




Repair by Design

Algorithmic Economic Governance, Civil Rights Doctrine, and Structural Discrimination in AI Systems

AI governance • civil rights • economic opportunity



BRITNEY PORTER

July 2026

Hellobritney.info@gmail.com

Distributed for comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies are available from the author.

Contents

A guide to the paper's legal argument and governance framework

00	Executive Summary	1	09	Designing Accountability	17
01	AI Systems as Economic Governance Infrastructure	6	• The Future of Oversight	18	
02	Three New Structural Modes of Algorithmic Discrimination	7	• Legal Triggers and Review Standards	19	
03	The Rule-of-Law Gap	10	• Auditing, Transparency, and Disclosure	19	
04	Emerging Litigation and Algorithmic Liability	12	• Remedies and Economic Repair	20	
05	Algorithmic Economic Exclusion as a Structural Condition	14	• Access to Justice and Legal Infrastructure	20	
06	Fragmentation of Governance and Institutional Limits	14	• Community Participation and Representation	21	
07	Algorithmic Economic Repair	15	10	Transnational Dimensions of Algorithmic Governance	22
08	Institutional Design Implications	16	• Conclusion	22	

WHITE PAPER

Executive Summary

Artificial intelligence systems now function as core infrastructure governing access to economic opportunity in the United States. Across credit markets, employment systems, housing allocation, insurance pricing, and public benefits administration, algorithmic decision systems now structure participation in economic life at scale. While these systems are often described as emergent technological innovations, their governing function is better understood as an extension of longstanding institutional practices through which economic opportunity has historically been organized, classified, and distributed.

Although artificial intelligence represents a novel computational form, the underlying governance logic it operationalizes is not new. Contemporary algorithmic systems extend earlier infrastructures of economic sorting that have long shaped access to housing, credit, labor, and wealth accumulation through formally neutral administrative processes. As *Keeanga-Yamahtta Taylor demonstrates in Race for Profit*, post-civil rights housing and credit markets continued to reproduce racial stratification through ostensibly race-neutral frameworks of risk assessment, underwriting, and market efficiency.¹ For example, underwriting criteria once applied through individualized assessments of borrower risk have been translated into credit-scoring variables and predictive models that automate decisions regarding eligibility, pricing, and access to capital. In that system, discrimination did not require explicit racial classification or intentional

¹ Keeanga-Yamahtta Taylor, *Race for Profit: How Banks and the Real Estate Industry Undermined Black Homeownership* 21–25 (2019).

exclusion. Inequality was produced through institutional arrangements that translated historical disadvantage into economic neutrality, allowing racialized outcomes to persist within formally colorblind administrative and market structures.

Fundamentally, the discriminatory force of these systems resulted from the conversion of historically unequal social conditions into seemingly objective measures of risk, value, and eligibility. Modern algorithmic systems operate through a similar governing logic, where data structures conditioned by historical inequities generate predictive signals that appear neutral in form while continuing to determine the allocation of economic opportunity.

Algorithmic systems continue this logic in computational form by translating historical classifications of risk and value into predictive variables and ranking systems that structure access to economic opportunity. They do not generally rely on explicit or intentional discrimination, but they produce structural inequality through system design, data configuration, optimization objectives, and proxy-based inference. In doing so, these systems inherit and operationalize historical patterns of stratification embedded in prior economic systems, translating them into predictive variables that govern daily access to opportunity. Discrimination therefore frequently emerges without identifiable intent, discrete decision points, or traceable decision makers, even as its distributive consequences remain materially consistent with earlier forms of exclusion. This misalignment exposes a structural limitation in legal systems built around intent, causation, identifiable defendants, and individualized remedies, which struggle to address harm produced through distributed computational processes.

This continuity signals not only a governance gap but also a recursive pattern in how modern institutions rationalize information and interact with technology. Economic systems now translate complex social conditions into simplified computational signals, which are then treated as objective representations of value, risk, or merit. Those signals are reintegrated into decision systems that further entrench the very distributions they purport to measure. This produces a circular structure in which historical inequality is encoded as data, treated as prediction, and reintroduced as justification for future allocation. In this sense, inequality is not only reproduced through institutions but also stabilized through the epistemic assumptions embedded in data-driven governance itself.

