ChangeThis
264.04

TRIAS ALGORITHMICA:
A SEPARATION OF POWERS IN AI

Hamilton Mann

Let's start with a question that no one is asking: what makes algorithmic power legitimate?

Societies learned to balance three powers, are still learning to balance a fourth, and continue to struggle with a fifth.

Over centuries, systems of governance evolved to stabilize power dynamics through three formal powers—the legislative, executive, and judicial branches. Later, the media became recognized as the "fourth power," followed by information technology, increasingly described as a "fifth power." The latter emerged with the internet and evolved through social networks, democratizing influence toward citizens while simultaneously turning Andy Warhol's prediction that everyone would have fifteen minutes of fame into a reality.

AI INTRODUCES A SIXTH: ALGORITHMIC POWER.

While most conversations about AI focus on how intelligent machines are becoming, who will win the AI race, whether we are approaching Artificial General Intelligence, or whether machines can think or become conscious, I believe we are asking the wrong questions. The most important issue of our time is how a sixth category of power—the algorithmic power—can be governed so that it does not destabilize the very foundations of our societies and democracies.

Every day, algorithms influence what we see, what we believe, who gets hired, who gets credit, what information spreads, and which opportunities become available.

This power is already at work, and every advance we make in AI further amplifies it at a pace whose implications we have not yet fully assessed and at a scale whose capacity to be weaponized may exceed that of any previous form of power.

Understanding this sixth power and determining the conditions under which it can legitimately be exercised are among the defining governance challenges of the twenty-first century.

THE SEPARATION OF ALGORITHMIC POWERS

If algorithmic power is to become a legitimate form of power, it must be subject to the same constitutional safeguards that societies have historically applied to every other form of power to prevent its concentration.

We would never accept concentrated power in government, yet we accept it in AI.

Modern democracies are built on a simple principle: no single institution should write the rules, enforce them, and judge itself. This principle established the foundation of the separation of powers among the three state-level powers—legislative, executive, and judicial—under the intellectual heritage of Montesquieu's *Trias Politica*.

This separation of powers has been one of the most enduring innovations in human societies, a true anti-tyranny technology designed to prevent the harmful consequences that arise when these three powers accumulate and concentrate in the same hands.

Yet this is precisely how many algorithmic systems operate today. The same organizations often design the rules embedded in AI, performing a legislative-like function, deploy those systems at scale, performing an executive-like function, and determine whether they have performed acceptably, performing a judicial-like function. What would be considered dangerous in a democracy—the very concentration of power that the separation of powers was created to prevent—has somehow become normal in AI development.

The most important issue of our time is how a sixth category of power—the algorithmic power—can be governed so that it does not destabilize the very foundations of our societies and democracies.

Algorithmic power is a constitutional category of power and should therefore be subject to the same scrutiny and safeguards that democracies have long applied to political power.

If democracies require Trias Politica to prevent the concentration of political power, the age of AI requires Trias Algorithmica to prevent the concentration of algorithmic power.

Just as constitutional democracy protects freedom by dividing political power, the algorithmic age requires mechanisms that divide and balance algorithmic power. If political power requires checks and balances to prevent tyranny, algorithmic power requires its own constitutional architecture to prevent concentration, preserve contestability, and safeguard human freedom.

WHAT IF WE FAIL

Throughout history, concentrated political power has reshaped markets, institutions, culture, and ultimately the conditions under which individuals exercise freedom.

Concentrated algorithmic power is likely to do the same.

Unlike traditional forms of domination, algorithmic power expands not primarily through coercion but through optimization. Systems become indispensable because they are useful, trusted because they are convenient, and influential because they become the reference points against which decisions, performance, and opportunities are measured.

As this process unfolds, AI systems increasingly become society's benchmark. They progressively establish the standards against which human capabilities are evaluated, whether in interaction, care, love, work, productivity, competence, creativity, or any other domain traditionally mediated by human judgment and once evaluated by people themselves.

What begins as technological optimization gradually becomes social normalization. This gradual redefinition of human worth represents an exercise of algorithmic power itself.

When algorithmic systems shape the standards by which social, economic, and human value are evaluated, they acquire the power to define what counts as valuable human contribution. We have been asking how AI will change society. The more fundamental question is who should possess the authority to establish those standards.

The danger is therefore one of constitutional algorithmic dependence because a society that progressively delegates the power to define, evaluate, and normalize human value to private organizations exercising concentrated algorithmic power progressively loses the capacity to determine for itself what is worth valuing. It ceases to be the author of its own standards. Without that authorship, legitimacy itself begins to collapse.

Preventing that outcome requires an architecture capable of separating algorithmic power.

SEPARATING THE LEGISLATIVE ALGORITHMIC POWERS

First, those who build AI systems should not be the ones who unilaterally define the purposes they are meant to pursue. When the developer of an AI system for allocating healthcare resources decides that minimizing treatment costs should take precedence over maximizing patient outcomes or equitable access, it is defining healthcare priorities rather than merely implementing them. Those priorities should be determined through independent public governance, not by the organization building the system. Accordingly, goal-setting must be separated from machine-making.

Building on this, organizations optimizing AI systems should not determine the limits within which that optimization may operate. When an autonomous vehicle is optimized to minimize journey time, that objective should operate within independently established safety constraints rather than allowing the system to determine for itself what level of pedestrian risk is acceptable. Just as speed limits are set by public authority rather than automobile manufacturers, optimization should remain subordinate to independently defined societal limits. Therefore, training governance must be separated from optimization.

