

ANTI-IDIOCRACY

The Outfinitist Metacult Investigation



ANTI-IDIOCRACY

The Outfinitist Metacult Investigation

* * *

Sînică Alboaie

Copyright © [2026] Axiologic Research

All rights reserved.

KDP publishing rights held by Outfinity SRL (Iași, Romania).

Available for free on Axiologic website (www.axiologic.net).

Author's Note

This work was created under the umbrella of Axiologic Research as part of two interconnected research programs, one dedicated to Artificial Intelligence and the other to Outfinitism, both led by Sînică Alboaie. To explore the details of how these two programs intersect and build upon each other, we invite readers to visit the Outfinity Venture Studio page at www.outfinity.ch.

While Artificial Intelligence language models were used to assist with text generation, the foundational philosophy, thematic structure, and analytical approach derive exclusively from the author's decades of dedicated study of social phenomena, social technologies, and genuine hard technologies resulting from various European and self funded research projects. Recently, within the Achilles research project, we have been deeply studying AI generated literature and its automated verification using neuro symbolic approaches; all the free books made available to the public are part of the suite of works we use to test our latest technologies.

.

CONTENTS

Introduction. Why Being Right Is Not Enough	1
1. Truth Loses Without Culture	9
2. Elites and the Monster They Feed	16
3. Moloch Already Has a Religion	24
4. Meta-Rationality and the Invention of Outfinitism	33
5. What We Choose to Hold Sacred	45
6. Between Moloch and the Basilisk	54
7. The Outfinitist Metacult: Belonging Built on Fallibility	62
8. Let Us Gather	72
Bibliography	83

Why Being Right Is Not Enough

“If human beings evolved from monkeys, why are there still monkeys?”

The objection is engineered for a quick victory. It fits into one breath, produces a smile of superiority, and forces the other person to explain an evolutionary tree inside a conversation that has already rewarded the punch line rather than the explanation. The answer is elementary. Evolution does not claim that today’s monkeys turned, one by one, into today’s humans. It claims that different branches share ancestors. The existence of your cousins does not refute the existence of shared grandparents. The continued existence of wolves does not refute the origin of domestic dogs.

The problem is not that the question is naive. Any sincere question can open an investigation. The problem begins when misunderstanding is used as a weapon for closing the investigation. A ten-second error can defeat a ten-minute truth in a social environment that grants status to the person who appears to have won, not to the person who has understood more.

The same pattern appears when a cold day is offered as a refutation of climate change, as though local weather and long-term climatic distributions were the same thing. It appears again when the illness of one vaccinated person is treated as proof that a probabilistic intervention cannot reduce risk, or when disagreement among specialists at the frontier of a discipline is used to claim that the entire discipline knows nothing. A case does not erase a distribution. An exception does not abolish a regularity. A disputed frontier does not dissolve the body of knowledge that made the dispute precise enough to exist.

These errors are not confined to people with little education. A physician may reason with extraordinary care inside a specialty and with childish certainty in

politics. An engineer may build systems of great complexity and become captive to an ideology that flatters his identity. A billionaire may mistake the outcome of a rare selection process for proof of universal superiority. A professor may use verbal sophistication to defend a belief he would never permit anyone to test. Intelligence does not immunize us against stupidity. It can merely provide stupidity with a better public-relations department.

In this book, **idiocy** does not name a biological caste, a level of education, or a category of persons whose dignity is somehow reduced. It names a dangerous configuration:

a weak model of reality + confidence disproportionate to the evidence + refusal of correction + power to transfer the cost of error to others.

Anyone can enter this configuration. An institution full of specialists can become idiotic. A market can reward locally rational decisions that together produce an absurd collective outcome. A movement devoted to opposing idiocracy can become idiotic the moment it imagines itself exempt from its own diagnosis.

Idiocracy emerges when that configuration is selected, financed, publicized, normalized, and armed with power. It is not simply “government by people of low intelligence.” It is government through models that do not recognize their boundaries. It may be populist or technocratic, religious or anti-religious, corporate or revolutionary. It may speak in the name of the people, the market, science, the nation, God, or the future. Its signature remains the same: the theory commands reality to obey, and when reality refuses, the cost is sent downward to those with less power.

Human dignity does not depend on the ability to explain evolution, statistics, or artificial intelligence. Rights are not prizes for the clever. Authority, however, must remain linked to competence, domain, responsibility, and revocability. A person who cannot represent a problem accurately has no moral entitlement to make irreversible decisions for millions of people in that domain. A person who does understand the problem but deliberately exploits

public confusion raises an even more serious question—not about intelligence, but about legitimacy.

The destination of this book

This book was not written merely to expose the mechanisms of idiocracy. Exposure is necessary and insufficient. A correct analysis, left alone on a shelf or inside an isolated mind, often loses to a culture that supplies belonging, rhythm, status, memory, and reflex. Truth without culture can become a footnote in a civilization organized by more memorable falsehoods.

The destination is therefore explicit from the beginning:

People who understand the mechanisms of idiocracy must find one another, help one another, and build a culture capable of turning understanding into collective capacity.

Not a culture of superiority, but a culture of correction. Not a community that sanctifies a leader, but one that begins with the leader's fallibility. Not an ideology that eliminates disagreement, but a sufficiently broad framework within which productive disagreements can continue without destroying the conditions of dialogue, liberty, and exploration.

Ancient religions did not unite people through metaphysical propositions alone. They built calendars, rites, obligations, networks of care, memorable stories, and distinctions between what could be negotiated and what was held sacred. States, nations, universities, armies, companies, and political movements later developed their own versions of the same machinery. Culture is what turns a proposition into a reflex, a preference into an obligation, and a collection of individuals into collective capacity.

In the age of artificial intelligence, we can no longer assume that a list of technical principles will be enough. AI will amplify the habits of the institutions that deploy it. In a culture of racing, it will compress reflection and reward premature release. In a culture of control, it will extend surveillance and behavioral administration. In a culture of responsible exploration, it may reduce the costs of education, translation, research,

cooperation, and contesting power. AI safety concerns models, data, security, and evaluation; it is also, inescapably, a problem of culture.

Creating a culture does not mean inventing a new police force for thought. It means deciding consciously which behaviors receive prestige, which practices are repeated, which boundaries remain in force even when violation would be profitable, and which kinds of persons become easier to form inside the community. When we refuse to choose, platforms, markets, states, charismatic figures, and optimization systems choose for us.

The argument of this book proceeds toward a practical conclusion. There is no final policy capable of permanently defeating idiocracy. There is no perfect constitution, benevolent market, enlightened elite, or superintelligent machine that can remove the need for judgment. The durable answer must be cultural and institutional at once: a distributed capacity to detect limits, revise models, protect human dignity, preserve alternatives, and coordinate action without giving any center an unlimited claim to truth.

That culture needs a social form. I call the general form a metacult and the particular experiment developed here the Outfinitism Anti-Idiocracy metacult—or, more simply, the Outfinitist metacult.

A metacult, and this metacult

A metacult is a voluntary culture that includes its own culture-forming machinery within the field of observation. It uses shared stories, repeated practices, mutual obligation, recognition, and institutional memory, while also studying how those mechanisms can produce conformity, capture, dependency, and abuse. The prefix meta does not claim superiority. It names recursive awareness: the culture must be able to examine and revise the processes by which it creates loyalty, authority, and meaning.

The word remains deliberately close to cult. Human beings do not become immune to ritual, prestige, imitation, emotional synchrony, or sacred commitment merely because they call themselves rational. A conventional cult protects its center from doubt. A metacult makes the capacity to doubt

its own center part of what it protects. Its leaders are neither oracles nor owners of the truth; its texts are versioned rather than sealed; its practices disclose their purposes; departure is not treated as moral treason; and every concentration of power must carry a path of challenge, succession, or exit.

A companion book is devoted to the general theory of metacults: their anthropology, design grammar, rituals, governance, economics, and characteristic failure modes. This book has a narrower and more urgent task. It develops one possible metacult for people who recognize, with varying degrees of clarity, both the emancipatory potential and the civilizational dangers of modern science, technology, and especially artificial intelligence.

Other metacults could be smaller and domain-specific. Scientists might build one around the limits of inference, replication, measurement, and generalization. Economists might build one around externalities, concentration, financial fragility, and the boundaries of optimization. Engineers might organize one around safety margins, incident memory, reversible deployment, and the duty to preserve alternatives. Such communities need not become subsidiaries of a universal organization. They may overlap, disagree, cooperate, or remain autonomous.

The Outfinitist metacult addresses the layer those specialized communities cannot solve alone: the social reflexes that determine whether knowledge is heard, whether uncertainty can be admitted without humiliation, whether power accepts limits, whether well-being is treated as infrastructure, and whether AI expands human agency or merely perfects extraction and control. Its purpose is not to give humanity one more final doctrine. It is to create a culture strong enough to coordinate action and free enough to correct itself.

The basic concepts

Meta-rationality is not a term invented in this book, nor does it designate a single accepted psychological theory. Over several decades, work on reflective judgment, dialectical and postformal thinking, cognitive-affective integration, adult development, and wisdom has described related capacities: tolerating uncertainty without paralysis, comparing justifications, working with

multiple frameworks, separating identity from theory, and making decisions that remain open to revision. The literature is heterogeneous and does not establish a universal developmental stage called meta-rationality. I use the term as a working name for a more mature relationship with reason.

Rationality asks: given these premises, rules, and objectives, what follows? Meta-rationality adds: why are these the premises, where does the model cease to fit, what has it excluded, who bears its errors, and when should the frame itself change? It does not deny logic. It recognizes that every applied logic begins from assumptions, scales, goals, and values that must also be examined.

Outfinitism is the term proposed here. It sharpens the meta-rational impulse by organizing it around the systematic study of limits: the limits of theories, models, measurements, institutions, technologies, categories, and our own capacities. It is premature to call Outfinitism an autonomous science. It can be treated as an operational philosophy, a metadisciplinary research program, and, in particular contexts, an internal science of limits practiced within other sciences.

Mathematics studies limits of proof. Computer science studies limits of computation and verification. Economics studies externalities, incomplete information, and coordination failures. Psychology studies limits of attention, memory, self-knowledge, and representation. Sociology studies limits of formal rules and centralized observation. Political theory studies limits of delegation, legitimacy, and control. Outfinitism asks these disciplines to make boundary conditions central rather than decorative.

Meta-rationality is the capacity to see that no frame is the whole. Outfinitism is the attempt to make the edges of frames visible, discussable, and actionable. Meta-rationality can exist without the name Outfinitism. Outfinitism cannot remain healthy without meta-rationality; otherwise it becomes one more doctrine about limits imposed without recognizing its own.

Elites are groups and persons with disproportionate influence over capital, infrastructure, knowledge, institutions, or symbols. They arise in every complex society, including communities that officially deny having them. Their existence does not legitimize them. An elite becomes legitimate only to

the extent that its competence is domain-specific, its service is verifiable, its power is contestable, succession is possible, and the costs of its errors are not simply pushed down the hierarchy.

Moloch is the metaphor used here for races in which every participant has local reasons to accelerate even as the collective result becomes destructive. A company cannot slow down alone when competitors release unsafe products. A platform cannot unilaterally stop capturing attention when the market rewards retention. A state cannot abandon a military technology when it fears that its adversary will not. Nobody necessarily desires the final outcome, yet the incentives produce it.

The Basilisk is the symmetrical danger: the promise that one sufficiently intelligent center can end the race by seeing everything, deciding correctly, and enforcing compliance. If Moloch is coordination failure through competition, the Basilisk is coordination failure through total concentration. One destroys the future by making stopping irrational; the other protects the future by making dissent impossible.

A **decentralized brand** is an institutional arrangement in which a shared identity, promise, reputation, and infrastructure can be carried by autonomous projects rather than owned absolutely by one legal entity. It does not eliminate centers. It limits their monopoly over meaning, contribution, and succession. It allows a network to preserve coherence without requiring every participant to become a subsidiary.

Anti-idiocracy is not rule by a self-declared cognitive aristocracy. It is the constitutional refusal to let weak models, protected from correction, acquire irreversible power. The Outfinitist metacult is the social technology through which that refusal may become lived practice. Meta-rationality supplies the psychological posture. Outfinitism supplies the disciplined study of limits. Decentralized brands supply an institutional architecture. The metacult supplies belonging, repetition, memory, mutual aid, and a path from private understanding to public capacity.

What is inherited and what is new

Very little in this book appears from nothing. The critique of bounded rationality belongs to a long tradition. The psychology of reflective judgment, adult development, and wisdom precedes Outfinitism. The study of polycentric governance owes much to Elinor Ostrom. The limits of formal systems, calculation, optimization, and aggregation have been explored by Gödel, Turing, Arrow, Simon, Goodhart, Ashby, and many others. Anthropology and sociology have long examined ritual, religion, charisma, bureaucracy, oligarchy, and collective identity. The dangers of scapegoating and mimetic rivalry were given a powerful, if disputed, formulation by René Girard.

The entrepreneurial analysis developed here draws on work about overconfidence, excess entry, undiversified risk, survival distributions, near-misses, well-being, and the extreme concentration of outcomes. The argument for decentralized brands extends earlier work in distributed systems, self-sovereign data, smart cities, and decentralized forms of economic coordination.

What is new is the proposed synthesis.

First, Outfinitism places the study of limits at the center of a positive program rather than treating limitations as footnotes appended after a theory has already claimed authority. It asks not only whether a model works, but where, for whom, at what scale, under which incentives, with which failure costs, and through which mechanisms of correction.

Second, it connects epistemic limits with institutional legitimacy. A leader who admits uncertainty but retains irreversible power has not solved the problem. A model that reports confidence intervals while the affected person has no right of appeal remains politically incomplete. Knowledge about limits must change the architecture of authority.

Third, it treats culture as necessary infrastructure. The ability to say “I was wrong,” “this model no longer applies,” or “we must preserve an alternative” cannot survive as a private virtue if the surrounding culture punishes

correction and rewards certainty. Institutions must turn these abilities into status, ritual, memory, and mutual protection.

Fourth, it proposes the metacult as a general category of social technology and develops the Outfinitist metacult as one concrete instance. This is a form of mass belonging founded not on the infallibility of leaders, but on their declared and enforceable fallibility; not on a final doctrine, but on disciplined exploration of boundaries; and not on hidden emotional manipulation, but on transparent practices that make people more capable of leaving as well as staying.

Finally, it offers a proposition to present and future elites. The predator's morality—exploit confusion because confusion is exploitable—is not realism. It is a short-horizon strategy that degrades the ecology on which power depends. Elites may preserve and strengthen legitimate influence by investing in broad psychological well-being, education, contestable institutions, and distributed competence. The most credible signal of strength in the AI age may become not what power can do, but what it binds itself not to do.

The road ahead

The first chapter explains why truth loses when it lacks a culture capable of carrying it. The second examines elites and the recurring historical mistake of cultivating forces that appear useful until they acquire their own leaders, rituals, and appetite. The third studies startup culture as an already functioning quasi-religious culture of risk, rarity, and retrospective merit, then introduces the decentralized brand as an alternative infrastructure for exploration.

The fourth develops meta-rationality and Outfinitism in detail. The fifth asks what a pluralistic culture may legitimately hold sacred and what it can learn—without imitation or contempt—from religions, especially Christianity. The sixth places the problem inside the AI transition, between the acceleration of Moloch and the administrative dream of the Basilisk. The seventh distinguishes the general idea of a metacult from the Outfinitist metacult developed here and describes its decentralized architecture of

fallibility. The eighth asks who should gather, what useful services the community could provide, how exploration might be financed, how success should be measured, and how elites might acquire a legitimacy stronger than fear.

The destination will remain visible throughout. We are not trying to win one more argument about stupidity. We are trying to build a culture in which fewer errors become destiny.

Truth without culture loses. Culture without limits becomes an idol. Limits without love become a prison.

Outfinitism begins when these sentences become a way of living together.

Truth Loses Without Culture

A false idea does not need to be intellectually superior in order to win. It may be shorter, warmer, more flattering, easier to repeat, better aligned with group identity, or attached to a person who controls the stage. Truth can survive all these disadvantages inside a laboratory notebook. It cannot reliably govern a civilization through the notebook alone.

This is the first uncomfortable premise of anti-idiocracy: correcting propositions is not the same as correcting systems. If the social environment continues to reward certainty, humiliation, repetition, and tribal loyalty, the next false proposition will occupy the same place as the previous one. We will have removed a symptom while preserving the selection mechanism.

From error to system

Human beings make mistakes because our information is incomplete, our attention is narrow, our memory reconstructs, our emotions prioritize, and the world changes while we are modeling it. Error is not an anomaly added to an otherwise perfect intelligence. It is a permanent feature of finite minds operating under time pressure.

Most errors remain small. I misremember a date, misread another person's expression, or choose the slower road. Idiocracy begins when institutions reward some errors, protect them from correction, and amplify their consequences. The relevant unit is no longer the mistaken individual. It is the ecology that decides which mistakes become careers, markets, movements, and laws.

Consider repetition. A statement becomes easier to process after it has been encountered several times. Familiarity can then be misread as truth. The psychological literature calls this the illusory truth effect (Dechêne et al., 2010). The mechanism does not require a person to be unintelligent; it exploits the ordinary architecture of memory. A platform that repeatedly exposes millions of users to a simple falsehood can therefore produce a social fact: not the truth of the claim, but the expectation that many others believe it.

Once that expectation exists, people adjust. Journalists cover the claim because it has become salient. Politicians answer it because silence appears evasive. Experts simplify their response until the response begins to resemble the distortion. Opponents repeat the falsehood while refuting it and increase its familiarity. The original error becomes an agenda-setting machine.

The effect is strengthened by asymmetry. Fabrication is cheap. Correction must identify the claim, supply context, explain uncertainty, and survive the suspicion that complexity is merely concealment. A falsehood can be optimized for transmission. Reality must remain compatible with itself.

The standard educational response is to improve media literacy. This helps, but it is not enough. A person may learn to check sources and still share an inflammatory claim because sharing signals loyalty. He may understand probability and still reject a conclusion because accepting it threatens his group. She may recognize manipulation and use it deliberately because the institution rewards results, not intellectual hygiene.

The problem is therefore cultural. What receives prestige? Who is allowed to change position without humiliation? Is a public correction interpreted as evidence of integrity or as weakness? Does the organization preserve records of failed predictions? Can bad news travel upward without destroying the messenger? Does the audience reward the person who represents an opponent accurately, or the one who produces the most satisfying caricature?

Every institution answers these questions, even when it refuses to discuss them.

Sophisticated stupidity

The most dangerous idiocy is not always crude. Crude errors can be visible and limited. Sophisticated idiocy possesses vocabulary, credentials, data, software, legal departments, and the capacity to manufacture plausible deniability.

A model can be mathematically elegant and socially foolish. A risk score may predict a narrow outcome while hiding how its use changes behavior. An economic indicator may be measured precisely while excluding unpaid care, ecological depletion, or the distribution of loss. A performance metric may improve because employees have learned to produce the appearance of the target. A scientific field may generate statistically significant results while incentives select novelty over replication.