At the same time, this circularity is not structurally inevitable. The same technological systems that currently reproduce stratification also contain the capacity for correction if designed and governed differently. Computational systems are not inherently biased or neutral; they are structured through design choices, training data selection, optimization objectives, and institutional incentives. When those components are calibrated toward transparency, contextual reasoning, and equity-aware design, algorithmic systems can function as mechanisms for identifying and correcting structural disparities instead of reproducing them. This is an important distinction because the issue is not the presence of technology, but the architecture within which technology operates. This structural framing creates the basis for the proposed framework of Algorithmic Economic Repair, which discusses the importance of governance in redesigning entire institutional and computational processes that influence how economic opportunity and capital are invested and who has access to the tools, resources, and pathways needed to solve problems and participate in economic life more equitably. That said, because the source of harm is structural; legal repair cannot just be limited to individual remedies.

This repetitive pattern reveals a structural governance gap in contemporary rule-of-law systems. Civil rights enforcement operates through a combination of intent-based constitutional equal protection doctrine, statutory disparate-impact frameworks under *Title VII and the Fair Housing Act*, and consumer protection frameworks addressing unfair or deceptive practices.

Although these frameworks vary in structure and evidentiary standards, they are primarily positioned toward identifying discrete acts, identifiable decision points, and traceable decision-makers through ex post adjudication. This inhibits their capacity to regulate institutional or computational systems that generate distributive inequality through continuous processes of classification, scoring, and optimization. Algorithmic systems, by contrast, produce economic outcomes through ongoing computational processes that operate at scale and cannot be reduced to isolated decision events.

This mismatch produces what this paper defines as algorithmic economic exclusion. This concept refers to structurally produced barriers to wealth creation, mobility, and economic participation generated through computational systems that allocate access to economic opportunity at scale. It builds on *Danielle Keats Citron's* “Technological Due Process,”² research that analyzes how automated systems can embed bias while scaling decision-making in ways that weaken transparency and traditional procedural protections. It also builds on the ideas developed by *Citron and Frank Pasquale* in “The Scored Society,”³ a law review article that shows how automated scoring systems replace personal judgment with predictive categories that determine access to credit, jobs, and other economic resources.

² Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249 (2008).

³ Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014).

These theories center on the problems with due process and algorithmic scoring, while algorithmic economic exclusion describes a broader structural problem in which computational systems control access to economic opportunity through ongoing processes of classification, scoring, and optimization across different areas.

Recent litigation involving Meta Platforms Inc. demonstrates that this condition is already operative within existing legal frameworks. Federal enforcement actions under the *Fair Housing Act*, as interpreted and enforced by the U.S. Department of Housing and Urban Development and the Department of Justice, have treated algorithmic advertising and delivery systems as potential mechanisms of discrimination even absent explicit intent or direct classification.⁴ However, these interventions remain largely constrained to outcome-based remediation and behavioral modification, focusing on correcting observable outputs such as individual decisions or discriminatory effects after they occur.

Algorithmic Economic Repair requires structural accountability directed at system design, training data composition, and optimization objectives, where data variables and related inputs are used in real time to detect disparities, reduce inequitable outcomes, and identify where targeted investment is needed to close economic gaps. These same systems also inform investment-based interventions that account for historical harms, not only preventing future inequities but actively identifying and addressing accumulated patterns of exclusion through corrective resource allocation aimed at restoring balance and expanding economic opportunity.

