# SUPREMO AMICUS

## INDIA'S FIRST AI INTEGRATED LAW JOURNAL

**Peer Reviewed, Refereed and Open access Journal**

- Available in 331+ International Libraries
- Indexed at 32 Databases


ISSN NO. 2456-9704
**Volume 10 Issue 1**
www.supremoamicus.org

# DISCLAIMER

The information presented in this article is intended for general informational and educational purposes only. While every effort has been made to ensure that the content is accurate, up-to-date, and reliable at the time of publication, the editorial board and publisher make no representations or warranties of any kind, express or implied, regarding the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

The views and opinions expressed in this article are those of the author and are based on personal research, experience, and interpretation. They do not necessarily reflect the official policy, position, or opinions of any affiliated organization, institution, or entity.

This article is not intended to serve as professional advice of any kind. The editorial board and publisher shall not be held liable for any errors or omissions in the content, nor for any losses, injuries, or damages arising from the use of or reliance on this information.

# ABOUT THE JOURNAL

Supremo Amicus is an online, peer-reviewed international journal devoted to the interdisciplinary fields of law and science. In an era marked by rapid technological progress and evolving legal frameworks, the journal seeks to bridge the gap between these two dynamic domains by offering comprehensive and critical insights into their various aspects. The journal places a strong emphasis on contemporary advancements, emerging trends, and the complex challenges faced by both the legal community.

The primary objective of the journal is to encourage and promote original, high-quality research. It is committed to publishing well-researched, analytically sound, and thought-provoking articles that adhere to rigorous academic standards. Each submission undergoes a thorough peer-review process to ensure authenticity, relevance, and scholarly integrity. In doing so, the journal maintains its commitment to excellence and credibility.

In addition to fostering research, the journal aims to make complex ideas accessible and engaging for a diverse readership. It strives to present content that is not only intellectually enriching but also clearly written and reader friendly.

Furthermore, the journal is committed to promoting interdisciplinary collaboration and global engagement. It welcomes diverse perspectives from contributors across different regions and backgrounds, thereby enriching the quality and scope of discussions presented within its pages.

With this vision we proudly present Supremo Amicus to our readers.

**-Editorial Team**
**Supremo Amicus**

# AUTOMATED GOVERNANCE AND PROCEDURAL FAIRNESS: RETHINKING DUE PROCESS IN THE AGE OF ARTIFICIAL INTELLIGENCE

**By** *Apala Ghosh*
**PhD Research Scholar**, *Department of International Relations, Jadavpur University*

**Abstract**
The increasing integration of artificial intelligence (AI) into governance processes is reshaping how decisions are made, implemented, and experienced by individuals. Across domains such as welfare distribution, digital identification, and regulatory enforcement, automated systems are being used to assist or replace human decision-making. While these systems offer efficiency and scalability, they raise significant challenges for established principles of procedural fairness. Traditional legal frameworks assume that decisions are made by identifiable authorities, can be explained through reasoned justification, and are subject to contestation. AI-driven systems complicate these assumptions by introducing forms of decision-making that are data-driven, probabilistic, and often opaque.

This article examines the tension between automated governance and procedural fairness, with a particular focus on the Indian context. It argues that existing approaches to due process are not fully equipped to address the structural characteristics of AI-based systems. Rather than treating these challenges as isolated technical issues, the paper situates them within broader transformations in governance, where authority is increasingly exercised through systems rather than institutions. Drawing on examples from digital identification systems and algorithmic decision-making, the analysis highlights the limitations of current safeguards and explores the implications for accountability, legitimacy, and citizen-state interaction. The article concludes by outlining the need to rethink procedural fairness in ways that reflect the realities of automated governance.

**Keywords:** Artificial Intelligence; Procedural Fairness; Due Process; Algorithmic Governance; India; Automated Decision-Making

## 1. Introduction: When Decisions Are No Longer Human

The incorporation of artificial intelligence into public administration has altered the conditions under which state decisions are made. Across a range of domains—welfare allocation, identity verification, regulatory enforcement—automated systems are increasingly used to process information, assess eligibility, and guide administrative outcomes. These developments are often framed in terms of efficiency, scalability, and the ability to manage large volumes of data. While such benefits are not insignificant, they tend to obscure a more fundamental shift: the relocation of decision-making from identifiable human actors to complex socio-technical systems.

This shift raises questions that extend beyond technical performance or policy effectiveness. At stake is the structure of decision-making itself. Traditional administrative processes are premised on the assumption that decisions are made by human authorities who can be identified, whose reasoning can be articulated, and whose actions can be challenged. These assumptions are embedded within broader principles of procedural fairness, which seek to ensure that decisions affecting individuals are made in a manner that is transparent, reasoned, and subject to contestation.

Automated decision-making complicates each of these elements. When decisions are generated or shaped by algorithmic systems, the basis on which they are made may not be easily reducible to a form of reasoning that can be communicated or scrutinised in conventional terms. Decision-making processes may involve the interaction of multiple components, including data inputs, model architectures, and optimisation criteria, none of which correspond neatly

to the notion of a decision-maker as understood in legal frameworks. As a result, the relationship between decision, reasoning, and responsibility becomes less direct.

The growing use of such systems in governance contexts has been accompanied by debates over transparency, bias, and accountability. While these concerns are important, they often focus on specific attributes of AI systems rather than on the broader transformation of procedural structures. The central issue is not only whether automated systems produce accurate or fair outcomes, but whether existing frameworks of procedural fairness are capable of addressing the conditions under which those outcomes are produced.