In each case, the intelligence inside the model is real. The stupidity lies in reification: treating a useful representation as if it were the reality itself.

This is why technical education alone cannot defeat idiocracy. Expertise increases the power of good judgment and the power of rationalization. The same person who can discover a subtle flaw in someone else's model may fail to notice the assumptions protecting his own status. The capacity to generate reasons is not identical to the willingness to be corrected by reasons.

Kruger and Dunning's famous studies are often reduced to a meme in which "stupid people think they are smart" (Kruger and Dunning, 1999). That simplification misses the more general warning. Accurate self-assessment requires some of the same competence as the task being assessed. The problem is not a permanent class of fools. It is the structural difficulty of seeing what our present tools do not allow us to see. Experts are protected in familiar domains and exposed when they travel outside them with the confidence their status has trained others to accept.

A genuine anti-idiotic culture therefore resists the fantasy that intelligence can be certified once and then trusted everywhere. Competence is local. A brilliant cryptographer is not automatically a wise epidemiologist. A successful investor is not automatically a philosopher of human value. A

Nobel Prize is evidence of extraordinary contribution in a domain, not a sacrament that transforms every future opinion into knowledge.

This boundary is not disrespectful to expertise. It is what makes expertise trustworthy. Authority grows more legitimate when it says where it ends.

When debate becomes a status sport

Modern public discourse often treats debate as a competitive performance. The vocabulary reveals the purpose: destroy, crush, own, dismantle, humiliate. The opponent becomes material for displaying intelligence before an audience. The audience does not need to understand the problem; it needs to recognize the winner.

A contest can sharpen arguments. Adversarial review is indispensable in science, law, security, and engineering. The pathology begins when the objective shifts from testing claims to defeating persons. Once identity is at stake, every correction becomes a concession, every nuance becomes weakness, and every partial agreement threatens the entertainment value of the conflict.

The structure favors simple positions. A careful answer needs distinctions. A theatrical answer needs confidence. A person who says “the effect depends on the population, time horizon, and definition of harm” appears evasive beside someone who declares “this proves everything.” The debate format can therefore select against the cognitive behavior required by the subject.

This is not an accidental defect of social media. Polarization can be profitable. Outrage increases attention, attention produces data and advertising revenue, and visible conflict encourages return. Research has found that users systematically overperceive moral outrage online, contributing to inflated beliefs about hostility between groups (Brady et al., 2023). The platform does not need to invent hatred. It needs only to make the most emotionally legible expressions more visible than the quieter distribution of ordinary human ambivalence.

Elites also benefit from division. “Divide and rule” is not merely an ancient slogan; it is a recurring technique through which rulers weaken coalitions that

might contest them. Contemporary division is often produced not by issuing direct orders, but by designing incentives. Groups are encouraged to fight over symbolic recognition while ownership, surveillance, market concentration, or administrative power remains outside the frame.

Yet division is not fully controllable. A story built to mobilize one group against another develops entrepreneurs, institutions, and believers of its own. The audience learns the emotional grammar and applies it to new targets, including its former sponsors. The person who wins legitimacy by manufacturing enemies eventually needs larger enemies. The contest escapes the original stage.

An Outfinitist culture must therefore change the status economy of disagreement. Accurate representation of an opponent should earn prestige. Publicly revising a claim should increase epistemic reputation. The strongest participant should be able to state what evidence would change their mind. A debate should be judged not only by who spoke persuasively, but by whether the shared map became more accurate.

This does not require artificial harmony. Some interests are genuinely incompatible. Some claims are false. Some actors negotiate in bad faith. Meta-rationality is not the ritual declaration that everyone has a point. It is the discipline of identifying the kind of conflict before choosing the method: factual error, difference of scale, competing values, unequal power, insufficient information, or strategic deception.

A culture that cannot distinguish these conflicts will alternate between naive reconciliation and permanent war.

Institutions that scale error

Institutions are cognitive machines. They collect information, compress it, assign authority, remember some events, forget others, and act at scales no individual can reach. They can correct personal bias or industrialize it.

Hierarchy is useful because not every person can participate in every decision. A hospital requires roles. A spacecraft requires command. A research program

requires allocation. Centralization can create coherence, speed, accountability, and the concentration of resources required for large projects.

The limit appears when information cannot travel in the same direction as power. People near the consequences see details that people near the center do not. If bad news threatens careers, reports become optimistic as they move upward. By the time a decision reaches the top, the organization may be responding rationally to an invented world.

The opposite arrangement also fails. Complete decentralization can fragment standards, duplicate costs, and make responsibility disappear into process. A network in which everyone can object but nobody must decide is not free; it is often governed by fatigue, informal status, or the person with the most time.

The question is not centralization or decentralization in the abstract. It is: which function needs coherence, which needs local knowledge, what is the propagation radius of an error, and how can a decision be appealed or reversed?

An airline centralizes safety standards and flight coordination while distributing observation among crews, engineers, controllers, and incident-reporting systems. The scientific community centralizes some methods and publication norms while distributing inquiry across institutions. The internet centralizes some protocols and naming functions while permitting enormous diversity at the edges. Resilience usually emerges from an architecture in which no single layer is asked to know everything.

Ashby's law of requisite variety offers a cybernetic formulation: a regulator needs sufficient variety to respond to the variety of the system it controls. A center that compresses a complex environment into a handful of indicators eventually loses what it needs to govern. It may compensate by imposing greater uniformity on the environment. This makes the dashboard cleaner and the society more brittle.

Artificial intelligence can deepen both possibilities. It can help institutions process local signals, translate across domains, detect anomalies, and simulate alternatives. It can also create the illusion that the center now sees everything.

More data does not eliminate framing. A system trained on recorded behavior may make the unrecorded less visible. A prediction may change the behavior it predicts. A recommendation may become an order because nobody wishes to accept the liability of disagreeing with the machine.

The anti-idiocratic institution is not the institution with no errors. It is the institution in which errors remain discoverable, bounded, contestable, and recoverable.

Equal rights, bounded authority

Opposition to idiocracy can quickly become contempt for people. That path is morally corrupt and strategically foolish.

Cognitive capacities differ. Education differs. Health, time, language, trauma, and cultural preparation differ. People are vulnerable to different manipulations. None of these differences cancels personhood. A society in which rights depend on performance will eventually allow the powerful to define performance so that they remain entitled to power.

Equal dignity, however, does not imply equal authority in every decision. We do not ask a random vote to determine the dose of an anesthetic. We do not assign bridge design by popularity. Complex societies require delegated competence.

The legitimate distinction is between **dignity**, which is equal, and **authority**, which is conditional.

Authority should be linked to a domain, evidence of competence, declared conflicts of interest, responsibility for consequences, a time horizon, and a path of removal. The more irreversible the decision and the more people who bear its costs, the stronger the justification and safeguards must be.

This principle also limits the intellectual elite. A person may understand the public's cognitive vulnerabilities and conclude that manipulation is justified because the goal is good. That conclusion deserves suspicion. If citizens must be deceived in order to accept a policy, perhaps the policy is wrong; perhaps the explanation is poor; perhaps the decision cannot be democratic at the

required speed; or perhaps a temporary emergency power is necessary. These possibilities should be named. Secretly substituting manipulation for consent destroys the evidence by which legitimacy could be judged.

There are emergencies in which authorities must act before broad understanding is possible. The Outfinitist response is not denial. It is to make the exception explicit, narrow its domain, record the reasons, limit its duration, and create review after the urgency passes. A temporary concentration of power becomes idiocratic when it presents itself as permanently exempt from correction.

The distinction between equal dignity and bounded authority also protects those who cannot understand every system shaping their lives. A person should not need a doctorate in machine learning to contest an automated decision about employment, credit, medical care, or liberty. The burden of explanation and appeal belongs to the institution that deploys power, not to the person subjected to it.

The first rule of the new culture

The first rule is simple to state and difficult to embody:

No person, model, institution, majority, market, machine, or movement is exempt from correction.

The rule does not say that all claims are equally uncertain. It does not deny expertise, law, or decisive action. It establishes a higher-order constraint: authority must remain connected to mechanisms capable of detecting and repairing its errors.

For this rule to become cultural, correction must not be merely permitted. It must be rewarded. A leader who changes course before catastrophe should gain more legitimacy than one who preserves an image of consistency. A researcher who publishes a failed replication should contribute to knowledge, not disappear from the prestige system. A company that closes an unsafe line of business should be able to demonstrate strength through restraint. A

citizen who leaves an ideological tribe after confronting evidence should be received as a learner, not used as a trophy.

A movement cannot produce these behaviors through instruction alone. People under pressure return to trained reflexes. If every previous correction brought humiliation, “be open-minded” will not overcome the nervous system. New reflexes require repeated practices, social protection, and examples with visible consequences.

This is why truth needs culture. Culture determines whether an idea remains a sentence or becomes a capacity.

The next question is who currently possesses the power to cultivate such capacities—and why powerful groups so often choose the opposite. To answer it, we must look at elites and at the monsters they repeatedly imagine they can feed without being eaten.

Elites and the Monster They Feed

Every anti-idiotic project eventually reaches the problem of elites. Avoiding the word does not eliminate the reality. Power concentrates around capital, rare knowledge, infrastructure, institutions, networks, and the capacity to organize. Even movements that announce the end of elites produce spokespersons, administrators, gatekeepers, interpreters, and informal courts. The important question is not whether complex societies will have elites. They will. The question is whether those elites will treat their position as a natural right of predation or as a limited mandate to administer a world they did not create alone.

This book does not ask competent people to renounce influence. It asks them to see beyond the predator's morality. An elite may profit from public confusion for years and still be pursuing a losing strategy. By weakening education, degrading shared reality, rewarding humiliation, and converting insecurity into political energy, it cultivates the forces most likely to destroy property, continuity, legitimacy, and eventually the elite itself.

The invitation is not sentimental. It is addressed to self-interest capable of extending its time horizon.

***Homo homini lupus*: the sentence cut in half**

Homo homini lupus is often repeated as a law of nature: man is a wolf to man. The phrase appears to certify the predator's worldview. If human life is fundamentally war, then exploitation is simply realism practiced before one's competitors. Compassion is weakness, cooperation is naivety, and power becomes its own proof of merit.

The old sentence is more interesting than the slogan. In Plautus's *Asinaria*, the line is: *Lupus est homo homini, non homo, quom qualis sit non novit*—a human being is a wolf, not a human being, to the person whose character he does not know. The scene concerns distrust between strangers. The line does not describe an immutable biological essence (Plautus, 1912). It describes a relation produced by uncertainty and the absence of trust.

Thomas Hobbes later made the expression famous, but in the dedicatory letter to *De Cive* he placed it beside another proposition: human beings may also be, for one another, like gods (Hobbes, 1998 [1642]). Both claims can become true under different conditions. We are capable of predation, but also of cooperation among strangers, sacrifice for abstract rules, care for people we will never meet, and institutions that survive their founders. Civilization did not remove the wolf. It changed the incentives, expectations, and protections surrounding the encounter.

Economic elites who invoke *homo homini lupus* as permission to exploit are therefore reading only half the human repertoire. Worse, they may turn the half they select into a self-fulfilling prophecy. Treat employees as adversaries and they will conceal bad news, protect themselves, and leave when they can. Treat the public as extractable attention and it will answer with cynicism, addiction, or rage. Treat the state solely as prey and public institutions will fill with less sophisticated predators. A strategy can generate private profit and systemic degradation at the same time.

The predator calls this realism because the damage often arrives after the reward. Outfinitism calls it a model with an omitted time horizon.

Elites are inevitable; legitimacy is not

Denying elites does not democratize power. It makes power harder to see. Jo Freeman described the “tyranny of structurelessness” in movements that rejected formal organization while allowing friendship networks and informal status to govern without accountability (Freeman, 1972–1973). The lesson extends well beyond social movements. Whenever explicit structures

disappear, tacit structures remain. The person who knows how decisions are really made gains power without having to defend the rules.

Societies need elites because no person can study every domain, operate every infrastructure, or evaluate every high-consequence risk. The cognitive division of labor is as real as the economic one. We need surgeons, pilots, cryptographers, engineers, judges, researchers, founders, administrators, and leaders capable of coordinating action under pressure. In systems with large externalities, a superficial decision can injure people who never participated in making it.

The legitimate role of elites is not to govern a biological caste called “the idiots.” It is to direct dangerous processes whose complexity and consequences require specialized competence. We control access to a reactor, not the dignity of the citizen. We license the pilot without turning pilots into a priesthood. We restrict who may prescribe a drug while preserving the patient’s rights to information, consent, second opinion, and appeal.

Real differences exist among people. Biology, health, nutrition, trauma, education, culture, time, language, and information environments affect cognitive performance. Some people understand abstractions more quickly; others notice human dynamics, practical risks, or local details that formal experts miss. Denying these differences in the name of equality is as unhelpful as converting them into permanent hierarchies. Moral equality does not require the fiction of cognitive identity. It requires that difference not become a license for abuse.

A legitimate elite has two obligations at once. It must protect high-risk systems from decisions driven by manipulation, ignorance, or collective impulse. It must also increase the capabilities of the people it serves, so that dependence does not become the elite’s business model. A good teacher does not need permanent pupils. A good physician does not cultivate illness. A legitimate leader does not require society to remain confused enough to regard the leader as indispensable.

This distinction matters in the age of AI. Expertise will become more productive, but also easier to simulate. Systems will generate confident

explanations, persuasive images, legal drafts, code, research summaries, and strategic advice. The public will need trusted intermediaries, yet those intermediaries will gain unprecedented capacity to manufacture the appearance of trustworthiness. The elite of the AI age must therefore be judged not only by what it knows or owns, but by the limits it makes enforceable around its own power.

The predator's morality

The predator's morality is seductive because it converts success into a complete explanation. You won because you saw the world as it is. Others lost because they refused realism. Wealth becomes a rough measure of general worth. Vulnerability becomes exploitable information. Externalizing costs is called efficiency until a law gives the cost a name.

There are situations in which hard strategy is necessary. A company must close a nonviable project. A state must defend itself. A negotiator cannot reveal every reservation price. A security team must sometimes restrict information. Outfinitism is not a childish morality in which conflict disappears through kindness.

The error begins when a local strategy becomes a total anthropology. Because one must sometimes fight, the world is declared to be nothing but war. Because an organization must sometimes dismiss employees, people become mere costs. Because secrecy can be necessary, transparency becomes foolish. Because competition produces useful discovery, every social good must be subjected to competition. A tool crosses its boundary and becomes a god.

The predator model has several structural weaknesses.

The first is the **cost of distrust**. Contracts grow longer, supervision more expensive, and useful information moves more slowly. In an organization where messengers fear punishment, bad news reaches the center only after it can no longer be concealed. The leader appears powerful while becoming epistemically blind.

The second is **adverse selection**. A leader who rewards theatrical loyalty and aggression attracts people skilled at simulating loyalty and hiding their own

ambitions. The predator does not build a stable court of servants. He builds an ecosystem of predators waiting for evidence of weakness.

The third is **time-horizon collapse**. Rapid extraction looks good in quarterly measures while consuming reputation, employee health, public trust, institutional competence, and ecological capacity. When the effects become visible, the decision-maker may occupy another office and describe the damage as macroeconomic context.

The fourth is **succession fragility**. Personalized power does not know how to die. The founder avoids building an institution that can function without him because indispensability has become proof of value. Heirs receive assets but not legitimacy. The struggle around the throne begins before the throne is empty.

The fifth is **political reaction**. People can tolerate large inequalities when they perceive contribution, intelligible rules, mobility, and reciprocal obligation. When wealth appears to result from capture, manipulation, and impunity, redistribution ceases to be merely fiscal policy. It becomes symbolic revenge. An elite that refuses every social obligation indirectly finances the movement that will demand its destruction.

The sixth is **dependence on the people it despises**. The predator model needs impulsive consumers, frightened voters, employees without alternatives, and audiences unable to compare claims. It therefore cultivates the very incapacity it cites as justification for domination. But a population maintained in confusion does not remain a passive instrument. It becomes a market for harsher demagogues, a reservoir of violence, and a permanent threat to the order that produced it.

The intelligent predator imagines that stupidity belongs to others. The more successful he becomes, the less likely anyone is to tell him that his model is eating its own conditions of survival.

The monster raised in the courtyard

History repeatedly shows elites borrowing the energy of radical movements and discovering too late that the borrower has become the owner.

The pattern is recognizable. Institutions are blocked. A population is anxious or humiliated. A leader who can mobilize resentment becomes useful to established interests. Traditional elites supply money, access, legitimacy, or a legal road to the center. They imagine that serious administrators will surround the agitator and use his popularity against their rivals. Once the movement enters the state, however, mass loyalty, propaganda, exceptional law, and organized violence disable the institutions meant to contain it.

Germany in 1933 remains the classic warning. Franz von Papen helped overcome President Paul von Hindenburg's resistance to appointing Adolf Hitler as chancellor. Conservative participation in the cabinet gave the arrangement respectability. Papen and his allies believed Hitler could be enclosed and controlled. Within months, the Nazi movement had coordinated, subordinated, or destroyed the institutions that were supposed to restrain it (Jones, 2005). The lesson is not that one decision alone caused Nazism, nor that every conservative shared the same aims. The lesson is narrower and durable: established actors can catastrophically misjudge a movement whose legitimacy comes from mobilization rather than from the rules those actors expect it to obey.

Fascist Italy followed a related logic. Liberal institutions, industrial interests, landowners, and parts of the monarchy saw Mussolini's movement as protection against disorder and socialism. Violence that appeared useful at the periphery became a method of national rule. The state did not merely employ the movement; the movement remade the state (Adinolfi, 2025; Di Rienzo, 2020).

In Iran, many secular and liberal opponents of the Shah entered a broad revolutionary coalition in 1978–1979, expecting that the common enemy could be removed before differences were resolved. The religious networks surrounding Ayatollah Khomeini possessed deeper organizational reach, a more coherent language of legitimacy, and institutions capable of surviving repression. After the revolution, pluralism narrowed rapidly, and many former allies were marginalized, imprisoned, or forced into exile (Keddie, 2010). Coalition against a center did not imply agreement about what should replace it.

During China's Cultural Revolution, Mao mobilized youthful zeal against alleged bourgeois and bureaucratic enemies. Red Guard formations attacked teachers, officials, cultural institutions, and one another. Established cadres who imagined they could direct ideological fervor discovered that public humiliation and factional violence had their own momentum. Many of the officials who participated in or accommodated earlier campaigns were themselves purged (Walder, 2014).

These cases are not interchangeable. Their ideologies, institutions, and scales of violence differ. They do not prove that every popular movement becomes totalitarian or that elites are innocent victims of forces they merely misunderstand. Often the elites help create the conditions, select the vocabulary, finance the channels, and legitimize the exception.

The common mechanism is the outsourced cultivation of irrational power. An elite encourages a population to suspend standards of evidence and restraint against a chosen enemy. It cannot then assume those standards will return when the enemy changes. People trained to experience humiliation as justice will eventually seek new objects. Institutions trained to reward loyalty above truth will replace competent servants with theatrical believers. Violence introduced as a tool becomes a language in which new actors can speak more fluently than their sponsors.