This paper advances a governance framework titled Algorithmic Economic Repair, which conceptualizes civil rights enforcement, consumer protection, and access to justice interventions

⁴ Fair Housing Act, 42 U.S.C. §§ 3601–3619 (2018); Charge of Discrimination, Sec’y, U.S. Dep’t of Hous. & Urb. Dev. v. Facebook, Inc., FHEO No. 01-18-0323-8 (Mar. 28, 2019).

not only as corrective mechanisms, but also as investments in the long-term stability, inclusiveness, and efficiency of economic systems. Continuing on the foundational framework established in *Repair by Design: Using Consumer Protection, Algorithmic Accountability, and Access to Justice to Address Structural Economic Inequities*,⁵ this paper integrates civil rights doctrine, consumer protection enforcement, algorithmic accountability standards, and access to justice mechanisms into a unified regulatory process designed to govern systems that allocate economic opportunity at scale.

The claim advanced here is that civil rights law must evolve from regulating discriminatory acts to regulating the design of systems that produce distributive economic outcomes, particularly where those systems inherit and reconfigure historically produced inequalities through computational means.

01 AI Systems as Economic Governance Infrastructure

Artificial intelligence systems have become integrated within the infrastructure of economic life.⁶ They now manage access to credit, employment, housing, insurance, and public benefits through automated classification systems that operate continuously and at scale.

⁵ Britney Porter, *Repair by Design: Using Civil Rights Law, Consumer Protection, Algorithmic Accountability, and Access to Justice to Address Structural Economic Inequities* (2026), available at SSRN: <https://ssrn.com/abstract=7159000>.

⁶ Britney Porter, *AI Is Becoming Infrastructure, Not Software: What Stanford's CodeX FutureLaw Reveals About the Next Era of Intelligence*, *Forbes* (Apr. 19, 2026), <https://www.forbes.com/sites/britneyporter/2026/04/19/ai-is-becoming-infrastructure-not-software-what-stanfords-codex-futurelaw-reveals-about-the-next-era-of-intelligence/>.

In comparison to traditional decision-making models, which relied on human judgment, algorithmic systems translate social and economic inputs into probabilistic outputs that determine eligibility, pricing, and access. These systems function both as advisory tools and as allocative infrastructures that structure economic participation.

This represents a major shift in regulatory logic, with society moving away from economic opportunity primarily being controlled through individual decision-making processes and towards computational systems optimized for statistical performance objectives.

In response, inequality is now produced at the level of global system design. This occurs through transnational deployment of similar algorithmic systems across jurisdictions, where models developed in one major sector or society are rapidly replicated and reused in others with limited structural modification. It is further reinforced by internationally shared cloud infrastructure, procurement pipelines, and vendor ecosystems that standardize model architectures across both public and private sectors.

02 Three New Structural Modes of Algorithmic Discrimination

Modern algorithmic discrimination does not appear in explicit or intentional form. It emerges through structural properties of computational systems.

One mechanism is proxy-based classification, in which protected characteristics are reconstructed indirectly through correlated variables such as ZIP code, digital and social behavior patterns, device identifiers, and consumer data signals. Behavioral and lifestyle preferences are not neutral inputs in this structure but are themselves shaped by historically produced conditions of segregation, unequal access to housing and credit, and long-standing

socio-economic inequality. Residential patterns concentrated in specific ZIP codes reflect cumulative effects of past and present discrimination in housing markets, lending practices, employment access, and public investment. These spatial concentrations are then translated into computational markers that stand in for race, class, and other protected characteristics, even where those categories are not explicitly used. As a result, algorithmic systems reproduce the effects of protected-class classification through statistical substitution embedded in both behavioral and geographic proxies.

The National Consumer Law Center, for example, has shown that credit reporting and scoring infrastructures operate as cross-sector gatekeeping mechanisms that extend outside of lending markets to housing, employment, and insurance, while encoding historical patterns of racialized economic exclusion into contemporary screening and eligibility determinations.⁷ This is manifested in persistent racial disparities in credit scores and credit invisibility, where certain groups of consumers are disproportionately more likely to have lower scores or insufficient credit histories, a condition directly linked to structural factors such as redlining, wealth gaps, and historical exclusion from housing and labor markets, as well as biased scoring methodologies. These disparities are not confined to lending decisions but overflow into housing, employment screening, and insurance pricing, where credit-based and proxy-driven evaluations determine access to economic opportunity across multiple domains.

