesson is that such choices are always value-laden.

Governance begins not with code, but with the prior decision of which trade-off to privilege.

The book's chapters on opacity, bias and responsibility speak directly to the governance challenges raised elsewhere. When Vessy Mink describes a music licensing model that returns publishing rights to the artist, she is addressing a governance failure of the traditional industry: the concentration of decision-making power in intermediaries. Rafaela Nóbrega's subversive art projects in Brazilian schools are exercises in redistributing curricular authority. Steve Coulter's "punk operating system" is a manual for building parallel institutions when formal ones become extractive. All of these are responses to what Cid and Caetano Filho analyse as the problem of many hands – the fragmentation of moral responsibility in complex technological societies.

The book argues that no algorithm can dissolve this fragmentation; only transparent, auditable, and participatory governance can.

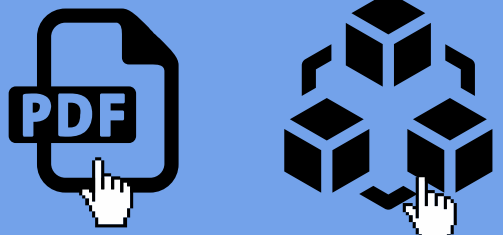


The book's appendix on the Brazilian AI Plan (PBIA 2024-2028) offers a rare case study of national governance of AI. It praises the plan's recognition that sovereignty and equity must shape technological development, but also warns that without robust mechanisms for public deliberation, transparency, and accountability, the plan may reproduce the very inequalities it seeks to overcome. This mirrors the magazine's concern with the gap between decentralisation rhetoric and actual power structures.

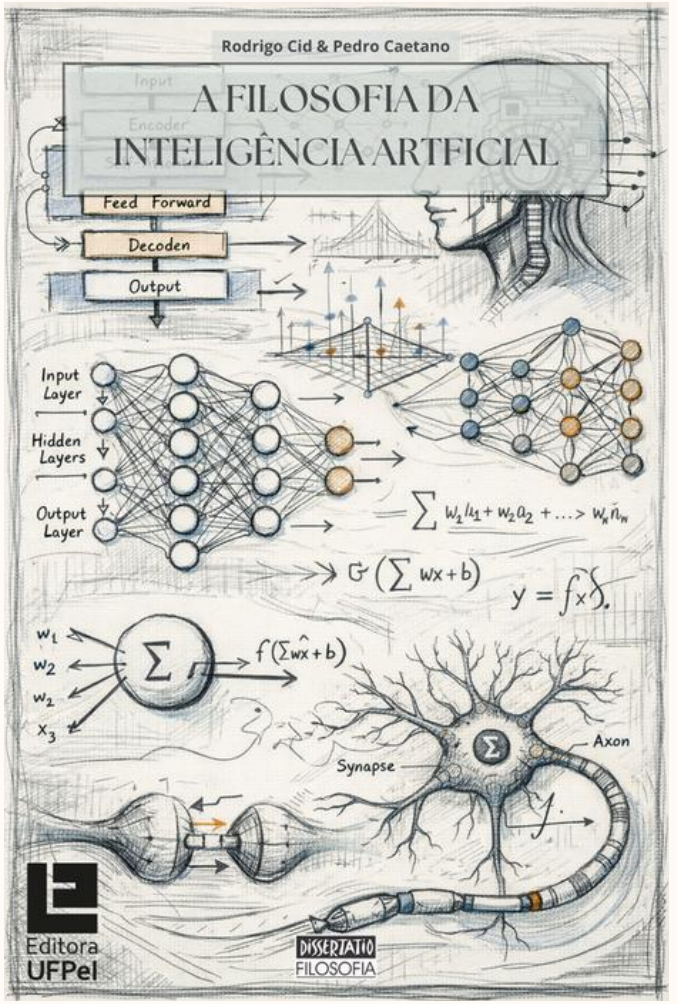
Finally, the book's own method – written in systematic collaboration with AI systems – embodies a governance experiment. The authors transparently declare their use of Claude, DeepSeek, Perplexity and ChatGPT, and they propose the CRediT-IA framework for responsible AI authorship. This is a governance protocol for academia, one that insists on human final authority while acknowledging the legitimate role of artificial assistants. In doing so, it aligns with the magazine's editorial call for bilateral transparency: if authors must declare their use of AI, then editors and reviewers must also disclose theirs.

The Philosophy of Artificial Intelligence does not offer a ready-made governance model. What it offers is more valuable: a conceptual toolkit to ask better questions. What kind of intelligence are we governing? Who is accountable when black-box systems fail? How do we embed fairness into algorithms without collapsing into mathematical impossibility? These questions are not optional. They are the very substance of governance in the age of intelligent machines.

THE BOOK IS AVAILABLE FOR FREE DOWNLOAD AT:



DOI: 10.5281/zenodo.20143966
ISBN: 978-65-998644-0-7



The readers are invited to explore it alongside the contributions to this issue – and to continue the conversation that, as the book's preface says, is the only way to build a future where technology amplifies, rather than diminishes, human philosophical depth.

VIRTUALIA


PALCO
DEGEN_
DO LIVRO AO NFT: TOKENIZANDO O MERCADO LITERÁRIO TRADICIONAL
GUTA NASCIMENTO
EMANUEL SOUZA
FER CAGGIANO
THE PHILOSOPHER
DEGEN