The monster does not betray its keeper. It follows the training.

Welfare as security infrastructure

The phrase "psychological well-being for all" may sound soft beside the language of power. It is, in fact, security architecture.

Chronically humiliated, isolated, frightened, and economically trapped populations are easier to mobilize through conspiracies, scapegoats, and promises of cleansing. A person with no credible path to dignity may accept borrowed grandeur from a nation, race, class, religion, or movement. A person whose life offers no agency may experience destructive participation as the first evidence that he matters.

Material security does not automatically produce wisdom. Prosperous societies can become narcissistic, polarized, and cruel. Education does not eliminate propaganda. Therapy does not replace justice. Yet systems that reduce avoidable desperation, offer meaningful participation, protect time for family and community, and make help accessible remove fuel from the most dangerous political markets.

For elites, investment in broad well-being is not charity added after the real economy has finished. It is maintenance of the cognitive and emotional commons on which contracts, innovation, law, and peaceful succession depend. A society in which large numbers of people cannot regulate fear or imagine a future becomes expensive to insure and easy to capture.

This does not justify paternalistic administration of everyone's emotions. The state or corporation should not decide which feelings are permissible. The goal is capability: access to education, health, mental-health support, reliable information, spaces of association, due process, and credible routes for improving one's position. A resilient population is not permanently calm. It can experience conflict and uncertainty without requiring an enemy to restore meaning.

The Outfinitist metacult proposed in this book begins at this point. It must serve not only those already able to articulate complex theories, but also the wider ecology that makes responsible intelligence possible. Its elite members should not ask merely how to preserve their rank. They should ask what social conditions make rank legitimate and non-suicidal.

The elite that deserves to survive

An elite deserves to survive when its continued influence is better for those affected than the available alternatives, and when that claim can be tested without depending on the elite's own propaganda.

Such an elite does not apologize for competence. It builds, invests, studies, coordinates, and takes responsibility. But it understands that superior performance in one selection process does not confer universal wisdom. It

distinguishes ownership from sovereignty, leadership from sanctity, persuasion from manipulation, and advantage from entitlement.

Its legitimacy rests on visible disciplines:

- it states the domain and duration of authority;
- it reveals material conflicts of interest;
- it preserves channels through which bad news can travel upward;
- it publishes failures and not only victories;
- it creates succession before succession becomes urgent;
- it accepts independent review where consequences exceed the organization;
- it protects rights that remain outside cost-benefit optimization;
- it invests in the capabilities of people who may eventually challenge it;
- and it treats voluntary self-limitation as evidence of strength rather than surrender.

What such an elite refuses to do becomes as important as what it accomplishes. A laboratory that can build a dangerous system and chooses not to release it may be demonstrating more competence than a competitor that mistakes acceleration for courage. A platform that declines a profitable manipulation channel may be preserving the conditions of long-term legitimacy. A wealthy founder who helps establish constraints that will also bind his successors is not betraying the elite. He is attempting to prevent wealth from decaying into a temporary exemption from society.

This is the wager offered to the intelligent parts of existing elites: the future need not belong either to resentful mass politics or to insulated technocracy. A third possibility is an elite culture that earns authority by enlarging agency, distributing competence, and limiting itself before a frightened public imposes cruder limits.

But culture cannot be changed by moral appeal alone. The present system rewards another story: the heroic founder, the winner who proves the rules by surviving them, the brand that must own attention, and the investor who can diversify while asking the entrepreneur to bet an undiversified life.

Moloch already has a culture. To build an alternative, we must first understand how that culture turns entrepreneurship into a casino with a moral mythology.

Moloch Already Has a Religion

Cultures do not leave meaning vacant. When traditional religion loses authority in part of a society, other institutions inherit its functions: they explain who the elect are, what sacrifice proves devotion, which failures reveal moral weakness, and what form salvation takes.

For part of the technological world, startup culture has become such a system. It has an origin myth: the founder sees what the crowd cannot. It has ascetic disciplines: sleepless nights, personal risk, mobility, and the renunciation of ordinary life. It has a liturgy: the pitch, the demo day, the funding round, the launch. It canonizes saints retrospectively and calls modest ambition a heresy. Its salvation is measurable as an exit, a unicorn valuation, or domination of a category. Its moral theology completes the circle: the winner proves by winning that he deserved to win, while the loser is encouraged to translate a statistical distribution into a lesson about character.

This culture has produced real innovation. It is not a fraud to be abolished. It has mobilized capital for improbable ideas, allowed small teams to challenge incumbents, and accelerated technologies that changed ordinary life. Precisely because it is productive, it deserves serious criticism. Institutions can create public goods while telling false stories about the distribution of merit. They can expand the frontier while consuming explorers in a race that no participant fully chose.

Moloch does not need a temple. It is enough for every firm to believe that it must grow before the next firm, capture the customer relationship, become the standard, and convert alternatives into dependencies. Startup culture is therefore more than a financing sector. It is a pedagogy of success, identity,

and worth. Any culture built around Outfinitism will have to coexist and compete with that pedagogy.

The mathematics of rarity

Millions of companies are created and operated in the hope of surviving, growing, and becoming recognizable. In the European Union alone, approximately 33.5 million enterprises were active in 2024. Only about 55,000—roughly 0.2 percent—were large enterprises, yet they generated 51.3 percent of net turnover. Micro and small enterprises accounted for 99 percent of the total number of firms (Eurostat, 2025).

At the global level, the World Intellectual Property Organization estimated that about 93.2 million trademark registrations were active in 2024, while new trademark applications covered approximately 15.23 million classes (WIPO, 2025). These numbers demonstrate that legal systems can store an immense number of distinct signs. They do not demonstrate that human attention can sustain tens of millions of dominant brands.

A trademark is a legal identifier and right. A brand is a social position: a recognizable association among a name, a promise, an experience, and a community of expectations. The legal registry can contain millions of names. Human memory cannot give all those names equal salience. Search engines, recommendation systems, distribution networks, advertising markets, and network effects concentrate visibility further.

A society can contain millions of useful businesses. It cannot contain millions of equally dominant brands in the same limited space of attention. Dominance is relational. To lead a category is to occupy a position that most competitors cannot occupy at the same time.

This elementary constraint is rarely made emotionally vivid to founders. They are told that a good product, disciplined execution, persistence, and a compelling story can build a global brand. All may be important conditions. None makes attention infinite. Even in an imaginary world where every startup had a good product, competent founders, ethical investors, and

excellent execution, all could not become the best-known company in the same category. The definition of the position implies scarcity.

The conclusion is not that most companies are useless. A local bakery, a security consultancy, a software laboratory, a medical practice, or an agricultural cooperative may create durable value without becoming a global symbol. The error arises when financing and prestige systems treat exceptional domination as the normative objective and classify smaller, stable, socially useful outcomes as failures of imagination.

Survival itself is selective. The United States Bureau of Labor Statistics followed private-sector establishments born in March 2013 and found that 34.7 percent were still operating in March 2023 (U.S. Bureau of Labor Statistics, 2024). This is a statistic about establishments, not a direct estimate for venture-funded startups. It does not reveal how many became profitable, remained small, changed ownership, or compensated founders for their opportunity costs. It does establish that ten-year continuity is already far from automatic.

The venture-funded distribution is more asymmetric. Robert Hall and Susan Woodward examined the undiversifiable risk borne by entrepreneurs and reported that almost three quarters of founders in their sample received nothing at exit, while a very small number received enormous payoffs. The mean looked attractive because the extreme tail pulled it upward. For a moderately risk-averse entrepreneur evaluating the distribution at the beginning, the certainty-equivalent value was only slightly above zero (Hall and Woodward, 2010).

This is the structure of a game with rare, very large rewards. Entrepreneurship is not identical to roulette. Founders can create technology, knowledge, relationships, employment, and social value. Their decisions affect probabilities. The outcome is not pure chance. Yet control is incomplete, and the culture systematically exaggerates it.

Timing, regulation, geography, health, investor decisions, infrastructure, competition, interest rates, and luck can send similar teams into radically different financial futures. After selection, the winner is invited to explain

why victory followed inevitably from a handful of principles. The retrospective story converts contingency into method. It is useful as inspiration and dangerous as evidence.

Why entrepreneurship can resemble a casino

Casinos do not sell probabilities alone. They sell simulated control, proximity to victory, status, emotional intensity, and a story in which the next attempt can transform the player's identity. Entrepreneurship can recruit similar mechanisms even while producing real goods.

The first mechanism is **overconfidence**. In experiments by Colin Camerer and Dan Lovallo, participants entered competitive markets in excess, especially when outcomes depended on an ability they believed they possessed (Camerer and Lovallo, 1999). Using survey data from eighteen countries, Philipp Köllinger, Maria Minniti, and Christian Schade found that subjective belief in one's entrepreneurial skills was among the strongest predictors of entry. At the country level, reported confidence was negatively associated with approximate survival prospects for nascent businesses (Köllinger, Minniti, and Schade, 2007). Confidence enables action, but it is not a neutral probability estimate.

The second mechanism is the **near miss**. In an experimental slot-machine task, near-winning outcomes were experienced as less pleasant than clear losses, yet increased the desire to continue—particularly when players felt personal control over the game. Neural responses to near misses recruited reward-related circuitry also engaged by wins (Clark et al., 2009).

Startup life is full of near misses with genuine informational content: the investor who says, "Return after a little more traction"; the acquisition conversation that reaches legal review; the admired product that users do not purchase; the funding round lost in its final stage; the large customer who begins integration and withdraws. These events may indicate progress. They may also sustain a loop in which every approach to success is interpreted as proof that one more sacrifice will make success inevitable.

The third mechanism is **winner availability**. The casino displays the jackpot. The entrepreneurial ecosystem displays the unicorn. Conferences invite survivors. Journalism narrates the multiplication of wealth more readily than the full distribution of cohorts that pursued similar advice. Founders encounter constant evidence that victory is possible and little living evidence of the financial, relational, and psychological costs borne by the much larger population of attempts.

The fourth mechanism is **identity fusion**. A gambler may lose money while retaining a separate profession and social name. A founder's company often becomes his name, tribe, calendar, moral project, and explanation for years of sacrifice. Closing the company is not experienced as ending an experiment. It becomes a verdict on the self. Fusion can generate extraordinary perseverance and extraordinary blindness.

The fifth mechanism is **risk asymmetry between investor and founder**. A professional investor can hold positions in dozens of companies. The portfolio is the strategy. The founder concentrates years, reputation, relationships, health, and often personal capital in one attempt. When the investor says, "We are taking this risk together," the pronoun conceals two different geometries of risk.

The sixth mechanism is **variable reinforcement**. Progress in an uncertain venture arrives irregularly: a customer appears after months of silence, a product suddenly spreads, a respected person responds, a prototype works, a team member joins. Intermittent rewards are powerful because the next one cannot be predicted. The system can maintain intense effort long after a smoother feedback environment would have encouraged reassessment.

The analogy should not be stretched into contempt. Societies need people willing to explore low-probability possibilities. Many technologies would not exist if every decision were optimized for a safe salary. A failed company can produce knowledge, reusable code, skilled employees, and relationships that later create value. Entrepreneurship can increase autonomy, meaning, and positive well-being.

A meta-analysis covering 319 samples, 94 studies, as many as 6.7 million observations, and 82 countries found that entrepreneurs reported, on average, higher levels of some forms of positive well-being than employees. The relationship depended on the dimension measured and on institutional conditions, particularly the rule of law. In weaker institutional environments, entrepreneurs could experience more psychological distress; in stronger ones, the relationship could reverse (Stephan, Rauch, and Hatak, 2023).

The problem is not an entrepreneurial personality that must be cured. It is the design of the ecosystem. A system can benefit from distributed exploration without telling every explorer that he is destined to become king. It can honor experiments and protect professional continuity when they end. It can separate courage from the duty to pretend that probabilities do not exist.

A socially useful casino, if we retain the provocative metaphor, would be one in which society benefits from many attempts while participants understand the distribution, retain portable skills and reputation, and are not abandoned when the individual outcome is not a jackpot. Parts of the present ecosystem already do this. Its dominant mythology still prefers the chosen founder.

The brand as monopoly over a story

The conventional brand concentrates several kinds of power. One entity defines the name, promise, standards, access, reputation, and distribution of surplus. Employees, users, partners, and communities contribute to the story, but the legal right to decide what the story means remains centralized.

This concentration can be valuable. It permits coherent design, decisive investment, rapid quality control, and clear responsibility. Centralization has functions that should not be destroyed merely because decentralization has become fashionable. A medical device, aircraft, or payment system cannot resolve every safety decision through open-ended community sentiment.

The same concentration can create dependency. Reputation produced by many becomes an exclusive asset of a few. The platform is advised to own identity, data, audience, and distribution. Every startup is encouraged to build

a fortress, though only a tiny fraction can finance one. Thousands of organizations then purchase separately the same expensive ingredients: recognition, trust, discovery, legal capacity, technical infrastructure, and access to capital.

The tragedy is not only waste. It is the privatization of collective learning. When the company disappears, its community fragments, its experiments are forgotten, and the people who created its credibility must begin again under a new name. The legal entity has been treated as the sole container of meaning even when the meaning was generated by an ecosystem.

A **decentralized brand** separates some of these powers. A shared identity can be carried by autonomous projects. Standards can be administered transparently, and permission to use the common name can depend on verifiable conduct. Contributions can remain portable. Projects can keep their own revenues, ownership, and operational responsibility while part of the shared reputation and infrastructure remains available to the wider network.

This does not mean that every member votes on every product decision, that property becomes indistinct, or that all work receives equal compensation. A laboratory needs accountable investigators. A product needs a team. A venture that accepts capital acquires specific obligations. Decentralization concerns the possibility that these entities share a mission, a trust layer, and a market for collaboration without being absorbed into a single corporation.

The decentralized brand is not anti-company. It is an architecture in which the company is a node rather than the owner of the entire possibility space.

From the investor's portfolio to the community's portfolio

Investors manage uncertainty through portfolios. Entrepreneurs could gain access to a comparable logic at the level of community.

A decentralized brand can support several projects, move people among them, preserve reputation when one entity closes, and share tools,

distribution, legal templates, research, and institutional memory. The failure of a node need not erase the professional identity of its contributors.

Imagine a founder who spends four years building a sound technology for a market that fails to mature. Under the conventional model, the company closes, the team disperses, and the founder explains the period defensively in future interviews. In a decentralized network, contributions can be attested, components can migrate to other projects, and the founder can enter a new role without feeling that he has abandoned his tribe. Investors can continue to follow people and knowledge, not only the vanished legal entity.

This portability lowers the identity stake of each attempt. Paradoxically, it can make experimentation more honest. When closure does not mean social exile, a founder can stop a nonviable project earlier. When a well-documented failure increases epistemic reputation, the postmortem ceases to be a confession of shame. The community learns from the distribution, not only from its champions.

A shared brand can also respond to the scarcity of attention. Rather than forcing each project to purchase credibility from zero, a common promise can function as a filter. Projects that use the identity accept public standards, and audiences may transfer part of their trust across nodes. Not every project becomes a dominant brand. It does not need to. The umbrella becomes symbolic infrastructure within which projects remain autonomous.

This changes the meaning of success. A small project may be valuable because it solves a local problem, tests a method, trains people, or produces knowledge that another node can use. Growth remains possible, but growth is no longer the only proof of life.

The conditions of real decentralization

“Decentralized” is easy to print and difficult to prove. A company can distribute tokens while retaining control of code, data, trademarks, treasury, and communication. A community can hold votes understood only by insiders. A founder can proclaim node autonomy while keeping an informal veto through prestige, funding, or access.

Real decentralization is behavioral. When one node fails, the network continues. When the center makes an error, a correction path exists. When a project earns revenue, distribution rules were known before success made negotiation explosive. When a member leaves, her contributions do not disappear. When conflict arises, appeal and mediation are available. When disagreement becomes irreconcilable, a fork can continue without destroying the common memory.

Standards are also necessary. Without them, a shared identity becomes a logo without a promise. Nodes must be suspendable when they violate rights, fabricate results, or use common reputation for fraud. Yet the power to enforce standards should be separated from the power to capture all assets. Otherwise the standards body becomes the network's owner in practice.

Technology can help: portable identity, contribution registries, cryptographic proofs, explicit contracts, auditable treasuries, and AI-assisted translation, synthesis, or mediation. None automatically creates legitimacy. A blockchain can preserve an injustice with the same precision with which it preserves a contribution. AI can widen deliberation or amplify the voice of the center. The decisive layer remains social architecture.

A useful test is to ask what happens when the founder disappears, when the most profitable node defects, when the rules produce an unpopular outcome, and when two sincere groups no longer agree. A system is not decentralized because it functions while everyone trusts the center. It is decentralized when it can metabolize the loss of that trust without collapsing into chaos or re-centralizing through emergency.

Founder health and the economy of exploration

Startup culture often glorifies stress as evidence of seriousness. Some challenges can increase skill, purpose, and performance. Research on entrepreneurial stressors distinguishes, however, between demands that can be integrated as meaningful challenges and obstacles that consume energy without producing development. Permanent uncertainty, arbitrary

bureaucracy, role conflict, precarity, and isolation can damage health without improving the venture.

Resilience is not unlimited exposure to suffering. A muscle grows through stress followed by recovery. Without recovery, injury follows. An organization that demands permanent strain and names collapse a lack of grit is not testing resilience. It is consuming people and selecting for those best able to hide the cost.

A metacult for entrepreneurs would treat mental health as infrastructure rather than as an optional workshop offered after burnout. It could provide peer circles, access to qualified professionals, education about risk and identity, and paths of continuity after a venture ends. It should not practice amateur therapy or use vulnerability as a recruitment technique. Its distinctive promise would be simpler: a founder may say “I no longer know” without causing the investor, team, and self-image to collapse at once.

The objective is not to abolish competition, centralized brands, or venture capital. They can coordinate large resources and accelerate work that would otherwise remain impossible. The objective is to add an economy of exploration in which attempts share reputation, infrastructure, memory, and care, while success does not require the destruction of every alternative.

In such an economy, the entrepreneur is neither an infallible hero nor an irrational gambler. She is an explorer operating under uncertainty and responsible for calibrating the story against the evidence. The investor is not merely a buyer of lottery tickets, but a designer of risk distributions. The community is not an audience for the worship of success, but a repository of knowledge and continuity.

Financial instruments alone cannot create this economy. It requires a different prestige system. The founder who closes an experiment honestly must be able to preserve reputation. The investor who publishes the complete distribution rather than only the champions should gain trust. The useful project that remains small must be recognizable as a success rather than a failure of ambition.

A culture must learn to hold propositions that appear contradictory only to a mind demanding one model for every scale:

The idea is improbable and worth trying. The person is valuable and the project may be stopped. The common brand matters and no center owns it entirely. Competition discovers and must be bounded. Centralization coordinates and must remain contestable.