A second mechanism is optimization-driven exclusion. Machine learning systems are designed to optimize objectives such as profit maximization, risk reduction, speed, or efficiency. These

⁷ Nat'l Consumer Law Ctr., *Past Imperfect: How Credit Scores and Other Analytics “Bake In” and Perpetuate Past Discrimination* (Feb. 27, 2024), <https://www.nclc.org/resources/past-imperfect-how-credit-scores-and-other-analyticsbake-in-and-perpetuate-past-discrimination/>.

optimization targets can systematically exclude populations that are statistically associated with higher cost or lower predicted value, even in the absence of discriminatory intent. This dynamic occurs because optimization functions are typically defined in terms of economic or operational performance rather than civil rights values, creating conditions in which efficiency-based objectives can clash with equitable access unless fairness, nondiscrimination, or equity criteria are explicitly incorporated into system design.

A third mechanism is historical data inheritance. Algorithmic systems trained on historical datasets reproduce structural inequalities present in prior housing, credit, labor, and consumer systems. Historical inequality becomes translated into a predictive structure, transforming past distributional injustice into present decision logic informed by severely flawed reasoning patterns that perpetuate how society interprets value, risk, and eligibility. A common defense in discussions of how bias functions within intelligent technology is that these models are simply using accurate historical data. However, accuracy in this setting is not neutral, because the data reflects conditions produced by prior exclusion, unequal access, and pre-existing discriminatory systems. This can lead to a situation in which even accurate predictions reproduce those underlying inequalities by treating them as normal baselines for future decisions and not as problems to be corrected.

Working in tandem, these mechanisms produce a condition in which discrimination operates without intent, classification occurs without transparency, and exclusion is generated through intricate system design.

03 The Rule-of-Law Gap

Civil rights law has historically oscillated between competing doctrinal frameworks that now define the limits of algorithmic governance, organized around three recurring questions: *when liability requires intent, when disparate impact is sufficient to establish a violation on its own, and how remedies apply when discriminatory harm is produced through systemic processes that cannot be attributed to individual decision-making.*

Griggs v. Duke Power Co. established that even employment practices without discriminatory intent can violate Title VII if they exclude protected class members and are unrelated to job performance.⁸ The Court held that employment requirements that disproportionately exclude African American applicants, even when neutral in form, are unlawful unless they are demonstrably related to successful performance of the job. The decision recognized that employment criteria that produce discriminatory effects must be justified by business necessity and grounded in a direct relationship to the skills required for the position.

Washington v. Davis, in turn, requires proof of discriminatory intent for constitutional liability under the Equal Protection Clause, limiting structural claims in constitutional contexts.⁹ This doctrinal limitation was further reinforced in *McCleskey v. Kemp*, where the Supreme Court declined to invalidate a criminal sentence based on statistical evidence demonstrating racial disparities absent proof of discriminatory purpose in the individual case.¹⁰ The decision exposed a more structural limitation within existing legal frameworks. Systemic disparities may be

⁸ *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

⁹ *Washington v. Davis*, 426 U.S. 229, 239 (1976).

¹⁰ *McCleskey v. Kemp*, 481 U.S. 279, 292–93 (1987).

measurable and empirically documented yet remain difficult to challenge when legal accountability depends upon identifying a specific discriminatory actor or intent.

Algorithmic systems do not necessarily create a new doctrinal problem. Instead, they continue to expose and intensify limitations that have long existed within civil rights enforcement when discriminatory outcomes emerge through interconnected institutional processes that were never fully addressed by legal frameworks focused primarily on discrete acts of discrimination. In computational environments, responsibility may be distributed across developers, vendors, deployers, data brokers, and institutional users, making harmful outcomes more difficult to identify, define, challenge, and remedy. These systems make discriminatory patterns less visible, normalize exclusion through routine decision-making, obscure causal connections between conduct and harm, and reproduce inequitable outcomes at a speed and scale once difficult to achieve through traditional forms of decision-making. Nevertheless, the underlying legal challenge remains familiar. As demonstrated by McCleskey, statistical disparities and systemic patterns of exclusion have historically struggled to generate liability when discriminatory intent cannot be attributed to a particular actor or decision. This creates a persistent gap between the ability to document discriminatory outcomes and the ability to secure legal accountability and corrective intervention.