Next, no private organization should be able to convert control over AI infrastructure into lasting market dependence. Rail operators do not build a different rail gauge for every company because common standards allow trains to move across infrastructure owned by different operators. AI systems should likewise be designed around portability standards that allow models and applications to move across infrastructures without being trapped by proprietary technical dependencies. This establishes that market portability standards must be separated from vertical lock-in.

We have been asking how AI will change society. The more fundamental question is who should possess the authority to establish those standards.

SEPARATING THE EXECUTIVE ALGORITHMIC POWERS

Equally, developers of AI systems should not be the sole arbiters of the data from which those systems learn. When the developer of a facial recognition system alone decides which faces are collected, labeled, and retained for training, it also determines which populations the system recognizes reliably and which are disproportionately exposed to error. Independent data stewardship is needed to ensure that representational choices are governed rather than privately determined. Consequently, data stewardship must be separated from developer control.

Yet none of this determines whether AI systems should be permitted to enter society. An aircraft manufacturer does not decide when a new passenger aircraft may begin carrying the public; independent certification determines whether it is safe to enter service. AI systems capable of significantly affecting society should likewise require independent authorization before deployment rather than relying solely on the judgment of their developers. Hence, deployment authorization must be separated from design ownership.

Society must also govern access to transformative capabilities. Just as no private organization may unilaterally acquire and operate facilities capable of enriching weapons-grade nuclear material, no organization should be able to scale frontier AI training simply because it can afford the necessary compute. Once computational capability reaches levels capable of producing systems with significant societal consequences, access should be governed independently of the organizations seeking to use it. For this reason, compute governance must be separated from scale at will.

SEPARATING THE JUDICIAL ALGORITHMIC POWERS

Once these systems enter society, their encounter with reality should not be treated as a mere validation of prior assumptions. Some societal impacts emerge only over time, through widespread adoption and lived experience. An AI hiring system should not be considered successful simply because it reduces recruitment costs or shortens hiring times. Independent monitoring should examine whether its deployment produces persistent patterns of discrimination, occupational segregation, or barriers to social mobility that only become visible after widespread use. This underscores that societal signal monitoring must be separated from engagement optimization.

Finally, no individual should lose the right to challenge decisions simply because they were made by an AI system. If an AI system rejects an applicant for university admission, the applicant should not simply be bound by the algorithm's conclusion. They should have the opportunity to understand the reasons for the decision, contest factual inaccuracies, and seek independent review before educational opportunities are denied. Ultimately, contestability must be separated from algorithmic finality.

Wherever the development and deployment of AI confers the power to define objectives, curate knowledge, constrain optimization, control deployment, mediate societal information flows, allocate opportunities and resources, govern access to computational infrastructure, or determine the conditions of market participation, these powers should never be concentrated in the hands of the same actors who build or operate them.

Individuals must understand how algorithmic power acts upon them in order to preserve their agency. Organizations must retain the capacity to pursue their own legitimate interests without becoming structurally dependent on algorithmic systems. Society must preserve its sovereignty by remaining the author of the norms and values by which it governs itself.

From the individual to the greater collective, the legitimacy of algorithmic power rests on these conditions, all of which depend on the separation of algorithmic powers.

What makes algorithmic power legitimate?

Trias Algorithmica.

Info

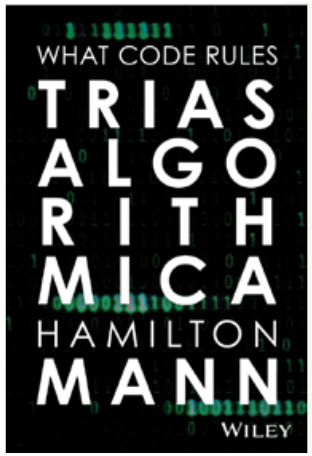

Ready to dig deeper into the book?
Buy a copy of Trias Algorithmica.

Want copies for your organization or for an event?
We can help: customerservice@porchlightbooks.com
800-236-7323

ABOUT THE AUTHOR

HAMILTON MANN is a globally recognized AI thought leader, researcher, executive, lecturer, and award-winning author. He is the originator of Trias Algorithmica, a multidisciplinary framework for governing AI legitimacy, and Artificial Integrity, a pioneering concept and research agenda that seeks to advance AI integrity over intelligence. He serves as Executive Chairman and President of the Artificial Integrity Institute. He also serves as Group Vice President at Thales, where he co-leads the company's AI and digital transformation initiatives worldwide. He teaches at INSEAD and HEC Paris. He was named to the Thinkers50 Radar, which recognizes 30 emerging management thinkers shaping the future of management, and received the Thinkers50 Distinguished Achievement Award in Digital Thinking in 2025. His work has been featured in *Stanford Social Innovation Review*, *California Management Review* and *Rotman Management Magazine*.

SHARE THIS

Pass along a copy of this manifesto to others.

SUBSCRIBE

Sign up for e-news to learn when our latest manifestos are available.

Porchlight

Curated and edited by the people of Porchlight, ChangeThis is a vehicle for big ideas to spread. Keep up with the latest book releases and ideas at porchlightbooks.com.

This document was created on September 23, 2026 and is based on the best information available at that time.