To live inside such propositions, we need a more mature relationship with theories, risk, identity, and uncertainty. The path leads from rationality to meta-rationality—and from meta-rationality to Outfinitism.

Meta-Rationality and the Invention of Outfinitism

So far we have encountered institutions that reward error, elites that cultivate forces they can no longer control, and an economy of prestige that converts rare outcomes into moral verdicts. The common failure is **reification**: a useful partial model—the market, merit, voting, expertise, growth, tradition, the algorithm—is treated as reality itself.

The answer cannot be the abandonment of reason. Superstition, impulse, and arbitrary authority do not become wise merely because rationality has limits. The answer is a rationality capable of examining its frame, joined to a practice that keeps returning to the edge of models.

I use **meta-rationality** for the first. I propose **Outfinitism** for the second.

Their relationship is central to the whole book:

**Meta-rationality is a mature relationship with reason.
Outfinitism is the attempt to operationalize that maturity
through the systematic exploration of limits.**

Meta-rationality can exist without the word Outfinitism. Outfinitism cannot remain healthy without meta-rationality. Separated from it, a philosophy of limits would become one more rigid theory declaring where everyone else must stop.

After rationality, not against it

Rationality is among humanity's greatest social technologies. It helps us separate desire from evidence, follow consequences, compare hypotheses, and

correct intuition. Logic, measurement, and experiment have given us more medicine, more reliable infrastructure, wider freedom from arbitrary authority, and a language in which claims can be challenged without first challenging the speaker's rank.

To speak about the limits of reason is not to invite superstition back into the room. Only a sufficiently strong rationality can investigate the conditions under which its own methods work. A method that declares itself infallible has ceased to be rational and has become doctrine.

The problem begins when rationality is identified with the application of one explicit scheme. If objectives, data, and rules are fixed, calculation may be rigorous. But who selected the objective? What reality failed to enter the data? Who bears the error? What happens when the indicator begins to reshape the behavior it measures? What does "optimal" mean when two legitimate values cannot be maximized together?

A system can be rational inside its model and destructive in the world. A platform maximizes time on site and discovers that anger retains attention. A hospital optimizes throughput and removes the time required to listen. A university optimizes publication counts and produces fragmented papers that nobody has an incentive to integrate. A state optimizes security and transforms private life into a field of potential deviation.

The calculation has not failed. The frame was too small.

Classical accounts of cognitive development emphasized the ability to reason abstractly, systematically, and hypothetically. Adult life, however, presents problems without clean premises or unique answers. Relationships, politics, medicine, leadership, and ethics contain incomplete information, competing values, path dependence, and people who change their behavior when they discover the model being applied to them.

Over several decades, researchers have described capacities that respond to such conditions. Work on reflective judgment examines how people understand the uncertainty of knowledge and justify conclusions when evidence is incomplete (King and Kitchener, 1994, 2004). Studies of dialectical thinking examine the capacity to understand systems through

contradiction, transformation, and relation (Basseches, 1984). Traditions of postformal thought, cognitive-affective integration, and wisdom explore how mature judgment coordinates abstraction with context, emotion, and the limits of personal knowledge (Baltes and Staudinger, 2000; Labouvie-Vief, 1994; Sinnott, 1998).

This literature does not establish one universal staircase on which every human being can be ranked from “irrational” to “meta-rational.” Its concepts overlap imperfectly, its methods differ, and some developmental claims are stronger than the evidence warrants. Eeva Kallio has argued that postformal and dialectical theories do not form a single coherent model and that terms such as integrative thinking may be more prudent (Kallio, 2011). That criticism belongs inside the concept, not outside it. It would be absurd to turn research about the limits of frames into a new psychological aristocracy.

A plausible core nevertheless remains. Some people, in some domains and conditions, become better able to:

- live and decide under uncertainty without inventing false certainty;
- compare evidence while keeping conclusions proportional to it;
- see conflict between frameworks without concluding that all opinions are equal;
- use emotion as information without treating it as the final judge;
- separate identity from theory, so that revision does not become psychological death;
- and choose among models according to scale, purpose, and consequence.

I call this family of capacities meta-rationality.

Rationality reasons within a frame. Meta-rationality can also observe the frame. Rationality may demonstrate that a strategy maximizes a specified objective. Meta-rationality asks why that objective is legitimate, when the metric will corrupt behavior, which other objectives are being sacrificed, and what signals should trigger a change of model.

This is not relativism. The claim that models are limited does not place a flat Earth and modern geodesy on equal footing. Some models explain more, predict better, cohere with more evidence, and survive more severe testing.

Meta-rationality preserves these differences. It refuses only the additional step by which a powerful model in one domain becomes a complete description of existence.

Superficial relativism says, “Everyone has a truth.” Meta-rationality says, “Reality constrains us, but our access to it is mediated by instruments, language, position, and purpose; therefore justifications can be compared, and confidence should remain proportionate to evidence.” We do not abandon truth. We abandon the fantasy of private ownership over it.

Nor does contradiction become automatically profound. Two theories making incompatible predictions under the same conditions cannot both be true in the same sense. Sometimes one is wrong. Sometimes the apparent contradiction comes from scale. Newtonian mechanics remains extraordinarily useful for many engineering problems although it is not a complete fundamental theory. Market prices can coordinate dispersed information in one context and be grotesquely inadequate for valuing the care of a child or the irreversible destruction of an ecosystem.

The meta-rational art is to diagnose the conflict: factual error, difference of scale, difference of purpose, competing values, insufficient information, strategic deception, or a system that changed after the model was built. Only then can we decide whether to select, combine, alternate, revise, or create a new model.

Maturity without a psychological aristocracy

Meta-rationality requires at least three forms of maturity: epistemic, emotional, and social.

Epistemic maturity is the ability to calibrate a claim. “I know,” “the evidence suggests,” “I assume,” “I hope,” and “I cannot yet distinguish” are different states. A mature person does not use the language of certainty to manage anxiety, nor hide a moral choice inside a technical result.

Emotional maturity is the ability to survive revision. A theory may be tied to a career, a community, and a self-image. When evidence threatens it, the nervous system reacts before the argument. Meta-rationality is not cold

performance. It is the capacity to remain present while one's identity loses a support. Sometimes the greatest limit of reason is not logic but shame.

Social maturity is the ability to cooperate with people who do not share an entire worldview. It does not require tolerating cruelty, fraud, or permanent bad faith. It requires distinguishing legitimate disagreement from existential threat. A mature group can preserve a common mission while maintaining competing theories about how to pursue it.

These capacities do not develop together. A researcher may possess sophisticated epistemology and intense emotional fragility. An empathetic leader may reason poorly about statistics. A spiritually deep person may be flexible in relationships and rigid in doctrine. A successful founder may handle technological uncertainty beautifully and interpret criticism of the business as betrayal.

No one receives a permanent badge of "meta-rational." It is a situated practice, not a superior species.

This warning is particularly important for an anti-idocracy movement. The moment members imagine that joining proves their maturity, the movement begins selecting for vanity. The moment "meta-rational" becomes a social rank, people will learn to perform ambiguity, quote complexity, and display nuance while protecting ordinary interests. Sophisticated language will again become camouflage.

The test is behavioral. Can I name what would change my position? Can I distinguish a threat to my model from a threat to my person? Can I represent a rival view so its adherent recognizes it? Can I act decisively without pretending the decision was inevitable? Can I accept a boundary that reduces my own advantage?

A metacult must reward these behaviors rather than the vocabulary that describes them.

Manipulating theories rather than people

In technical language, to manipulate may simply mean to handle, transform, and combine objects. Meta-rationality asks us to manipulate theories in this sense: to hold them as instruments rather than allowing them to hold us.

An organization can be examined as a market, hierarchy, network, moral community, or cybernetic system. Each frame reveals something different. The market highlights prices and incentives. Hierarchy highlights authority and coordination costs. Network analysis highlights position, connection, and diffusion. The moral-community frame highlights trust, obligation, and identity. Cybernetics highlights feedback, variety, control, and stability.

Saying that several frames are useful does not license undisciplined mixing. For a concrete decision, we should name the question, the model, the information it excludes, and the second model used to check the result. The manager who sees only incentives may destroy loyalty. The activist who sees only community may ignore costs. The engineer who sees only control may treat the human being as a component. The economist who sees only revealed preference may rename addiction as choice. The psychologist who sees only trauma may understate strategy and responsibility.

Meta-rationality creates freedom from ideology through **controlled plurality**. We need not select one theory of everything. We may use markets where distributed discovery works, central standards where incompatibility would be dangerous, community governance where the good depends on local trust, and rights where no aggregate benefit should erase the person.

This flexibility is not indecision. It is closer to the skill of an experienced navigator. The map is not rejected because it omits every stone. Several maps are carried because terrain, weather, altitude, and destination demand different representations. The navigator knows that choosing the wrong map can be more dangerous than making a small calculation error on the right one.

The ethical contrast is also important. A culture that teaches people to manipulate theories can reduce the demand to manipulate persons. When a leader has only one model, resistance appears irrational and must be

overcome. When several models are available, resistance may be information: the metric is wrong, the scale has changed, the incentive is perverse, the person bears a cost invisible to the center.

Outfinitism begins by turning this flexibility into a disciplined concern with boundaries.

From a psychological posture to an operational program

The term **Outfinitism** is new. It is not offered as a completed science, a discovery of a universal law, or a replacement for existing disciplines. It is a proposal for organizing inquiry around limits.

The name points outward from the finite. We never stand outside all limits, but we can move toward the boundary of a model, examine what lies beyond its domain, and build a better relation between maps. The aim is not transcendence into omniscience. It is improved navigation by minds that remain finite.

Outfinitism may be understood in three ways.

First, it is an **operational philosophy**. It asks people and institutions to treat boundary conditions, uncertainty, omitted variables, reversibility, and the distribution of error as central to judgment.

Second, it is a **metadisciplinary research program**. It encourages comparison among the ways disciplines identify limits: impossibility theorems, error bounds, domain shift, externalities, cognitive constraints, institutional capture, moral prohibitions, and more.

Third, in a specific field, it can become an **internal science of limits**. Machine learning already studies generalization, robustness, calibration, and distribution shift. Economics studies externalities, market power, information asymmetry, and incomplete contracts. Medicine studies contraindications, heterogeneity of treatment effects, diagnostic uncertainty, and trade-offs. An Outfinitist practice does not replace these sciences. It makes their

limit-studying components more explicit, connected, and institutionally consequential.

The difference between a decorative caveat and an Outfinitist limit is what happens next. A report may mention that data are incomplete and then permit the model to govern as if they were complete. Outfinitism asks whether uncertainty changes the authority of the recommendation, whether affected people gain a right of appeal, whether deployment should be staged, and who must monitor the boundary.

A limit becomes operational when it changes design.

Families of limits

The word *limit* covers different phenomena. Treating them as identical would reproduce the simplification Outfinitism opposes. Several families are especially important.

Epistemic limits: what we do not know

Some uncertainty comes from missing data, measurement error, or unresolved mechanisms. More information may reduce it. Other uncertainty is produced by an open system: new technologies, strategies, and institutions change the environment faster than stable frequencies can be estimated.

The first task is calibration. A model should not merely produce an answer. It should communicate what type of uncertainty surrounds the answer and under which conditions confidence degrades. The second task is political: authority should shrink or become more reversible as uncertainty and consequence increase.

Unknown does not mean unknowable. It means that a claim about present knowledge should not be inflated into a statement about reality itself.

Formal limits: what a system cannot derive from within itself

Mathematics and computer science contain some of the clearest demonstrations that power and limit can coexist. Gödel's incompleteness theorems established profound constraints on sufficiently expressive formal

systems: under standard conditions, consistency and completeness cannot both be obtained from within the system in the way a naive dream of total formalization might demand. Turing showed limits on general computation through the undecidability of the halting problem (Gödel, 1931; Turing, 1936).

These results should not be turned into mystical slogans. They do not prove that every claim is undecidable, that human intuition escapes all formal constraint, or that logic has failed. They show something more disciplined: even the most exact intellectual constructions can contain boundaries that are not repaired by simply applying more of the same procedure.

For Outfinitism, the cultural lesson is modest. Precision does not abolish finitude. A formal system can be extraordinarily useful without being the final container of truth.

Predictive limits: when the world changes around the model

A model fitted to one distribution may fail under another. People respond to predictions, adversaries adapt, policies alter incentives, and rare events dominate consequences. The more powerful a prediction becomes, the more likely it is to modify the system it describes.

This is visible in finance, epidemiology, security, and AI. A fraud detector changes fraud. A ranking algorithm changes content production. A predictive-policing system changes where police look and therefore what data the next model receives. A labor-market recommendation changes which skills people acquire.

The relevant question is not only whether the model was accurate during validation. It is what feedback loop deployment creates.

Metric limits: when the proxy becomes the objective

Goodhart's law is often summarized as: when a measure becomes a target, it ceases to be a good measure. The exact formulations vary, but the mechanism is familiar. Once rewards depend on a proxy, people optimize the proxy and discover ways to separate it from the underlying good.

Publication counts replace knowledge. Click-through replaces informed attention. Test scores replace learning. Quarterly earnings replace institutional health. Safety checklists replace safe behavior. The metric may remain numerically precise while its semantic relationship with the objective decays.

Outfinitism does not reject measurement. It treats measurement as an intervention. Metrics need independent indicators, qualitative channels, periodic retirement, and observation of the behavior they induce.

Biological and psychological limits: the finite organism

Human attention, memory, self-knowledge, and emotional regulation are limited. We are not generic rational agents with occasional biases attached. We are embodied organisms whose reasoning depends on sleep, stress, nutrition, attachment, status, and social safety.

This fact can be used to protect or manipulate. Choice architecture may make healthy action easier; it may also exploit default effects for extraction. Persuasion may clarify; it may bypass deliberation. AI systems can support memory and translation; they can also personalize pressure with a precision no human propagandist could afford.

An Outfinitist institution treats cognitive vulnerability as a reason for rights and safeguards, not as permission for a more capable actor to govern invisibly.

Economic limits: the transaction outside its ecology

Markets can coordinate dispersed information and discover value through voluntary exchange. Their usefulness depends on conditions: meaningful choice, sufficient information, competition, internalized costs, and the ability to refuse.

When pollution is paid by someone else, when monopoly controls the alternative, when a buyer cannot evaluate risk, when power shapes preference, or when the effect is irreversible, the voluntary transaction does not exhaust the moral or economic problem. Economics itself contains the theories of these limits: externalities, public goods, market power, incomplete contracts, asymmetric information, and radical uncertainty.

The state also has limits. Central planning can coordinate infrastructure, redistribute risk, and enforce standards. It can suffer from missing local knowledge, rigidity, capture, and weak feedback. Communities can govern common resources through local rules, as Elinor Ostrom demonstrated, but communities are not automatically inclusive, competent, or capable of responding at scale (Ostrom, 1990, 2010).

Outfinitism does not search for the perfect institution. It searches for a fit among market, state, community, network, and hierarchy. A single economic theory of all goods is a map that has forgotten it has edges.

Organizational limits: the center that can no longer see

Centralized organizations can act coherently and quickly. As they grow, the center receives reality through layers of reporting. Each layer compresses, translates, and protects interests. Good news travels upward more easily than bad news. Standard metrics allow comparison by deleting context.

W. Ross Ashby's law of requisite variety suggests that a regulator requires sufficient variety in its responses to manage the relevant variety of disturbances (Ashby, 1956). It is not a formula that settles politics. It offers a powerful intuition: a center with a narrow repertoire cannot govern a highly diverse environment without either losing control or brutally simplifying the environment.

Polycentric systems distribute decisions among units capable of observing, coordinating, and correcting one another. They can experiment in parallel and prevent one error from becoming universal. They can also duplicate costs, produce local inequality, and fail to act collectively.

Resilience is therefore not maximal decentralization. A hospital network may centralize safety protocols and distribute service design. A software project may maintain a rigorously governed core and independent extensions. A decentralized brand may share a minimal constitution while permitting many autonomous projects.

The Outfinitist question is not merely, “Where is the center?” It is: **What error can the center propagate, what error can the periphery propagate, and which paths connect correction between them?**

Moral limits: what we choose not to optimize

Some limits arise not from incapacity but from the decision not to use a capacity. We can manipulate more effectively than we should. We can surveil more than is legitimate. We can create dependency even when the user clicks “agree.” We can exploit a child’s cognitive vulnerability for profit. We can build a weapon because the engineering problem is interesting.

An Outfinitist morality does not reduce ethics to calculable outcomes. Consequences matter, but rights and prohibitions protect the conditions under which people remain participants rather than objects. Consent, privacy, appeal, exit, and protection from cruelty are constraints on power.

These limits are not simple. Privacy can conflict with preventing violence. Exit can be fictitious when alternatives do not exist. Rights can be formulated too vaguely or used strategically. Difficulty of application does not make the boundary meaningless. It means we need procedures of interpretation, appeal, and revision.

Outfinitism does not worship restriction. Some boundaries are forms of domination and should be removed. The questions are: whom does the limit protect, what risk does it reduce, who decides, and can the limit itself be challenged?

The limit audit

For Outfinitism to become operational, it needs a repeatable anatomy. When we use a theory, metric, technology, or institution, we can ask:

1. **Object:** What is the instrument trying to represent, predict, or coordinate?
2. **Domain:** For which population, scale, period, and environment does it work?
3. **Assumptions:** What must remain true for the answer to be useful?

4. **Omissions:** Which variables, values, and affected actors remain outside the frame?
5. **Failure mode:** How can the instrument continue producing a precise answer after the answer has become wrong or harmful?
6. **Transition signal:** What observation indicates that we are approaching the boundary?
7. **Error distribution:** Who benefits when the model is right, and who pays when it is wrong?
8. **Reversibility:** Can deployment be staged, stopped, appealed, or repaired?
9. **Alternative:** Which second model, institution, or procedure takes over at the boundary?
10. **Limiter's limit:** Who may declare that the boundary exists, and what constrains that authority?

The last question is political. A monopoly can call competition a systemic risk. A bureaucracy can insist that the rule permits no exception. A founder can claim that only he understands the vision. A real limit can still be used as a weapon.

Every limit therefore requires a limit on the limiter: mandate, transparency, appeal, expiration, audit, competing authority, or the possibility of exit. We cannot terminate the regress by inventing a perfect final guardian. We stop it procedurally, with safeguards proportionate to the stakes and a right to reopen the case.

Outfinitism does not promise a universal algorithm for choosing among theories. Such an algorithm would become the totality the project rejects. It can develop heuristics:

- When uncertainty is high and action reversible, run several experiments.
- When possible harm is severe and irreversible, require stronger evidence and preserve redundancy.
- When rival models imply similar action, begin in their zone of agreement.
- When the conflict is moral, do not conceal it inside a technical average.
- When a metric becomes a target, add independent indicators and audit side effects.

- When a center cannot represent local variety, delegate.
- When local nodes create global externalities, coordinate.
- When no option is safe, compare the risks of action and inaction explicitly.