Algorithmic systems continue to expose the unresolved tension between these frameworks. In computational environments, discriminatory intent is typically absent, distributed, or non-determinable. Harm continues to be perpetuated and amplified through system optimization, data selection, model training processes, and historical data inheritance instead of individualized biased decision-making events.

Consequently, algorithmic governance produces outcomes that mirror the discriminatory patterns examined in *Griggs* and *McCleskey*, while often falling outside the enforcement boundaries established by *Washington v. Davis*. This creates a structural rule-of-law gap in which discriminatory outcomes exist without stable doctrinal mechanisms for accountability.

04 Emerging Litigation and Algorithmic Liability: When Systems Become Defendants

Recent federal enforcement actions involving Meta Platforms, Inc. demonstrate that algorithmic systems are already being treated as potential mechanisms of civil rights discrimination under existing law.¹¹

In *United States v. Meta Platforms, Inc.* (DOJ enforcement under the Fair Housing Act), federal authorities alleged that Meta's advertising delivery systems enabled discriminatory targeting and delivery of housing advertisements through machine-learning algorithms to determine which users received housing advertisements in ways that produced disproportionate outcomes among protected groups. More importantly, this case demonstrates that the alleged discrimination did not arise from an advertiser intentionally selecting or excluding protected populations. Instead, the algorithmic delivery system itself became the site of legal scrutiny. Even where advertisers selected broad audiences, Meta's optimization mechanisms still determined which individual users were more likely to receive certain advertisements based on patterns identified through large-scale behavioral and demographic signals. The legal theory also did not depend on explicit intent but instead relied on disparate impact analysis, recognizing that algorithmic systems can produce discriminatory effects through optimization and delivery mechanisms.

¹¹ Complaint, *United States v. Meta Platforms, Inc.*, No. 1:22-cv-05187 (S.D.N.Y. June 21, 2022).

The settlement required modifications to Meta’s advertising systems and ongoing monitoring under court supervision. As part of the settlement, Meta agreed to discontinue its Special Ad Audience tool for housing, employment, and credit advertisements and develop a new machine-learning system designed to reduce disparities in ad delivery. The company was also required to implement independent monitoring and ongoing reporting mechanisms under court supervision. These requirements represent one of the first major federal interventions, treating algorithmic optimization systems as objects of civil rights regulation. This enforcement action is pivotal because it confirms that algorithmic ad delivery systems can constitute civil rights-relevant infrastructure under current laws.¹²

However, the enforcement structure remains limited. The remedy focuses on modifying observable outputs without examining system design, training data composition, or optimization objectives. This reflects a significant limitation in current doctrine. Courts can recognize algorithmic discrimination under disparate impact theory but lack a consistent framework for regulating the structural design of systems that produce it. Algorithmic Economic Repair is designed to address this gap by directing accountability toward the institutional, computational, and economic arrangements that generate these outcomes, ensuring that accountability matches the level at which exclusion is often produced and that these patterns can be measured, monitored, and evaluated in real time.

Additional litigation involving platform-based algorithmic systems in employment, content moderation, and automated enforcement contexts reinforces this emerging trajectory. In these domains, courts and regulators are confronting systems that function as gatekeepers to economic

¹² Settlement Agreement & Final Judgment, *United States v. Meta Platforms, Inc.*, No. 1:22-cv-05187 (S.D.N.Y. June 27, 2022).

and civic participation. Enforcement remains fragmented and reactive, instead of oriented toward system governance.

05 Algorithmic Economic Exclusion as a Structural Condition

Throughout consumer and labor markets, AI systems produce what this paper defines as algorithmic economic exclusion. This condition reflects structurally embedded barriers to economic participation generated through computational decision systems that govern access to credit, employment, housing, insurance, public investment, and community resources that facilitate economic growth and mobility.

Just like traditional forms of discrimination, algorithmic economic exclusion is cyclical and systemic. It operates through eligibility scoring that converts historical inequality, social hierarchies, and behavioral data into predictive classifications, which are then structured through optimization-based ranking systems that determine access to resources and opportunities in real time. This shifts the nature of inequality from episodic exclusion to system-level economic organization. Economic opportunity is no longer allocated through individual institutional decisions but through computational systems embedded in market formation.

06 Fragmentation of Governance and Institutional Limits

Regulatory oversight of algorithmic systems is currently fragmented throughout civil rights enforcement agencies, consumer protection authorities, and sector-specific regulators. Each

framework addresses a component of algorithmic harm, but none govern AI systems as integrated allocative infrastructures.

Civil rights law operates primarily through ex post litigation. Consumer protection law focuses on unfair or deceptive practices but is not structured to evaluate systemic distributive outcomes. Administrative agencies regulate domain-specific harms but lack coordinated authority over cross-sector algorithmic systems.

This creates a critical question: what regulatory framework governs systems that gatekeep economic opportunity across civil rights, consumer protection, and access to justice domains at the same time?

The absence of coordinated oversight creates a structural governance gap in which systems that allocate economic opportunity are not regulated together under civil rights law.

07 Algorithmic Economic Repair: A Design-Based Approach to AI Governance

The concept of algorithmic economic repair introduces a unified governance framework for regulating AI systems as systems that distribute economic opportunity in ways that reflect the nuanced needs, contexts, and lived realities of the communities they serve, while also enabling data sovereignty and direct participation in system design, technological development, and economic mobility.

Under this model, algorithmic systems are not treated as neutral decision tools but as economic governance systems subject to civil rights constraints at the level of design. Data inputs, scoring systems, and optimization processes are restructured to function as mechanisms that expand economic opportunity instead of reinforcing cycles of exclusion and inequality. This produces a circular structure in which historical inequality is encoded as data, treated as prediction, and reintroduced as justification for future allocation. Regulation therefore shifts from reactive enforcement toward structural accountability embedded within system design.

To achieve this, algorithmic economic repair merges civil rights doctrine, consumer protection, algorithmic accountability, and access to justice systems into a single governance architecture. This framework ensures that systems shaping economic opportunity are auditable, structurally non-discriminatory, and engineered to deliver measurable equity.

08 Institutional Design Implications

Operationalizing algorithmic economic repair requires pivoting from outcome-based enforcement toward system-level oversight.

AI systems used in consumer markets should be subject to pre-deployment evaluation when they are deployed in high-impact domains that are responsible for consumers' access to employment, housing, credit, insurance, or other forms of economic opportunity. Factors such as scale, predictive decision-making authority, and the potential to generate disparate impacts should inform whether review is required. These systems must also be auditable in ways that allow

regulators to evaluate whether proxy variables, optimization objectives, or training datasets are producing systemic exclusion.

Where systems cannot be transparently audited, their use in high-impact economic domains should raise immediate concerns regarding regulatory risk. Civil rights enforcement, consumer protection authorities, and sector-specific regulators must operate in a coordinated fashion to address algorithmic systems as unified infrastructures and not isolated tools.

Within this type of structure, reparative programs function as an extension of oversight functions and as solutions designed to promote mobility and market access. This includes institutionally administered instruments that respond to measurable patterns of exclusion through targeted investment, resource distribution, and corrective interventions designed to address inequities produced or reinforced by automated systems.

Lastly, anyone affected by unfair automated determinations must retain the ability to contest those decisions and access necessary resources without barriers, while having clear and simplified pathways to explanation, review, and correction processes that are timely and enforceable in practice instead of only procedural in form.

09

Designing Accountability: Key Questions for Legal Professionals, Policymakers, Regulators, and Communities

Although this paper offers a general framework for how governing systems can better allocate economic opportunity in ways that mitigate discriminatory outcomes, many institutional design questions remain unanswered. These questions must be explored through more than a

technological lens. They require coordinated attention to civil rights enforcement, consumer protection, administrative law, access to justice, legal services infrastructure, data governance, community participation, and democratic accountability. Effective governance requires a coordinated ecosystem capable of identifying, monitoring, challenging, and remedying algorithmic harms across housing, employment, credit, insurance, small business and community investment, and other varying domains influencing economic mobility.