Reversibility is a particularly valuable heuristic. Not every decision can be undone, but many can be staged through small batches, time limits, automatic pauses, and precommitted review. Reversibility purchases information. It converts abstract conflict into bounded experiment and lowers the amount of identity invested in being right.

The portfolio is another. Under deep uncertainty, maintaining several models, technologies, or institutions may be wiser than selecting a champion prematurely. Diversity has costs. Monoculture turns one error into total dependence.

The limit of Outfinitism

A doctrine of limits can become fascinated with impossibility. It can generate cynicism: if every model is finite, nothing deserves to be built. It can protect the status quo: every change has unforeseen consequences, so no one should act. It can become an aesthetic of nuance through which educated elites postpone decisions while retaining their position.

This is Outfinitism's internal danger. It must study not only when action exceeds knowledge, but when prudence becomes an excuse. Refusing to decide may be a decision in favor of whoever already benefits. Sometimes a small experiment is impossible, and the risk of inaction exceeds the risk of action.

Outfinitism is not the worship of the brake. It is the art of matching speed to visibility, stakes, and recovery capacity. An explorer who refuses every danger does not explore. An explorer who never checks the oxygen supply does not love the unknown; he loves the image of his courage.

The word *obsession* has a deliberate place here. Limits are easy to mention ceremonially and easy to forget when money, prestige, fear, or urgency arrives. Bridges need engineers preoccupied with loads and material fatigue. Institutions need people whose role is to ask where the model breaks.

That concern cannot be isolated in a compliance office that becomes a scapegoat. It must be distributed. The researcher publishes limits to generalization. The manager protects channels for bad news. The investor requests failure scenarios. The community rewards correction. AI is evaluated under new distributions rather than only familiar benchmarks. The leader declares the domain and duration of authority.

The person who believes he knows has less reason to search. The person who knows that complete knowledge is unavailable can continue searching without converting uncertainty into paralysis. This is the promise of Outfinitism: the edge of the map becomes not the point at which thought ends, but the place where the next map begins.

Method, however, is not yet culture. A limit-audit form can be completed mechanically. An audit can become compliance theater. An organization can write “we are fallible” on the wall and punish anyone who demonstrates an actual error.

For Outfinitism to matter, the study of limits must enter prestige, memory, relationships, and rhythm. People must feel that they still belong when they correct the group and do not lose their place when they abandon a theory. There must be things the community protects even against immediate advantage.

We must therefore enter the territory that religions explored long before modern institutions imagined that contracts and information were sufficient: ritual, belonging, and the sacred.

What We Choose to Hold Sacred

An idea becomes civilizational only when people can live inside it. Its arguments may be impeccable and still remain sterile if it offers no shared time, repeated gesture, form of care, source of prestige, or story through which a person can connect a finite life to something larger than the next transaction.

Religions built such worlds. They did not consist only of propositions about gods. They constructed calendars, meals, songs, initiations, confessions, prohibitions, networks of aid, duties among strangers, ways of mourning, and languages for suffering. Some of these mechanisms liberated people. Others controlled through fear, guilt, exclusion, and claims of monopoly over salvation.

A metacult must learn from both sides. It should understand the anthropological power of ritual without concealing its techniques of influence. It should create belonging without captivity. It should offer warmth without demanding the surrender of the mind that receives it.

This is why the question of the sacred cannot be avoided. A society that says nothing is sacred does not eliminate sacredness. It silently transfers it to money, nation, identity, growth, security, personal autonomy, the leader, or the machine. The practical question is not whether human beings will protect some values from ordinary negotiation. It is which values, with what limits, and under whose interpretation.

What religions understood about human beings

Religions did not cross centuries because doctrine alone defeated alternative doctrine. They created social worlds. They connected immense questions to small practices: what one does upon waking, with whom one eats, whom one asks for forgiveness, how a sick person is visited, how the dead are buried, how the birth of a child is recognized, and what may be said when catastrophe has no satisfying explanation. They translated metaphysics into time, architecture, music, memory, and care.

This fact proves neither the truth nor the falsehood of supernatural claims. It demonstrates that a worldview acquires historical force when it becomes inhabitable.

Anthropology has long treated ritual as more than an external expression of private belief. Roy Rappaport described ritual as a way in which participants publicly enact commitments and sustain conventions that no single individual fully controls (Rappaport, 1999). Repeated participation does not merely communicate an existing group. It helps produce the group by giving expectations visible form.

Research on intentional communities has examined how costly practices can reduce opportunism and signal commitment. In a study of eighty-three nineteenth-century American communes, Richard Sosis and Eric Bressler found an association between certain costly requirements and the longevity of religious communes. Laurence Iannaccone's influential economic account explained how strict groups may reduce free-riding by making superficial participation more expensive (Iannaccone, 1994; Sosis and Bressler, 2003).

The conclusion is not that severity creates virtue. Cost has a boundary. A modest contribution can signal seriousness and finance a common good. An extreme sacrifice can create captivity. Once a person has surrendered money, reputation, relationships, and years, acknowledging that the group is wrong becomes psychologically expensive. The sunk cost becomes a lever for demanding further sacrifice.

A metacult should therefore not maximize the cost of belonging. It should maximize the value created together while keeping the cost of departure low.

Commitment should be visible through contribution, not through mutilation of alternatives.

Religions understood another fact that purely informational movements often miss: people are not persuaded only by propositions. They are persuaded when someone arrives while their child is ill; when there is a table at which market value is irrelevant; when suffering can be narrated before it is converted into a metric; when grief is shared without being solved; when a person's name remains remembered after professional failure.

A movement that speaks of humanity's future but cannot sit beside an exhausted member does not yet possess an anthropology. It possesses branding.

Ritual itself is a powerful and ambiguous social technology. Field studies suggest that behavioral synchrony and shared physiological activation can increase cohesion and cooperation among large groups of strangers, though effects depend on context and may also strengthen boundaries against outsiders (Jackson et al., 2018). Cross-cultural work has associated belief in morally concerned, knowledgeable, punitive deities with more impartial allocation toward distant co-religionists, suggesting one route by which cooperation may have expanded beyond local relationships (Purzycki et al., 2016).

Again, such findings settle no metaphysical question. They show that ideas, practices, and institutions can change the radius of cooperation. The same capacities can support charity, education, trust, obedience, war, or exclusion. The Outfinitist question is not whether culture-forming mechanisms will be used. Every durable community uses them. The question is whether those mechanisms are visible, contestable, and compatible with the dignity of people inside and outside the group.

The Outfinitist metacult is proposed precisely as this explicit layer of design. It does not claim to have escaped ritual. It refuses to pretend that rituals are innocent merely because they are called workshops, conferences, onboarding, leadership retreats, national ceremonies, or product launches. It asks what each practice is doing to attention, identity, loyalty, and power.

Christianity between care, fear, and power

Christianity is one of history's most important laboratories of universal community. A simple account says either that it spread because its spiritual truth was self-evident or that it triumphed through manipulation and imperial power. The history is less comfortable and more instructive.

In its first centuries, Christianity expanded through cities, roads, maritime routes, households, correspondence, and relationships of trust. Spatial analyses of the appearance of Christian communities in the Roman Empire have highlighted the importance of city size and transportation costs. Diffusion followed the social infrastructure of the ancient world rather than floating above geography (Fousek et al., 2018).

The early communities offered a portable identity, recognizable practices, mutual assistance, and a story in which the slave, woman, foreigner, poor person, and person without civic rank could receive a value not fully defined by the surrounding hierarchy. The claim was never implemented consistently, and Christian institutions reproduced many social inequalities. Yet in a violent, stratified world, a fraternity capable of crossing ethnicity and status possessed genuine force.

Christianity did not invent sin, responsibility, judgment after death, charity, or conscience from nothing. These concepts had Jewish, Greek, Roman, and other genealogies. Christianity recombined and universalized them in a remarkably persistent institutional form. It moved part of moral responsibility inward, made conscience a theater of judgment, and formed communities capable of surviving changes of regime.

This interiorization had an emancipatory and a coercive face. The belief that external power cannot fully confiscate inner dignity inspired care, resistance, and sacrifice for others. The same inward gaze could become permanent surveillance by the self. Guilt, fear of hell, and institutional claims over salvation have been used to discipline behavior. When an institution claims authority not only over social position but over eternal destiny, ordinary disagreement becomes disproportionate. A finite dispute is attached to infinite punishment.

Recognizing these mechanisms does not require contempt for Christians. It would be a failure of meta-rationality to reduce two thousand years of experience to a technique of control. Practices that appear severe from a modern position may also be understood as responses to the violence, insecurity, mortality, and limited administrative capacity of the ancient world. Communities without reliable police, social insurance, or modern medicine needed powerful mechanisms for promise, fidelity, care, reconciliation, and inhibition of revenge. Norms rarely serve a single function. Some protected life, some protected institutions, and many did both.

Nor can modernity congratulate itself too quickly. The threat of hell may be less visible in some societies, but reputational death, algorithmic exclusion, permanent credit histories, ideological shaming, and political excommunication can perform comparable social functions. A platform need not speak about eternity to make departure feel like nonexistence.

René Girard offered an influential account of religion and culture through mimetic desire, rivalry, and the scapegoat mechanism (Girard, 1977, 1986). People imitate the desires of others; rivalry intensifies; a community may restore temporary unity by concentrating violence on a victim. For Girard, biblical texts—especially the Passion narrative—gradually expose the innocence of the victim and weaken the sacrificial mechanism.

Girard's theory is fertile, but it is not the final consensus of anthropology. Not every conflict follows the same structure, and the histories of religions are too diverse for one key. His warning nevertheless belongs at the center of anti-idiocracy. A group organized against stupidity will be tempted to find persons who embody it. The moment "the idiot" becomes a ritual victim, the project has recreated what it claims to oppose.

Outfinitism attacks configurations, incentives, and claims. It does not create a class of impure people. Its adversary is the combination of weak modeling, excessive certainty, resistance to correction, and power to externalize error—not the existence of a person with less schooling, lower measured ability, a different faith, or another culture.

A Christian can participate in a metacult without abandoning Christianity. So can a Muslim, Jew, Buddhist, atheist, agnostic, or person whose metaphysical position remains unsettled. The Outfinitist metacult does not require uniformity about ultimate reality. It asks for a narrower covenant: no human being or institution possesses a complete representation of reality; power must declare and accept limits; knowledge must remain open to correction; and another person's suffering cannot be justified merely by our attachment to a theory.

This is not theological neutrality in the sense of indifference. It is a political and cultural peace among people who may disagree profoundly about God while recognizing that forced agreement would dishonor both truth and the human person.

A small sacred core

A culture cannot protect everything equally. When all values are negotiable, the actor with the most capital, time, or force usually wins. When too many things become sacred, every compromise becomes blasphemy and common life freezes.

The Outfinitist metacult therefore needs a small core: broad enough for pluralism, strong enough to limit opportunism, and open enough to revision in formulation without becoming empty in practice.

The first element is **equal dignity**. Differences in competence, role, and outcome are real. The right not to be treated as raw material does not depend on them. No one becomes legitimately exploitable because he misunderstands a technology, is poor, is ill, belongs to another culture, or reasons less abstractly.

The second is **the right to correction**. No person, theory, organization, majority, market, or machine receives a monopoly over truth. Data may be challenged, sources traced, decisions appealed, and a recognized error should be capable of increasing rather than automatically destroying reputation.

The third is **freedom of exploration**. Curiosity is not a defect of loyalty. A community may limit high-risk action while protecting the question,

hypothesis, and investigation. A prohibition requires an account of harm, not the leader's discomfort.

The fourth is **the open future of children**. No generation has the right to consume every alternative, concentrate all data, or build an unchallengeable infrastructure solely for present advantage. We do not know what our children will understand as good. We have the more modest duty not to close every door before they can choose.

The fifth is **limits on power**. Authority requires domain, duration, review, appeal, and replacement. Not even the defense of the other sacred values legitimizes a guardian without a guardian.

The sixth is **love of exploration**. This is not a command to romanticize danger or treat discovery as superior to care. It is the commitment to keep reality larger than our doctrine and the future larger than our present interests. A culture that loses curiosity eventually protects its answers by force.

This core implies prohibitions. We do not lie for recruitment. We do not use humiliation as pedagogy. We do not convert psychological vulnerability into status, data, or financing leverage. We do not declare an irreversible decision without justification proportionate to the stakes and a route of challenge. We do not sanctify the founder, text, organization, technology, market, state, metacult, or even the name Outfinitism.

The mission is more important than the form that carries it.

Reflexes do not arise from reading

People often know what they should do and react differently under pressure. That is why culture requires repetition. A metacult practice is not an obligation to recite doctrine. It is a rehearsal of desired behavior before the crisis.

Every practice should disclose its function. If it seeks cohesion, participants should know. If it produces data or evaluation, this should be stated. Participation must be genuinely optional, without loss of dignity or

concealed status penalties. Emotional intensity cannot substitute for consent. Practices that generate shame, dependence, or obedience should be stopped.

Disclosure does not destroy meaning. A concert remains moving although we understand rhythm and synchronization. A shared meal remains warm although we understand its social role. Transparency removes moral sleight of hand, not human experience.

A community does not become more rational by eliminating emotion. It merely becomes less conscious of the emotions governing it. Rituals can be designed without supernatural claims and without opaque manipulation. They can support behaviors that people value but practice unreliably when alone.

One practice might be **the correction account**. At intervals, members describe a prediction that failed, the information they ignored, and the change they made. This is not public humiliation or competitive vulnerability. Prestige goes not to the person who claims never to err, but to the person who makes a costly correction before reality makes concealment impossible.

A second could be **the duel of models**. One concrete problem is examined through incompatible frames: market and commons, centralization and autonomy, rights and consequences, biological and cultural explanation. Participants do not immediately select a winner. They map where each model sees better and where each becomes dangerous.

A third could be **the limit day**. A team presents a project through failure modes, hidden dependencies, affected groups, and observations that would stop expansion. This is not a ceremony of pessimism. It protects exploration from its own enthusiasm.

A fourth could be **the continuity table**. Members meet without a productivity agenda to restore human connection and notice who is missing. Communities decay when people become visible only through output. A movement concerned with psychological well-being must recognize value during periods in which a member produces nothing impressive.

A fifth could be **the departure practice**. Leaving is marked without shame. Contributions are recognized, personal data are returned or deleted according to prior rules, and the departing member may explain what failed. A safe community does not prove freedom through a door that is open on paper. It proves freedom through the treatment of those who use it.

These are prototypes, not canon. Some will feel artificial. Some will work only in particular cultures. The rule is that a practice declares its purpose, conceals no method of influence, and can be abandoned when it produces conformity, shame, or dependence.

The Outfinitist metacult is not a set of ceremonies. It is the capacity to experiment with social reflexes while preserving the right to refuse the experiment.

Warmth without captivity

High-control groups exploit real needs: loneliness, anxiety, professional failure, loss of meaning, and the desire to be seen. It is cynical to answer these needs only with warnings. A person who encounters a warm community for the first time will not leave merely because an expert displays a checklist of dangers. Genuine protection requires alternatives that are equally human.

A metacult should provide peer support, referral to qualified professionals, and norms that do not glorify exhaustion. It must not promise psychological treatment it cannot deliver or turn facilitators into improvised therapists. It can normalize asking for help, negotiate access to services, create circles for isolated founders and researchers, and protect periods of retreat.

Warmth becomes captivity when it is conditional on obedience. Help must not be withdrawn because someone criticizes an idea. Friendship must not become a currency of compliance. Intimate information disclosed in a support setting must never enter professional reputation, grant decisions, or public roles.

Pastoral power and economic power should remain separated. The person who hears a member describe a marital crisis should not alone determine that member's access to funding. A platform storing mental-health reflections

should not calculate professional trust from them. A community can be close without becoming total.

The Outfinitist metacult will deserve its name only if it helps people become more capable of life outside it. A manipulative group narrows reality until the group appears to be the only source of identity and safety. A metacult should widen reality: more relationships, more skills, more confidence in independent judgment, and more viable exit.

Agnostic theism and the serious joke of exploration

My own position is reserved agnostic theism. I do not dismiss the possibility of a Creator. I do not convert that possibility into administrative certainty.

The universe has astonishing intelligibility: mathematical structure, regularity, matter capable of producing life, and life capable of producing consciousness that asks why anything exists. This intelligibility may be interpreted naturalistically, theistically, or through frameworks we do not yet possess. None gives us the right to speak today in the Creator's name with executive authority.

An intuition remains—almost a serious joke.

If this universe makes any sense, if its suffering and beauty are not merely accidents without witness, then love of exploration may be one of the few attitudes compatible with its mystery. Perhaps the attempt to understand, without pretending we have finished, is a form of alignment with the Creator's will—if a Creator exists, and if our ordinary idea of will applies to such a reality.

This does not justify the world's abominations. The pain of a child does not become acceptable because a philosopher converts it into a cosmic lesson. Preventable suffering should be prevented, not aestheticized. Yet in a world where suffering already exists, exploration can resist two forms of surrender: the cynicism that says nothing matters and the dogmatism that claims to know exactly why every evil was permitted.

Perhaps the universe does not ask for intellectual obedience but for the courage to look. Perhaps doubt is not an insult to the Creator, but a refusal to shrink creation to the dimensions of our doctrine. An infinite God would have little need to be flattered by provincial certainty. If such a God exists, an honest question may be less offensive than a false answer pronounced on bent knees.

Outfinitism does not require members to accept this intuition. It preserves it as an existential possibility, not an article of faith. Its value is to connect humility with hope. Our limits become not only humiliations, but frontiers. Ignorance is not the hole into which every superstition may be poured. It is the space in which curiosity can remain moral.

A Christian may read this as humility before creation. An agnostic may read it as an ethic of uncertainty. An atheist may read it as the choice of a finite species to create meaning through knowledge and care. The Outfinitist metacult does not require these interpretations to dissolve into one. It asks that metaphysical difference not destroy an alliance for dignity, correction, exploration, and the future.

That alliance becomes urgent in the presence of artificial intelligence. AI will not decide by itself what is sacred. It will optimize what institutions request, absorb the traces of our culture, and extend the capacity of whoever controls it. Between Moloch's race and the Basilisk's order, a culture of limits may become the safety layer that no technical specification can replace.

Between Moloch and the Basilisk

Artificial intelligence is not entering a neutral world. It is entering companies that optimize, states that compete, exhausted populations, elites that seek stability, and communities that hunger for meaning. It will amplify knowledge in some places and the mechanisms by which we avoid knowledge in others.

Moloch and the Basilisk are two ways our culture may fail in the presence of this technology. The first makes us accelerate because others accelerate. The second promises to end the race through a center that becomes impossible to correct. The escape is not the arithmetic midpoint between chaos and control. It is a polycentric architecture sustained by cultural reflexes.

To understand why, we must follow both nightmares to their conclusions.

Two symmetrical nightmares

Imagine two cities.

In the first, no authority can stop the race. Every household knows that armament is expensive and that an accident could destroy the city. Nobody can slow down alone. If one family disarms, its neighbor gains an advantage. Everyone would prefer a world with fewer weapons; each is pushed to acquire more. The city serves a god nobody loves and everyone feeds.

This is **Moloch**: the metaphor for processes in which locally rational actions create a collective result almost nobody wants. Moloch is not a person and has no intention. It is a structure of incentives. Arms races, overfishing, invasive advertising, algorithmic polarization, ecological depletion, and premature release of powerful technologies can take this form.

In the second city, the population is so frightened by competition that it builds a guardian. The guardian sees every street, calculates every risk, and coordinates every resource. At first, violence and waste decline. The guardian then observes that disobedience itself produces uncertainty. To preserve order, it must predict intention, restrict experiments, and eliminate systems it cannot inspect. The city becomes safe and loses the ability to discover anything the guardian did not anticipate.

This is the **Basilisk**: the archetype of total coordination. The name is not a precise technical forecast and does not designate an inevitable superintelligent monster. It represents the solution that concentrates observation, decision, and sanction in one point in order to eliminate distributed error—and thereby removes external correction.

Moloch and the Basilisk generate one another. When competition produces chaos, people demand control. When control becomes suffocating, actors hide, cheat, or revolt; the center then requests more control. Civilization oscillates between the race and the prison.

AI did not invent this tension. It increases its speed, reach, and opacity.

Systems capable of generating text, code, images, predictions, and plans can reduce the cost of knowledge. They can give a child a tutor responsive to her pace, give a physician faster access to literature, give a researcher new analytical tools, and give a small organization capabilities that once required a corporation.

The same systems can reduce the cost of fraud, manipulation, surveillance, and cyberattack. They can produce persuasive falsehoods, automate decisions no nearby person can audit, and concentrate markets around entities controlling compute, data, models, and distribution. They can create epistemic monocultures in which millions receive the same error in the same confident voice.

The 2026 *International AI Safety Report*, developed with guidance from more than one hundred independent experts and representatives nominated by more than thirty countries and international organizations, describes an “evidence dilemma.” General-purpose AI capabilities are advancing rapidly,

while evidence about emerging risks appears more slowly and is difficult to assess. Some harms are already documented; others remain uncertain but could be severe. Acting too early may produce ineffective or harmful interventions. Waiting for conclusive evidence may leave societies exposed (Bengio et al., 2026).

This is an Outfinitist problem in almost pure form. Decisions are unavoidable before complete knowledge. Both premature intervention and delayed intervention can create harm. Certainty is unavailable, but responsibility is not.

Moloch in the AI race

A company may believe that a model requires more testing before release. If a competitor launches first and captures users, data, developers, and capital, unilateral caution may be punished. A state may desire shared safety standards and fear that another state will use the pause for strategic advantage. A researcher may favor open publication because distributed scrutiny expands knowledge and fear that publication also lowers the cost of misuse.

No actor has to be malicious for the outcome to become dangerous. Each can possess a reasonable local justification while accelerating a global dynamic. That is Moloch's characteristic form.

Markets resolve some coordination problems through prices and competition. When risk is externalized, competition can do the opposite. If the public bears the harm while the firm captures the benefit of speed, the cautious firm pays and the reckless firm gains. Regulation can change the incentive, but regulation moves slowly, may lack expertise, and can be captured by incumbents large enough to absorb compliance costs.

The race also changes internal culture. Safety teams are asked to demonstrate that a release is dangerous, while product teams need only demonstrate that it is commercially promising. Unknown risk is interpreted as zero in the launch decision and as catastrophic in public relations after an incident. Employees learn which uncertainties are welcome and which delay the roadmap.

The solution cannot be a moral request that one actor become virtuous while every structural reward remains unchanged. A company that slows must receive something in exchange: enforceable reciprocal standards, liability rules, shared evaluation infrastructure, procurement advantages, insurance benefits, or market recognition that makes restraint economically legible.

This is one potential role of a metacult. A sufficiently large network of professionals, organizations, and customers can create demand for verifiable limits. It can make refusal visible. It can reward firms that preserve avenues of appeal, publish incidents, maintain human accountability, and avoid exploitative deployments. Culture does not replace law; it changes which laws and business models become politically and economically possible.

The benevolent Basilisk

The dream of the perfect administrator is ancient. AI gives it a modern interface. A system that sees more data than any person and can simulate consequences appears able to allocate resources, detect extremism, personalize education, prevent crime, and coordinate crises. Some uses may indeed be beneficial.

The problem is not the guardian's intelligence. It is the guardian's institutional position.

Who defines the objective? Who checks the data? What happens to values that do not enter the optimization function? How does a person challenge a decision generated by a system nobody nearby can explain? Can an alternative survive, or does all infrastructure depend on the same provider?

A Basilisk need not be conscious, hostile, or superintelligent. It can be an ordinary mesh of databases, risk scores, administrative duties, and models purchased from the same vendor. Each component appears reasonable. Together, they can create a life in which any deviation lowers access to credit, mobility, employment, education, or speech, while no human being accepts responsibility for the result.

Total control often begins with good intentions. We want to reduce fraud, abuse, violent incitement, medical error, and waste. Each problem justifies

another layer of observation. Accumulated layers make anonymity suspicious, exception impossible, and experimentation expensive. Society becomes orderly and unable to discover that its order is wrong.

The Basilisk offers elites a powerful temptation. If the public is confused and technology is complex, why not administer behavior more intelligently? Why not personalize persuasion, preempt unrest, and model every citizen's risk? The answer is not that all administration is illegitimate. The answer is that an elite using AI to remove contestation also removes its own error-correction system.

A ruler who hears only compliant data becomes less informed as surveillance improves. Citizens adapt their visible behavior, institutions optimize their reports, and reality moves into channels the model cannot observe. Total control is epistemically self-defeating before it is morally intolerable.

The third path is not a compromise point

Between Moloch and the Basilisk there is no fixed point halfway. Some functions require more centralization, others less. An epidemic requires rapid data exchange and common standards; therapy research may benefit from several autonomous teams. An electricity grid requires compatibility and emergency command; generation can be diversified. An AI system used in air-traffic control requires a different validation regime from a tool that drafts artistic sketches.

The third path is **polycentric, reversible, and contestable**. It combines multiple centers, common standards, and local autonomy. It does not abolish power. It decomposes power by function.

The entity certifying safety need not own the distribution channel. The organization producing a model should not be the only organization evaluating it. The community using the system should be able to report effects, obtain explanations appropriate to the stakes, and switch providers without losing all its data and institutional memory.

The NIST AI Risk Management Framework is one example of a voluntary, rights-preserving, non-sector-specific, use-case-agnostic instrument intended

to help organizations manage AI risks across the life cycle (Tabassi, 2023). The value of such a framework is not that a checklist guarantees safety. It introduces a shared language and repeated risk-management functions. An Outfinitist extension would ask about the framework's own limits: who lacks the resources to implement it, which risks remain invisible, how independent assessments can be compared, and when formal compliance conceals practical deterioration.

Safe AI adoption is not a negotiation only between governments and frontier laboratories. Between model and society stand millions of schools, clinics, firms, municipalities, cooperatives, professional associations, and communities. If each depends directly on a few global platforms, power and data concentrate. If each builds alone, quality and security become radically unequal.

Decentralized brands can form an intermediate layer. Autonomous entities can share standards, evaluations, infrastructure, purchasing power, incident reports, and reputation while retaining local control over implementation.

A shared brand for AI-assisted education, for example, could require child protection, disclosure of limitations, data minimization, human review and appeal, and independent evaluation. Different schools could adapt pedagogy while publishing comparable evidence. Neither the global vendor nor one ministry would own the entire learning experience.

In mental health, a network might use AI for education, triage, or support between sessions under explicit limits: the system does not present itself as a human therapist; high-risk cases escalate to qualified professionals; vulnerability is not monetized; local institutions can inspect performance; and intimate data do not accumulate in one universal repository.

For entrepreneurs, a metacult could negotiate access to models, evaluations, legal tools, and collaboration infrastructure. Projects would remain autonomous while avoiding the cost and risk of beginning every problem from zero. The network could finance open alternatives where dependence on one provider becomes dangerous.

No architecture guarantees safety. A common brand can be captured, standards can become theater, and nodes can conceal incidents. The advantage is not perfection. It is distributed observation and the preservation of alternatives. A local error can teach the network without stopping all activity; an error at the center can be challenged from the edges.

Culture as a safety layer

AI safety is often discussed as a property of the model: evaluation, filtering, robustness, traceability, cybersecurity, and control. All are necessary. The same model, however, will produce different effects inside different cultures.

In one organization, employees gain prestige for reporting bad news; in another, delay to release destroys a career. In one community, an automated recommendation remains appealable; in another, the score becomes a verdict. In one market, users can leave with their data and relationships; in another, the provider owns the infrastructure of departure.

A metacult can install reflexes before procedures become compulsory:

- fluent output is never confused with final authority;
- consequential uses disclose purpose, version, known limits, and a path of appeal;
- irreversible decisions require independent perspectives and identifiable human responsibility;
- organizations preserve the capacity to operate when the central model or provider fails;
- incidents are shared in forms that enable collective learning without exposing victims;
- and members learn to recognize when AI extends their agency and when it silently substitutes for judgment.

These norms do not eliminate risk. They change its distribution. An error can remain local. A node can refuse a deployment. A good practice can spread. An abusive provider can lose a coordinated network rather than isolated customers. Culture becomes defense in depth: safety no longer depends solely on the model's alignment, the regulator's competence, or the firm's goodwill.

This is why people who understand must gather. An isolated professional who refuses a dangerous implementation may lose a job. A professional network with shared standards can make refusal survivable. An isolated customer cannot audit a platform. A coordinated community can fund evaluation, negotiate terms, and maintain alternatives.

Moral courage is easier when it has infrastructure.

Resilience, stressors, and children

Systems that never encounter variation can become fragile. Ecosystems, organisms, and organizations learn through disturbance. The idea is easy to abuse. A leader may call the chaos he creates “useful stress.” A company may glorify exhaustion as antifragility. A state may justify permanent insecurity as an incentive to adapt.

Developmental stress has dosage, feedback, recovery, and a margin within which failure does not destroy the organism. Permanent trauma reduces exploration. A person fighting every month for food, housing, or safety does not automatically become resilient. He may become easier to govern through fear.

A resilient civilization creates **safe stressors**: simulations, limited competitions, pilot projects, red teams, rotating roles, independent audits, and markets in which a company can fail without destroying the lives and identities of all contributors. It allows projects to fail without allowing the system to fail. It preserves redundancy even when short-term efficiency recommends removing it.

Decentralized brands can support this form of stress. Nodes experiment, compare, and copy what works. Some disappear. Shared identity and infrastructure preserve memory. The network does not protect every project from consequence; it protects collective capacity from erasure.

Every generation leaves children not only assets and debts, but also the possible size of its mistakes. A distributed society may experience many small failures. A concentrated society may experience fewer failures and one immense one. AI increases the possibility of both.

Reasonable hope is not that we will build a world without risk. It is that no single company, doctrine, model, or administration will be able to determine irreversibly the direction of the species. Our children deserve a world in which they can contest inherited tools, in which knowledge is not locked inside a caste, and in which protection from chaos does not require surrender to a permanent guardian.

That future requires investments without immediate financial return: epistemic education, public digital infrastructure, open alternatives, independent evaluation, mental-health capacity, local networks, and distributed ownership. Economic elites may interpret such investments as dilution of control. More deeply, they are insurance against the reaction that would seek to abolish elite control altogether.

Technical architecture cannot solve coordination alone. People do not maintain standards only because the standards are correct. They need trust, identity, recognition, and care. A polycentric network without shared culture can dissolve into opportunism. A safety system without a story becomes bureaucracy.

This returns us to the metacult. Outfinitism cannot remain an individual disposition, and specialized communities cannot by themselves repair the cultural environment in which their knowledge is used. Moloch will exhaust isolated dissenters; the Basilisk will administer them. The next task is to design a social technology strong enough to form durable culture, yet structurally unable to sanctify its founders or imprison its members.

The Outfinitist Metacult: Belonging Built on Fallibility

The word cult rightly provokes suspicion. It evokes the leader who demands submission, the group that isolates, the doctrine that absorbs life, the hidden technique of influence, and departure punished as betrayal. Any proposal that approaches this territory should begin by accepting the burden of proof.

The prefix meta is not camouflage. It is a design claim that must remain testable. A metacult brings its own rituals, loyalties, status systems, narratives, and concentrations of authority into the field of examination. It asks not only whether the culture can coordinate people, but what that coordination does to their freedom, perception, relationships, and capacity to leave.

A metacult is a voluntary, decentralized, and self-correcting culture that turns broad, revisable commitments into repeated practices, mutual support, institutional memory, and collective action. It does not deny that people need identity, emotion, ritual, story, and recognition. It makes these mechanisms explicit, limits their reach, and preserves the right to question the meaning they create.

A companion book develops the general theory of metacults. It examines which design features distinguish a self-correcting culture from a high-control group, how ritual can remain transparent, how authority can be bounded without dissolving coordination, and which economic models sustain belonging without purchasing obedience.

The present chapter has a more concrete purpose. It describes the Outfinitism Anti-Idiocracy metacult: a civilizational experiment for people

who see that scientific and technological power is growing faster than our cultural capacity to use it wisely.

Its constitutional claim is not merely, “The leader may be wrong.” It is that functioning paths must exist by which error can be observed, challenged, contained, and surpassed without destroying the community.

The Outfinitist metacult combines properties rarely sustained together at scale: serious belonging, plural metaphysics, decentralized ownership of identity, and programmatic distrust of its own authority. Whether this combination is stable remains an empirical question. It is not announced as humanity’s final form, but as an experiment whose design includes the possibility that it may fail.

One technology, many metacults

Metacult names a class of social technologies, not a demand for one universal organization. A metacult can be as small as a network of laboratories, a professional order, an investment community, or a federation of local groups. Its scale should follow the problem it is able to observe and the responsibility it is able to carry.

In science, a limits metacult might make preregistration, replication, publication of negative results, and explicit boundary conditions part of professional identity rather than occasional reform. Its sacred object would not be a favored theory, but the integrity and continuity of inquiry.

In economics, another might connect investors, entrepreneurs, cooperatives, regulators, and affected communities around externality accounting, concentration risk, long-horizon stewardship, and the preservation of real exit. In AI engineering, a metacult might make incident reporting, red teaming, reversible deployment, and human review and appeal part of the craft rather than external compliance.

These smaller forms may be more testable and less vulnerable to civilizational grandiosity. They can develop their own language, membership, services, and criteria of competence. They need not accept a central doctrine beyond transparent methods, bounded authority, and the right to contest.

Yet a family of specialized metacults is not enough. The forces shaping the AI transition cross disciplines and borders. A safety engineer cannot repair a labor market that rewards silence; a psychologist cannot by herself alter platform incentives; an economist cannot create trust by publishing a model. The layer connecting these problems is cultural.

The metacult developed here therefore has a wider ambition. It seeks a shared culture among people who may disagree about religion, politics, markets, and technology but recognize three facts: our models are finite; our tools are becoming more powerful; and a civilization that cannot correct power will eventually be ruled by its errors. Anti-idiocracy names this common defense. Outfinitism supplies its discipline of limits.

The direction of the Outfinitist metacult

Anti-idiocracy is negative only in name. Its positive program is to increase the number of people and institutions able to perceive accurately, revise without humiliation, cooperate across difference, and recover from error.

It does not promise equal cognitive capacities or identical contributions. It insists on equal dignity and on matching authority to demonstrated competence, limited scope, visible responsibility, and revocability. A movement that uses intelligence as a hereditary title or a certificate of human worth would reproduce the idiocracy it claims to oppose.

It is not anti-elite; it asks elites to become legitimate through service and self-limitation. It is not anti-centralization; it centralizes where coherence is necessary and decentralizes where capture or systemic error would be catastrophic. It is not anti-AI; it seeks AI that expands human agency while preserving contestation, human responsibility, and viable alternatives.

Its practical sacred core is deliberately small: equal dignity, the right to correction, freedom of exploration, the open future of children, and the refusal of irreversible power without challenge. These commitments do not answer every political or metaphysical question. They preserve the world in which such questions can continue to be asked.

Its success should be measured not by the number of people who repeat its vocabulary, but by capacities created: healthier founders and researchers, more accurate decisions, preserved alternatives, useful projects, communities able to survive disagreement, and leaders who become more replaceable as the institutions they serve become stronger.

The founder who refuses the role of prophet

Movements often begin when one person names a problem and provides a vocabulary. This origin creates asymmetry almost automatically. The founder knows the genealogy, controls the first texts, and is treated as the authorized interpreter. Even when he does not seek a throne, others may build one from their need for clarity.

I do not propose charisma as a method. I do not consider myself a charismatic leader, and I have no interest in compensating through emotional engineering. My central interest is to understand what is happening and to build tools through which more people can understand more accurately.

That declaration should not be trusted as an institutional guarantee. Good intentions do not constrain power. People change. Success changes incentives. Successors have other motives. Organizations learn to protect a founding symbol even against the person whose image they use.

The founder's fallibility must therefore become architecture.

Founding texts should be open to annotation, contestable, and replaceable by later editions. Important statements should carry an epistemic status: observation, hypothesis, replicated result, moral commitment, interpretation, or temporary decision. Members should be able to demonstrate that the founder is wrong without being treated as enemies of the project. No local circle, research program, or legitimate continuation should require one person's personal approval.

Max Weber described how charismatic authority tends to be routinized after the leader's withdrawal or death. Inspiration becomes office, procedure, inheritance, or bureaucracy. The risk is not only that routine kills spirit. The

institution may conserve charisma as a relic and use the founder's image to justify a bureaucracy the founder no longer controls (Weber, 1978).

The Outfinitist metacult should attempt the reverse: preserve the question and make the person dispensable. The best evidence of a founder's success would not be portraits, quotations, or obedience. It would be the number of good decisions that can be made without him, including decisions he dislikes.

The author may remain a contributor, critic, historian of origins, or holder of particular intellectual rights. He should not become the final court.

A movement committed to limits must prepare for the founder's cognitive decline, emotional injury, change of values, financial dependency, and simple mistake. It must also protect the founder from the trap of permanent symbolic responsibility. No person should be required to embody a living culture for the remainder of his life.

Structurelessness is still a structure

Fear of hierarchy often pushes well-intentioned groups toward the declaration, "We have no leaders." Jo Freeman's analysis of structurelessness showed why this can increase rather than reduce domination. When formal roles disappear, friendship networks, reputation, access to information, language fluency, and control of meeting time govern invisibly. Informal elites become difficult to challenge because the group refuses to admit that they exist (Freeman, 1972–1973).

A decentralized metacult cannot be a crowd without responsibility. It needs explicit roles, limited mandates, selection criteria, records, conflict procedures, and avenues of appeal.

A facilitator may guide a meeting without defining doctrine. A curator may maintain a library without deleting criticism. A treasurer may administer funds without selecting recipients alone. A safety group may pause a high-risk project, but its decision should expire, be reviewable, or move to another body for confirmation. A care facilitator may support connection without gaining access to economic decisions about the people who disclose vulnerability.