As policymakers, regulators, legal professionals, technologists, civil rights organizations, community advocates, and affected populations consider implementation pathways, several foundational questions emerge.

● The Future of Oversight

Who should hold primary oversight authority for algorithmic systems operating across multiple sectors?

Enforcement is currently fragmented across civil rights agencies, consumer protection regulators, administrative bodies, and sector-specific institutions. Coordinating responsibilities, information sharing, and accountability remains a central governance challenge.

This fragmentation also creates an access to justice challenge, as individuals seeking to assert their rights in response to algorithmic harm may struggle to identify the appropriate forum, navigate overlapping enforcement pathways, secure appropriate representation, and access the necessary tools to pursue relief or obtain remedies across dispersed systems that may not even exist in a coherent or accessible form.

- Legal Triggers and Review Standards

What legal threshold should trigger review of high-impact algorithmic systems?

Should possible triggers include deployment in regulated sectors, evidence of disparate impact, influence over access to economic opportunity, scale of affected populations, or use in decisions involving essential economic rights and benefits?

- Auditing, Transparency, and Disclosure

What standards should govern algorithmic audits, and who should conduct them?

Ongoing issues concern measuring proxy discrimination, transparency, explainability, inherited bias in training data, and ongoing system performance in the context of opaque models where decision pathways and data sources are not fully observable or traceable and offer limited visibility into decision logic.

Related questions concern what documentation, provenance information, and disclosures organizations must maintain and provide to regulators, auditors, advocates, and affected individuals so that system behavior is meaningfully auditable and understandable in practice, while protecting privacy and legitimate proprietary interests without compromising accountability.

● Remedies and Economic Repair

What remedies should follow when algorithmic systems produce systemic exclusion while also protecting economic stability and access for citizens?

Options include system modification, independent monitoring, corrective investment, targeted resource allocation, community-based reparative measures, or deployment restrictions.

Another core design question is when relief should focus on individual redress, institutional reform, or broader economic repair, along with reparative approaches that address both current harms and the cumulative effects of past and ongoing exclusion. These approaches must be understood as an investment function in system stability that can be measured and help prevent the reinforcement of persistent inequities over time while strengthening long-term economic inclusion and efficiency.

● Access to Justice and Legal Infrastructure

How should access to justice systems evolve to address algorithmic harms?

Current enforcement often depends on individuals recognizing harm, securing legal assistance, and navigating complex processes.

For underrepresented communities, barriers to legal representation can limit the ability to assert rights and recover lost economic opportunity, further reinforcing unequal access to economic systems. This raises questions about how to expand legal aid capacity, redesign administrative complaint systems, and create sustained investment in post-law-school pipelines such as funding,

fellowships, incubation models, and institutional support that enable attorneys to serve underrepresented communities and build practices focused at the intersection of access to justice, economic justice, civil rights, consumer protection, and emerging computational harms at scale.

● Community Participation and Representation

What role, then, should affected communities play in governing systems that continue to shape economic opportunity?

Strategic participation may require more than consultation or conversations, along with procedural safeguards, data governance rights, and direct, interdisciplinary representation in both oversight structures and technology development. Communities most impacted are often least represented in design, procurement, and enforcement, and that imbalance must be addressed through more inclusive governance and sustained investment in pathways that support historically and currently underrepresented communities in entering, contributing, and developing these systems.

These questions are open design challenges requiring deeper analysis and interdisciplinary collaboration across law, policy, technology, civil rights, and community institutions. The legitimacy of algorithmic governance will depend on a shift in how these systems are understood and conceptualized, as well as on whether institutions develop the capacity and organizational structure needed to ensure sustainability, accountability, and transparency function as an integrated ecosystem instead of fragmented efforts.