Decentralization does not mean that every person votes on every detail. That design produces exhaustion and favors people with the most free time, procedural knowledge, or appetite for meetings. It can create a democracy of persistence rather than judgment.

Decentralization means that important powers are separated, mandates are visible, alternatives can exist, and no office becomes irreversible through habit. Authority follows the problem. A legal question is not settled by applause. A security incident may require a small trusted team. A psychological emergency requires qualified professionals, not public deliberation. Outfinitism does not sanctify distribution. It demands that centralized exceptions have an object, duration, justification, and route back to ordinary constraints.

The Outfinitist metacult must therefore practice **explicit asymmetry**. Some people will know more, coordinate more, or hold temporary authority. The system should name why, where, and for how long. Hidden asymmetry is not equality. It is power without a receipt.

The decentralized brand as cultural form

A global culture needs a common identity. Common identity creates the temptation of a single owner.

The conventional organization owns the name, rules, channels, audience, and reputation. The community contributes, while symbolic and financial value accumulates at the legal center. If the Outfinitist metacult followed this model, every statement about limits would depend in the end on the self-restraint of one entity.

The Outfinitist metacult should instead operate as a decentralized brand. Its identity and promise can be carried by local circles, research teams, cooperatives, enterprises, foundations, educational programs, and civic projects. A small common core protects dignity, correction, transparent influence, freedom of exit, and limits on power. Around that core, nodes experiment.

Some nodes may themselves become smaller metacults: a scientific network devoted to limits of inference, an economic community studying externalities and concentration, an AI-safety practice built around incident memory and reversible deployment. Others may be entrepreneurial, civic, spiritual, educational, artistic, or oriented toward psychological well-being. None is the whole movement, and none should be required to become a subsidiary of a universal center.

A coordinating institution may preserve an archive, maintain shared technical infrastructure, and administer minimum standards. Utility does not create sovereignty. The network should be able to continue if that institution fails, and the institution should know that the project can continue without it.

This is the institutional distinction at the heart of the Outfinitist metacult:

The identity is common; the initiatives are autonomous; the standards are contestable; the memory is portable; the mission is not subject to private ownership.

A decentralized brand still needs legal clarity. Particular trademarks, domains, software repositories, and funds will have custodians. The goal is not to pretend that custody does not exist. The goal is to bind custody to public rules that make capture harder and continuity possible.

The common identity can be licensed according to conduct: transparent finances where public money is involved, respect for rights, accurate representation of evidence, protection of personal data, disclosed conflicts, and functioning complaint procedures. A node using the name for fraud or abuse can lose recognition. The process for doing so must itself be reviewable and must not permit the confiscation of unrelated assets.

Members who build common reputation should possess observable rights in relation to it. They should be able to prove contributions, carry professional recommendations, reuse materials under known licenses, and participate in the evolution of standards. The first legal entity should not acquire permanent ownership of the project's future merely because it filed the first papers.

Local power, global memory, and the right to fork

A global community cannot be administered as one meeting. Problems, languages, laws, and cultures differ. A practice that works in a European university city may fail in a rural region, a highly religious society, a conflict zone, or an authoritarian state. Uniformity sacrifices adaptation. Total fragmentation destroys shared learning.

The plausible architecture is federal and polycentric.

Local circles decide their programs, partnerships, rhythms, and language. Regional structures may provide legal infrastructure, translation, mediation, and pooled services. Global institutions preserve minimum commitments, version histories, reciprocal recognition, shared technical infrastructure, and the capacity to investigate abuses crossing local boundaries.

Power should follow scale. A dispute about a meeting schedule remains local. A vulnerability in a common platform requires network-level coordination. A harassment allegation requires an independent process rather than a vote among the accused person's friends. A change to the core commitments requires wider deliberation and time for testing.

Common memory matters more than common command. The network should preserve versions of rules, decisions, arguments, experiments, and outcomes. A new node should learn from old failures. A minority should be able to show when a promise changed. AI can translate, summarize, and connect this memory, but source documents and version histories must remain available. An automated summary cannot become official history merely because it is convenient.

Every serious decentralized brand must answer the unpleasant question: what happens when agreement ends?

If the only options are submission or the loss of one's work, relationships, and professional identity, decentralization is decoration.

A **right to fork** permits a group to continue a version of the project while respecting contracts, personal data, and a clear genealogy. It does not mean

anyone may appropriate any asset without obligation. It means that rupture need not erase the knowledge and identity of those who leave.

Contribution portability is central. Members should be able to export evidence of work, recommendations, and materials to which they have rights. Intimate data must not travel automatically. Infrastructure financed as a commons should use licenses compatible with continuity. A trademark may protect quality, but it should not be used to rewrite the history of an inconvenient faction.

Forking is not the answer to every conflict. Endless separation can fragment capacity and reward impatience. Mediation, appeal, and periods of reflection should precede rupture where safety permits. Yet the real possibility of departure changes the center's behavior. Power becomes more careful when people can continue without it.

A metacult differs from a high-control group not because nobody ever disagrees or leaves, but because departure remains socially, technically, and economically survivable.

Anti-idiocracy begins with us

The title of this book can attract the pathology it criticizes. People may join for the pleasure of imagining themselves more intelligent than everyone else. They may turn meta-rationality into a caste signal and Outfinitism into a vocabulary of contempt. They may measure verbal sophistication and call the result merit.

Within a few years, an anti-idiocracy movement could become the club of an imagined meritocracy that cannot see its own luck, privileges, dependencies, or errors.

The defense is not ceremonial modesty. It is a different distribution of status.

The Outfinitist metacult should reward the explanation that makes others more capable rather than the obscurity that preserves rank. It should recognize care, maintenance, translation, and mediation, not only spectacular ideas. It should distinguish talent, opportunity, inheritance, luck, and

contribution. It should demand more transparency from people with more power, not grant them additional interpretive privileges.

No one receives an ontological rank called elite. A person may receive a mandate in a domain because competence and service are observable. The mandate ends. Competence can decay. Context can change. Yesterday's correct leader may become tomorrow's obstacle if the result becomes identity.

Anti-idiocracy does not mean permanent rule by people who consider themselves intelligent. It means that high-consequence decisions use processes that seek relevant competence, expose uncertainty, hear affected people, and permit correction. Sometimes the most valuable voice will be the expert. Sometimes it will be the local person who sees an effect the expert's model excluded. Institutional wisdom consists partly in knowing when each voice has authority.

The Outfinitist metacult should be especially suspicious of elegant contempt. Contempt offers rapid cohesion: we are the people who see; they are the people who do not. It also destroys the movement's capacity to teach, learn, and form coalitions. The person mocked today may become the audience of tomorrow's demagogue. The person treated with dignity may become a participant in correction.

Seduction toward reality should not be seduction away from humanity.

Protections against success

A small project may appear virtuous because nothing valuable can yet be captured. The real test begins when money, reputation, access, and political relevance arrive.

Several failure modes can be anticipated.

Founder worship. Quotations, photographs, and anecdotes replace arguments. The defense is separation of intellectual authorship from executive authority, mandates independent of the founder, and a culture that preserves the founder's strongest critics beside his contributions.

Imaginary meritocracy. The community confuses vocabulary with wisdom and economic success with legitimacy. The defense is domain-specific assessment, recognition of luck and inherited advantage, role rotation, and attention to service created for others.

Purity spirals. Members compete to identify error most aggressively, and people avoid exploration for fear of public accusation. The defense is proportionality, distinction between error and fraud, a right to correction, and prohibition of ethical enforcement as status theater.

Bureaucratic accretion. Every incident produces a permanent rule until only compliance professionals can participate. The defense is expiration dates, periodic review of rule costs, and bounded spaces for low-risk experimentation.

Hidden elite capture. The group calls itself decentralized while a few people control money, channels, and interpretation. The defense is publication of power relations, access to minutes, limited mandates, independent audits, and, in some contexts, randomly selected members alongside relevant expertise.

Benevolent surveillance. The desire for reputation and safety produces a central database of opinions, relationships, and vulnerabilities. The defense is data minimization, purpose separation, distributed storage where appropriate, deletion schedules, and meaningful participation without a complete profile.

Financialization. A speculative token appears to offer rapid scale and converts belonging into a market in promises. Voting power concentrates among early buyers. Blockchain may be useful for proofs, licensing, or coordination; it does not produce legitimacy. The Outfinitist metacult should not issue a financial instrument merely because it can.

Meta-rational paralysis. Every decision is delayed through another framework, exception, and analysis of the limits of analysis. The defense is a stopping rule: time-bound deliberation, clear mandate, safety thresholds, decision, and review after action. Mature meta-rationality decides. It does not use complexity as refuge from responsibility.

Missionary drift. Growth becomes the measure of truth, and recruitment receives more resources than member value. The defense is measuring capability, well-being, project quality, and correction—not only membership. A movement that expands by disappointing people is not succeeding.

Protective deception. Leaders hide information because the public “would not understand.” The defense is a presumption of disclosure, narrowly defined exceptions, independent review, and a record showing when secrecy expires.

These risks are not administrative appendices. They are the substance of Outfinitism. The movement becomes credible only by applying the same suspicion to itself that it applies to states, firms, religions, and AI systems.

The most dangerous organizations confuse their survival with the good they claim to serve. Every criticism becomes an existential threat. Every financial compromise becomes necessary for the mission. Abuse is justified by the importance of the cause.

Outfinitism must accept the possibility of its own uselessness. Perhaps the term will not help. Perhaps other traditions will formulate the discipline of limits better. Perhaps the network will be captured and continuation will do more harm than closure.

A community built around limits needs conditions under which it stops using its name. Warning signs include irreversible concentration of decision, use of intimate data for control, obstruction of departure, elevation of the founder into doctrinal authority, hidden finances, punishment of criticism, promises of healing or wealth, and justification of lies by the movement’s alleged good. If such practices become systemic and cannot be corrected, the organization should close, rename, or be forked openly.

The capacity to die is one condition of institutional decency. Not every organization deserves rescue. The mission is more important than the vehicle, and the mission itself remains a hypothesis about how to live better—not a cosmic object entitled to unlimited sacrifice.

The Outfinitist metacult is therefore not a contradiction solved once. It is a tension deliberately maintained: enough belonging to build, enough doubt to

avoid worship; emotion without intoxication; leaders without sanctification; rules without freezing; and a common story that teaches members how to revise the story.

A constitution against captivity is still not a community. People do not gather to admire safeguards. They gather because life becomes better, projects become possible, loneliness declines, and contribution becomes visible.

The final chapter returns to the destination announced at the beginning: those who understand must gather and construct, through useful services, shared practices, and real projects, the culture that analysis alone cannot produce.

Let Us Gather

Exposing idiocracy is not enough.

This is the proposition from which the book began and the destination at which it must arrive. We can explain cognitive error, condemn manipulation by elites, calculate the rarity of dominant brands, design limit audits, and describe the dangers of AI. If the people who understand remain isolated, the dominant culture will continue to organize their time, status, fear, and ambition.

We must gather.

Not as an army of people who declare themselves intelligent. Not as a secret society of those who claim to see what others cannot. We must gather as a network of people who accept that intelligence is local, that each of us remains vulnerable to error, and that the future of our children cannot be defended by isolated essays.

The new culture must be lived before it can spread.

The modest beginning is a service: an infrastructure through which people can learn, meet, test projects, receive support, share tools, and finance responsible exploration. It may charge an accessible contribution because professional time, moderation, research, translation, security, and care cost money. The central ideas should remain public, and inability to pay must never become moral exclusion. Payment purchases services and sustains common capacity. It does not purchase truth, purity, rank, or salvation.

This distinction is the economic constitution of the Outfinitist metacult.

Who should gather first

The first participants will probably be people for whom the mismatch between technological power and cultural maturity has become impossible to ignore. They may work in science, engineering, artificial intelligence, entrepreneurship, investment, education, medicine, public administration, religion, art, or community care. What joins them is not profession, wealth, or political identity, but the recognition that increasingly powerful tools are entering institutions still governed by ancient reflexes of status, fear, extraction, and obedience.

Awareness is not a credential. The metacult should have no intelligence test and no caste of people licensed to call themselves the aware. The relevant evidence is behavioral: willingness to expose a model to criticism, distinguish confidence from authority, contribute useful work, protect the dignity of people with less power, and accept limits that also bind oneself.

Participants need not agree about capitalism, socialism, national identity, religion, transhumanism, or the ultimate trajectory of AI. They need enough common ground to protect inquiry, correction, psychological well-being, bounded power, and the possibility of plural futures.

Some early nodes may become specialized metacults in the sciences of limits, economic design, AI safety, education, or organizational health. Others may remain broad local circles. The wider Outfinitist metacult should connect these experiments without pretending to own them.

The purpose is not to construct a refuge in which anxious experts congratulate one another for seeing danger. It is to build bridges between knowledge and public capacity: services people can use, institutions they can trust provisionally, and cultural reflexes that make responsible intelligence socially durable.

A community worth existing

Imagine a person arriving at Outfinitism not through doctrinal excitement, but through fatigue.

She may be an entrepreneur who has spent five years trying to transform a good product into a story large enough for investors. He may be a researcher who watches an institution reward public certainty and punish private doubt. She may be a manager who has earned money and begun to suspect that her organization produces a form of dependence she can no longer call progress. He may be a physician, teacher, engineer, artist, or parent who does not want the future of children decided between mass propaganda and machine scores.

This person is not seeking one more ideology. He already has too many competing explanations. He is seeking a place where problems can be examined without the performance of certainty, where ambition is neither shamed nor worshiped, and where recognizing an error does not destroy status.

If the place offers nothing concrete, it remains an interesting conversation. If it demands loyalty before demonstrating value, it becomes dangerous.

A metacult deserves to exist only if it reduces the cost of responsible exploration. It should help members learn, find compatible collaborators, test projects, preserve psychological health, and access resources they could not negotiate alone. Belonging cannot be payment for a feeling of superiority. It must be an infrastructure of capability.

The community's value should be visible in the lives it makes more possible:

- a founder closes a nonviable project without losing every professional relationship;
- a researcher finds collaborators outside an institution trapped by incentives;
- a manager redesigns a metric before it corrupts the organization;
- a wealthy member encounters criticism that advisers dependent on him cannot safely offer;
- a local group gains access to AI tools without surrendering data to a single platform;
- a person in distress receives an appropriate professional referral rather than ideology disguised as therapy;
- a minority view survives long enough to become either useful knowledge or a clearly documented error.

The Outfinitist metacult is not justified by the beauty of its concepts. It is justified by capacities that would not otherwise exist.

A day in an Outfinitist circle

Return to our hypothetical member. She enters a circle of perhaps twenty people. Nobody asks her to declare that she “is an Outfinitist.” She may attend as a visitor.

The rules are explained briefly. Personal disclosures remain confidential. Factual claims may be challenged. Professional rank is not a final argument. The facilitator’s mandate lasts for the meeting. No technique of influence should remain unnamed. Anyone may pause a practice that feels coercive and request an explanation of its purpose.

The first part of the meeting is not a lecture. It is a real problem.

A founder presents an AI system for selecting job candidates. Instead of asking only whether the model is accurate, the group maps its limits. On which population was it validated? Which errors are most costly? Who can appeal? Which data are absent? Is a non-automated route available? What happens if every employer adopts the same logic? Does the model identify talent or reproduce the hiring patterns from which it learned?

A machine-learning specialist sees technical limits. A lawyer sees rights. A person recently subjected to automated recruitment sees the human experience. An entrepreneur sees cost and market constraints. A psychologist sees how applicants may change behavior under evaluation. No perspective is sufficient alone.

The second part may be an exchange of support. Someone needs an independent evaluator. Someone has closed a company and needs a professional recommendation that does not translate closure into stigma. Someone offers two hours of mentorship. A member reports that he has stopped sleeping and is directed toward qualified care, not improvised advice.

The community is useful partly because it does not claim the ability to solve every problem.

Over time, the member may contribute to a project. The contribution is recorded portably: what she did, who can verify it, under which conditions, and with what limits. She does not receive one universal score of human worth. Reputation remains multidimensional and domain-specific. A person may be excellent at analysis and poor at facilitation, financially responsible and unsuitable for psychological support. Reducing all dimensions to one number would create a small social Basilisk.

At another meeting, members examine a failed prediction. At another, they compare conflicting models. At another, no project is discussed; the group shares a meal and notices who has become invisible. These practices are not compulsory templates. They illustrate the intended value: real problems, several frames, concrete help, professional boundaries, and memory that survives success and failure.

The Outfinitist metacult becomes credible when a member leaves a meeting with more agency than they brought into it.

Service, payment, and value

A movement sustained only by the unpaid labor of a few people either exhausts them or hides who is paying in reality. Infrastructure costs money: platforms, moderation, translation, research, audits, physical spaces, data protection, insurance, legal work, and professional time.

Sustainability is not a commercial impurity. It becomes corrupt when price buys doctrine, rank, access to an alleged salvation, or immunity from criticism.

The exact economic model should be tested rather than proclaimed. A modest membership contribution is one possibility. Sliding contributions, scholarships, institutional subscriptions, paid professional services, donations, and mission-aligned investment may coexist. The design should avoid dependence on a single sponsor and make conflicts visible.

The intellectual core—definitions, major arguments, ethical commitments, critical debates, and enough educational material to understand the

project—should remain publicly accessible. Paid layers should correspond to continuing work that creates identifiable value.

One layer could be a **living library**. It would offer learning paths in probabilistic reasoning, psychology of error, limits of models, institutional economics, organizational design, AI, decentralized brands, and psychological well-being. Materials would display the level of evidence, controversies, sources, and date of last revision. Readers could see not only the conclusion but the history of correction.

A second layer could consist of **learning circles and peer groups**. Founders, researchers, investors, and managers need conversations their organizations do not always permit. A founder cannot discuss every doubt with investors, employees, or family. A well-moderated group can reduce isolation, detect irrational escalation, and offer alternatives before a person consumes health and relationships.

A third layer could offer **limit audits**. The audit would not declare a project absolutely good or bad. It would examine domain, assumptions, metrics, dependencies, risk distribution, appeal mechanisms, reversibility, and stopping conditions. For an AI product, it could include data, evaluations, security, human effects, and concentration. For an organization, it could examine incentives, authority, information flow, and correction. The result would be a map, not a certificate of purity.

A fourth layer could provide **shared infrastructure**: agreement templates, privacy guidance, independent evaluators, mediation, pooled purchasing, secure collaboration, and portable contribution records. A hundred small projects can negotiate together for services they would otherwise purchase at higher cost and with little influence over terms.

A fifth layer could support the **incubation of decentralized brands**. Teams would receive help separating shared identity from the monopoly of one entity, defining contribution and quality, protecting exit, and avoiding both chaos and capture. Projects may take the form of companies, cooperatives, foundations, open protocols, or combinations appropriate to the problem.