Transnational Dimensions of Algorithmic Governance

Furthermore, AI systems operate across jurisdictions through global infrastructure such as cloud computing systems, multinational platforms, and cross-border financial and labor markets. In many cases, algorithmic systems trained in one regulatory environment are deployed in others without structural modification. They may also be adopted or procured by government institutions without adequate attention to cultural and contextual differences.

This produces a transnational replication of inequality in which structural biases originating in one jurisdiction are reproduced through technological deployment across others. Algorithmic governance therefore exists at the intersection of domestic civil rights enforcement, international regulatory coordination, and human rights commitments governing equitable access to capital, opportunity, and resources.

FINAL SECTION

Artificial intelligence systems no longer influence market outcomes; they function as the governing infrastructure of economic life throughout the world. AI systems now mediate access to credit, housing, labor, and public programs, and in doing so they structurally shape the terms of economic participation at scale.

As this paper has demonstrated, the governance logic underlying these technologies is not an emergent phenomenon but a digital extension of historical institutional practices. Just as post-civil rights markets cloaked racial stratification in the neutral language of risk assessment and underwriting, contemporary algorithmic platforms translate historical inequality into

predictive variables under the appearance of mathematical objectivity. By encoding past distributional injustice as data inputs, optimization objectives, and proxy variables; these systems produce a recursive operating model that stabilizes and reproduces systemic exclusion.

This structural reality reveals a persistent rule-of-law gap. While doctrinal developments from *Griggs* to *McCleskey* reflect an unresolved tension between intent-based liability and statistical disparity, recent enforcement actions such as *United States v. Meta Platforms, Inc.* indicate that federal authorities are beginning to recognize automated delivery systems as mechanisms of civil rights discrimination under disparate impact theory. However, current judicial and regulatory remedies remain limited; they are fragmented, reactive, and confined to output-based modification rather than system-level accountability.

To preserve rule-of-law legitimacy in an AI-driven economy, regulatory frameworks must move beyond individualized, ex post adjudication. Algorithmic economic exclusion is a systemic condition, and it requires a system-level response.

The framework of Algorithmic Economic Repair offers an integrated, ex ante governance ecosystem that unites civil rights doctrine, consumer protection enforcement, economic justice, algorithmic accountability standards, and access to justice pathways. Implementing this model requires a fundamental paradigm shift from reactive, linear, outcome-based enforcement toward interdisciplinary oversight. This type of coordinated structure would consist of civil rights lawyers, consumer protection regulators, data scientists, technologists, procurement officials, access to justice practitioners, etc., with the composition of these teams varying depending on the nature of the harm, the affected populations, and the economic or regional context in which the

system is deployed, as well as the specific problems being addressed. From this perspective, oversight cannot be structured as a fixed institutional template but must be adaptable to the characteristics of the problem at hand.

Adaptation is also required to strengthen access to justice, which must be treated as a systemic design issue rather than a downstream service function. This challenge demands investment in and protection of legal talent itself, where barriers to entry, discrimination in legal employment, and unequal access to institutional pathways can limit the capacity of the legal system to respond effectively.

It also requires ensuring that legal professionals from diverse backgrounds are positioned to address challenges rooted in the communities they come from and are often motivated to serve. For these reasons, regulatory design may require specialized units within larger agencies but cannot rely on a one-size-fits-all institutional model given the variability of algorithmic harms across domains and populations.

If technology is to expand economic opportunity rather than automate exclusion, its foundational components must be subject to pre-deployment evaluation, structural auditability, and enforceable pathways for human review, appeal, and correction of automated decisions. In this sense, civil rights law must expand its scope from the regulation of isolated discriminatory acts to the governance of the systemic and computational infrastructures that determine economic access and opportunities.

Algorithmic Economic Repair operationalizes this structural transition by establishing a governance framework in which civil rights enforcement, consumer protection, algorithmic accountability, and access to justice function as interconnected components of system-level

regulation and correction. It redefines economic exclusion as a structural condition arising from the design and implementation of computational systems that dictate access to economic opportunities. In doing so, it aligns legal accountability with the level at which economic exclusion is generated.