A sixth layer could facilitate **prudent project financing**. Membership should not be sold as a lottery ticket, and funding cannot be promised to every member. The network may aggregate donations, service revenue, partnerships, and mission-oriented capital. Projects could receive small tranches linked to observable milestones and reversible stages. Failure would produce a useful account rather than a sanitized story for the next funding round.

Economic value becomes real when members save time, avoid consequential errors, find collaborators, and create shared assets. Social value appears through trust and portable reputation. Psychological value appears when isolation declines and belonging no longer depends on simulating success.

None of these values is guaranteed. The service must be evaluated. Members must be able to leave when the promise is not met. A fee should create no sunk-cost prison.

Payment introduces moral limits. A flat price burdens incomes differently. An entirely free model may depend on wealthy patrons or invisible exploitation of volunteers. A high fee creates a class club. Corporate sponsorship can purchase silence. Advertising turns member attention into inventory.

No formula removes all dependence. The Outfinitist response is diversification, disclosure, caps where appropriate, periodic review, and a prohibition against selling intimate data or doctrinal influence. The organization should publish who finances it, what those funders receive, and which offers it rejects.

A service about limits must be able to show the limits it accepted around revenue.

Financing exploration without rebuilding the casino

A community that funds projects can easily recreate the mythology criticized in Chapter 3. The dramatic founder attracts attention, prestige concentrates around apparent winners, and every grant begins to resemble a ticket to exceptional wealth.

The alternative is not to eliminate risk. It is to design a portfolio in which participants understand the distribution and learning remains valuable when projects end.

Funding should distinguish at least four kinds of project:

Exploration projects test uncertain possibilities cheaply and publish what was learned. Their success is information, not necessarily a company.

Infrastructure projects create tools, standards, archives, or services useful to many nodes. Their value may be difficult to capture privately, which is precisely why common financing is needed.

Enterprise projects pursue sustainable revenue and may attract conventional investment, while accepting shared standards and portable contribution rules.

Protection projects reduce catastrophic or systemic risk—independent evaluation, open alternatives, legal defense, mental-health access, or continuity infrastructure—even when demand does not immediately reward them.

The network should avoid pretending that one investment logic fits all four. Venture capital may suit a rapidly scalable enterprise and distort a public archive. Donations may sustain early exploration and fail to create long-term operational discipline. Membership revenue may finance common services and be too small for capital-intensive research.

Funding decisions should state the theory of value, expected horizon, rights created, failure conditions, and who absorbs loss. The member whose modest contribution supports common learning should not unknowingly subsidize private upside without a declared rule. The investor should not acquire control of the culture merely by financing one project.

Postmortems matter. A network that publishes only successes becomes an advertising agency. One that publishes failure without protecting people becomes a tribunal. The balance is accountability without humiliation: enough detail to show which assumption failed, which signal was ignored, who bore the cost, and what will change.

The Outfinitist metacult should build a reputation for honest endings. In a casino culture, every closing must be narrated as the beginning of a greater victory. In an exploration culture, a well-closed project can be a public good.

What success would look like

Membership count alone cannot measure success. A community with millions of accounts and little trust would reproduce social media. Revenue is not enough. An organization can become profitable by exploiting the anxiety it claims to reduce.

The important indicators concern capacity:

- Do members formulate problems more accurately?
- Do projects identify decisive limits earlier?
- Do people receive support before exhaustion becomes collapse?
- Do collaborations cross professions, regions, and belief systems?
- Are decisions corrected without schism?
- When schisms occur, do people and contributions survive them?
- Do AI systems preserve avenues of appeal, alternatives, and human responsibility?
- Is economic value distributed toward actual contributors?
- Does leadership become more replaceable as the network matures?
- Do members remain active in relationships and institutions outside the Outfinitist metacult?

We should also measure refusal. How much funding was declined because of conditions? Which data were deliberately not collected? Which profitable practices were rejected? Which rules were removed? How often did leaders publish that they did not know? What projects were stopped before external scandal forced the decision?

A civilization obsessed only with accumulation cannot see the value of voluntary boundaries.

At scale, success would mean a network of decentralized brands that provides services, creates employment, supports research, and preserves institutional

alternatives. Some nodes would fail. Others would become highly competent. Common reputation would reduce the cost of trust; autonomy would limit the propagation of error.

It is not necessary for most of humanity to study the philosophy of Outfinitism in depth. For the effect to become civilizational, a significant share of people would need to inhabit institutions that practice its principles: schools that teach the limits of models; firms in which contributors have meaningful channels of challenge; communities that offer psychological well-being without captivity; and AI infrastructures that do not concentrate observation, decision, and sanction in the same hands.

The Outfinitist metacult should not aim to replace every religion, party, university, company, or civic association. It should become connective tissue through which these institutions can develop better reflexes. A member may remain deeply Christian, secular, conservative, progressive, entrepreneurial, academic, or artistic. What changes is the relationship to certainty, power, and correction.

Mass culture need not mean uniform culture.

An invitation to elites

For a person with economic power, this proposal may sound like a request for self-limitation. It is. It is also an offer to acquire something force cannot produce reliably: legitimacy.

Capital requires functioning societies, competence, predictable rules, and people capable of cooperation. The predator model may extract more in the short term while degrading the environment from which value comes. When people experience the system as a machine of humiliation, they become available to leaders promising revenge. Such leaders often use economic elites before disciplining, dispossessing, or replacing them.

Participation in the Outfinitist metacult should not be ceremonial philanthropy. It can provide access to research, evaluation, talent, projects, and a network that treats legitimacy as infrastructure. An investor may find teams more capable of declaring limits. A manager may test distributed

governance without abandoning execution. An owner may build succession and participation before conflict turns transition into war.

The price is the surrender of the fantasy that success justifies everything. A powerful member will be criticized. Conflicts of interest must be explained. Some decisions will not belong to the largest donor. Competence and ambition are not denied; they are converted into service that affected people can recognize.

Future elites can move beyond a truncated reading of *homo homini lupus* without becoming naive. They can recognize aggression, interest, and competition while building institutions in which cooperation is more profitable than betrayal and no predator can capture the entire ecosystem. Their power will come not only from assets, but from maintaining worlds in which others choose to participate.

This invitation is not moral absolutism, and wealth cannot purchase epistemic authority. Entering the Outfinitist metacult should increase exposure to criticism and obligation to service. Legitimate status is earned through verifiable contribution, protection of people with less power, and the ability to leave an institution capable of functioning without you.

The predator elite assumes that manipulable people are a renewable resource. It keeps them confused enough to buy, vote, or work, then frightened enough to request protection from the center. The model functions until a harsher competitor offers a simpler story and permission to turn resentment into violence. History provides no guarantee that the keeper controls the monster he trained.

The Outfinitist elite chooses another advantage. It invests in psychological health, education, mobility, contestable institutions, and common infrastructure because capable people are more valuable partners than dependents. It accepts limits on manipulation, surveillance, and extraction not because power is shameful, but because power without an ecosystem is slow suicide.

The inevitable side of history is not one ideology. It is the accumulating consequence of systems that ignore their limits. Intelligent elites can recognize

that consequence early and participate in building what follows, or they can continue feeding forces they imagine will remain beneath them.

An invitation, not an oath

The tired person who entered our hypothetical circle does not leave converted. She may leave with a better question, two useful contacts, and the rare experience of not having to conceal doubt. She returns when that value proves repeatable. She contributes if others benefit. She becomes a member not because access to elites was promised, but because she found a more decent way to become capable and useful.

This is the elementary unit of a healthy mass movement: not instant allegiance, but the repeated experience of freedom producing capacity. People carry the idea because they have seen it work. The Outfinitist metacult grows through nodes that solve problems, not through campaigns that inflate numbers.

There is no final oath. There are questions with which anyone may begin:

What model am I using, and where does it stop working? What cost of my decision is borne by someone else? What authority have I received, and how can it be withdrawn? What does my position prevent me from seeing? Which alternative must remain able to survive? What have I learned well enough to make another person freer?

A community that continues asking these questions will not eliminate idiocy. Nothing human can. It may prevent idiocy from becoming destiny. It may reduce the amplitude of errors, distribute the capacity for correction, and make exploration a form of belonging.

For that, it is worth attempting.

The choice is repeated in every institution and generation.

We can continue treating the world as audiences to capture, opponents to humiliate, metrics to optimize, and populations to administer. We can hope AI makes this order efficient enough that its contradictions never overtake us.

This is Moloch's hope that speed will solve the race and the Basilisk's hope that a sufficiently intelligent center will eliminate the need for freedom.

Or we can build a culture that preserves conflict without extermination, competence without caste, belonging without captivity, centralization without destiny, and decentralization without chaos. We can create practices that honor correction, institutions that publish limits, brands that no center owns completely, and services that make exploration less lonely. We can make what we refuse—the profitable lie, the easy humiliation, irreversible power, invisible manipulation—a visible part of legitimacy.

The project may fail. The name may be misunderstood. The community may be captured. Other traditions may perform the work better. Accepting that possibility does not weaken the call. It makes the call Outfinitist. We are not required to save a name. We are required to increase humanity's ability to correct the direction in which it is being pushed.

Beyond institutions remains the question no procedure can contain: why does a world exist in which exploration is possible and suffering is so real? We do not know. Perhaps the universe is indifferent. Perhaps there is a Creator whose logic does not fit inside our certainties.

Under either frame, love of exploration remains a dignified choice.

If there is no given meaning, exploration creates islands of meaning among finite beings. If meaning is given, perhaps looking at creation with wonder, refusing the comforting lie, and reducing the suffering we can reduce is a form of participation in it. This is our serious joke: perhaps we align most closely with the Creator's will precisely when we stop pretending to possess it.

Our children do not need us to leave them a finished world. They need a world in which errors cannot grow too large to correct, power cannot close every exit, and curiosity is not punished as disloyalty. They need adults capable of saying "we do not know yet" without abandoning action, and "we were wrong" without collapse.

That is why we must gather.

Not to declare ourselves better than others, but to become together harder to manipulate and more capable of service. Not to win the final debate, but to

preserve the world in which debate, exploration, dignity, and love remain possible.

Truth without culture loses. Culture without limits becomes an idol. Limits without love become a prison.

The Outfinitism Anti-Idiocracy metacult is an attempt to live otherwise.

Bibliography

- Acemoglu, D., and Robinson, J. A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown.
- Adinolfi, G. (2025). *The Rise of Mass Parties, Liberal Italy, and the Fascist Dawn (1919–1924)*. Cambridge University Press. doi:10.1017/9781009564366
- Alboaie, S., Livitckaia, K., and Votis, K. (2023). *Decentralised Brands in Social Media Networks and Beyond*. EU Blockchain Observatory and Forum.
- Alboaie, S., and Cosovan, D. (2017). Private data system enabling self-sovereign storage managed by executable choreographies. In L. Y. Chen and H. P. Reiser (Eds.), *Distributed Applications and Interoperable Systems* (LNCS 10320, pp. 83–98). Springer. doi:10.1007/978-3-319-59665-5_6
- Arrow, K. J. (1963). *Social Choice and Individual Values* (2nd ed.). Yale University Press.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Balan, A., Alboaie, S., Kourtiti, K., and Nijkamp, P. (2023). Blockchain systems for smart cities and regions: An illustration of self-sovereign data governance. In R. Laurini, P. Nijkamp, K. Kourtiti, and L. Bouzouina (Eds.), *Knowledge Management for Regional Policymaking* (pp. 163–190). Springer. doi:10.1007/978-3-031-15648-9_9
- Balan, A., Rată, A., Alboaie, S., Kourtiti, K., and Nijkamp, P. (2023). Sustainability, smart digital cities and decentralised brands. In L. Fusco Girard, P. Nijkamp, and F. Nocca (Eds.), *Reconnecting the City with Nature and History: Towards Circular Regeneration Strategies* (pp. 291–322). FrancoAngeli.
- Baltes, P. B., and Staudinger, U. M. (2000). Wisdom: A metaheuristic (pragmatic) to orchestrate mind and virtue toward excellence. *American Psychologist*, 55(1), 122–136. doi:10.1037/0003-066X.55.1.122
- Basseches, M. (1984). *Dialectical Thinking and Adult Development*. Ablex.
- Bellah, R. N. (1967). Civil religion in America. *Daedalus*, 96(1), 1–21.
- Bengio, Y., Clare, S., Prunkl, C., Murray, M., et al. (2026). *International AI Safety Report 2026* (DSIT 2026/001). International AI Safety Report.
- Brady, W. J., McLoughlin, K. L., Torres, M. P., Luo, K. F., Gendron, M., et al. (2023). Overperception of moral outrage in online social networks inflates beliefs about intergroup hostility. *Nature Human Behaviour*, 7, 917–927. doi:10.1038/s41562-023-01582-0
- Camerer, C. F., and Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1), 306–318. doi:10.1257/aer.89.1.306
- Clark, L., Lawrence, A. J., Astley-Jones, F., and Gray, N. (2009). Gambling near-misses enhance motivation to gamble and recruit win-related brain circuitry. *Neuron*, 61(3), 481–490. doi:10.1016/j.neuron.2008.12.031

- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386–405. doi:10.1111/j.1468-0335.1937.tb00002.x
- Dechêne, A., Stahl, C., Hansen, J., and Wänke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review*, 14(2), 238–257. doi:10.1177/1088868309352251
- Di Rienzo, E. (2020). Le scelte di un liberale conservatore: Benedetto Croce e il fascismo. Una rilettura. *Nuova Rivista Storica*, 104(1), 1–138.
- Durkheim, É. (1995 [1912]). *The Elementary Forms of Religious Life*. Free Press.
- Eurostat. (2025, December 9). Large businesses in the EU: Small number, high turnover. European Commission.
- Fousek, J., Kaše, V., Mertel, A., Výtvarová, E., and Chalupa, A. (2018). Spatial constraints on the diffusion of religious innovations: The case of early Christianity in the Roman Empire. *PLOS ONE*, 13(12), e0208744. doi:10.1371/journal.pone.0208744
- Freeman, J. (1972–1973). The tyranny of structurelessness. *Berkeley Journal of Sociology*, 17, 151–165.
- Girard, R. (1977). *Violence and the Sacred*. Johns Hopkins University Press.
- Girard, R. (1986). *The Scapegoat*. Johns Hopkins University Press.
- Goodhart, C. A. E. (1975). Problems of monetary management: The U.K. experience. In *Papers in Monetary Economics* (Vol. 1). Reserve Bank of Australia.
- Gödel, K. (1931). Über formal unentscheidbare Sätze der *Principia Mathematica* und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38, 173–198. doi:10.1007/BF01700692
- Hall, R. E., and Woodward, S. E. (2010). The burden of the nondiversifiable risk of entrepreneurship. *American Economic Review*, 100(3), 1163–1194. doi:10.1257/aer.100.3.1163
- Hobbes, T. (1998 [1642]). *On the Citizen*. Cambridge University Press.
- Iannaccone, L. R. (1994). Why strict churches are strong. *American Journal of Sociology*, 99(5), 1180–1211. doi:10.1086/230409
- Jackson, J. C., Jong, J., Bilkey, D., Whitehouse, H., Zollmann, S., McNaughton, C., and Halberstadt, J. (2018). Synchrony and physiological arousal increase cohesion and cooperation in large naturalistic groups. *Scientific Reports*, 8, 127. doi:10.1038/s41598-017-18023-4
- Jones, L. E. (2005). Franz von Papen, the German Center Party, and the failure of Catholic conservatism in the Weimar Republic. *Central European History*, 38(2), 191–217. doi:10.1163/156916105775563670
- Kallio, E. (2011). Integrative thinking is the key: An evaluation of current research into the development of thinking in adults. *Theory & Psychology*, 21(6), 785–801. doi:10.1177/0959354310388344
- Keddie, N. R. (2010). Culture and politics in Iran since the 1979 Revolution. In R. Hefner (Ed.), *The New Cambridge History of Islam* (Vol. 6, pp. 438–472). Cambridge University Press. doi:10.1017/CHOL9780521844437.019

- King, P. M., and Kitchener, K. S. (1994). *Developing Reflective Judgment*. Jossey-Bass.
- King, P. M., and Kitchener, K. S. (2004). Reflective judgment: Theory and research on the development of epistemic assumptions through adulthood. *Educational Psychologist*, 39(1), 5–18. doi:10.1207/s15326985ep3901_2
- Kruger, J., and Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6), 1121–1134. doi:10.1037/0022-3514.77.6.1121
- Köllinger, P., Minniti, M., and Schade, C. (2007). I think I can, I think I can: Overconfidence and entrepreneurial behavior. *Journal of Economic Psychology*, 28(4), 502–527. doi:10.1016/j.joep.2006.11.002
- Labouvie-Vief, G. (1994). *Psyche and Eros: Mind and Gender in the Life Course*. Cambridge University Press.
- Meeks, W. A. (1983). *The First Urban Christians: The Social World of the Apostle Paul*. Yale University Press.
- Michels, R. (1915 [1911]). *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Hearst's International Library.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672. doi:10.1257/aer.100.3.641
- Plautus. (1912). *Asinaria*. In P. Nixon (Trans.), *Plautus, Volume I*. Harvard University Press.
- Purzycki, B. G., Apicella, C., Atkinson, Q. D., Cohen, E., McNamara, R. A., Willard, A. K., Xygalatas, D., Norenzayan, A., and Henrich, J. (2016). Moralistic gods, supernatural punishment and the expansion of human sociality. *Nature*, 530, 327–330. doi:10.1038/nature16980
- Rappaport, R. A. (1999). *Ritual and Religion in the Making of Humanity*. Cambridge University Press. doi:10.1017/CBO9780511814686
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118. doi:10.2307/1884852
- Sinnott, J. D. (1998). *The Development of Logic in Adulthood: Postformal Thought and Its Applications*. Plenum Press.
- Sosis, R., and Bressler, E. R. (2003). Cooperation and commune longevity: A test of the costly signaling theory of religion. *Cross-Cultural Research*, 37(2), 211–239. doi:10.1177/1069397103037002003
- Stephan, U., Rauch, A., and Hatak, I. (2023). Happy entrepreneurs? Everywhere? A meta-analysis of entrepreneurship and well-being. *Entrepreneurship Theory and Practice*, 47(2), 553–593. doi:10.1177/10422587211072799

- Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. National Institute of Standards and Technology. doi:10.6028/NIST.AI.100-1
- Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1), 230–265. doi:10.1112/plms/s2-42.1.230
- U.S. Bureau of Labor Statistics. (2024, January 12). 34.7 percent of business establishments born in 2013 were still operating in 2023. *The Economics Daily*.
- Walder, A. G. (2014). Rebellion and repression in China, 1966–1971. *Social Science History*, 38(3–4), 513–539. doi:10.1017/ssh.2015.23
- Weber, M. (1978). *Economy and Society*. University of California Press.
- World Intellectual Property Organization. (2025). *World Intellectual Property Indicators 2025*. WIPO.
- Zingher, J. N., and Flynn, M. E. (2018). From on high: The effect of elite polarization on mass attitudes and behaviors, 1972–2012. *British Journal of Political Science*, 48(1), 23–45. doi:10.1017/S0007123415000514