This article examines the tension between automated governance and procedural fairness, with particular attention to the Indian context. It does not proceed on the assumption that the use of AI in governance is inherently problematic or that it should be abandoned. Rather, it treats automation as a structural development that requires analytical engagement. The question is not whether AI should be used, but how its use interacts with principles that have historically governed administrative decision-making.

The argument advanced here is that automated systems challenge procedural fairness not by violating it in discrete instances, but by operating on assumptions that diverge from those underlying existing legal frameworks. Procedural fairness presumes that decisions can be explained, that those affected can meaningfully respond, and that responsibility can be assigned to identifiable actors. Automated systems, by contrast, often produce outcomes through processes that are statistical, distributed, and continuously evolving. This creates a mismatch between the logic of the system and the expectations of the law.

To develop this argument, the article proceeds in several stages. It begins by outlining the conceptual foundations of procedural fairness, situating it within both legal doctrine and broader governance practice. It then examines the rise of automated decision-making in public administration, identifying the characteristics that distinguish it from traditional forms of decision-making. The analysis subsequently turns to the points of tension between these systems and procedural norms, focusing on issues of reasoning, responsibility, and individual assessment. Drawing on examples from digital identification systems and automated welfare administration, the article illustrates how these tensions manifest in practice. It then considers the implications for governance, before concluding with an exploration of how procedural fairness might be rethought in light of these developments.

The aim is not to resolve the tension definitively, but to clarify its structure. In doing so, the article seeks to contribute to a more grounded understanding of how automation reshapes the conditions under which governance is exercised, and how principles such as fairness must adapt in response.

## 2. Understanding Procedural Fairness

### 2.1 Foundations of Due Process

Procedural fairness is often treated as a technical requirement within legal systems, associated with specific rules governing how decisions should be made. However, its significance extends beyond formal compliance. At its core, procedural fairness reflects a broader commitment to ensuring that the exercise of authority is constrained by standards of reasonableness, transparency, and participation.

In its most basic formulation, procedural fairness requires that decisions affecting individuals are not arbitrary. This principle is typically expressed through a set of minimum safeguards. These include the opportunity to be heard, the obligation to provide reasons, and the possibility of review or challenge. While the precise articulation of these safeguards varies across jurisdictions, their underlying rationale remains consistent: individuals should not be subject to decisions that they cannot understand, contest, or engage with.

These requirements are not merely procedural in a narrow sense. They are tied to deeper questions of legitimacy. A decision that is procedurally fair is more likely to be accepted, even when its outcome is unfavourable. This is because the process through which the decision is made signals that the individual has been treated as a participant in the system, rather than as an object of administrative action. Procedural fairness thus functions as a bridge between authority and acceptance, linking the exercise of power to its justification.

Importantly, procedural fairness assumes a particular model of decision-making. It presupposes that decisions are made by identifiable actors who can consider relevant information, weigh competing considerations, and articulate reasons for their conclusions. It also assumes that the process through which decisions are made can be reconstructed, either by the decision-maker or by an external body reviewing the decision. These assumptions are rarely stated explicitly, but they underpin the operation of procedural safeguards in practice.

**2.2 Procedural Fairness in the Indian Context**

In the Indian legal framework, procedural fairness is closely associated with the constitutional guarantees of equality and personal liberty. Judicial interpretation of Articles 14 and 21 has played a central role in embedding principles of fairness within administrative action. Over time, courts have moved beyond a narrow reading of these provisions to emphasise that state action must be not only lawful, but also reasonable and non-arbitrary.

This development is reflected in the incorporation of principles of natural justice into administrative decision-making. The requirement that individuals be given an opportunity to be heard, and that decisions be supported by reasons, has been treated as an essential component of fair procedure. These principles are not confined to formal adjudicatory settings; they extend to a wide range of administrative actions that affect rights, entitlements, and access to public resources.

At the same time, the application of procedural fairness in India has been shaped by practical considerations. Administrative systems often operate under conditions of scale, resource constraint, and urgency. As a result, procedural requirements have been applied flexibly, with courts recognising that the content of fairness may vary depending on context. What remains consistent, however, is the expectation that state action should not be arbitrary, and that individuals should have some meaningful avenue to understand and respond to decisions that affect them.

This balance between principle and practicality is particularly relevant in the context of automated governance. The Indian state has been an active adopter of digital systems in areas such as identity management and welfare delivery. These developments have been justified in part by the need to manage large-scale administrative processes more efficiently. However, they also raise questions about how procedural fairness is to be maintained when decision-making is mediated through systems that operate differently from traditional administrative structures.

**2.3 Why Procedure Matters in Governance**

Procedural fairness is sometimes treated as secondary to substantive outcomes. If a decision is correct or efficient, the process through which it is reached may appear less important. This view, however, underestimates the role that procedure plays in structuring governance.

First, procedural safeguards function as a constraint on the exercise of power. By requiring that decisions be justified and open to challenge, they limit the scope for arbitrary action. This is particularly important in administrative contexts, where decisions are often made by officials with discretionary authority and where the consequences for individuals can be significant.

Second, procedure contributes to the reliability of decision-making. The requirement to consider relevant information, to hear affected parties, and to

provide reasons creates a framework within which errors can be identified and corrected. While procedural fairness does not guarantee correct outcomes, it increases the likelihood that decisions will be made on a rational basis.

Third, procedural fairness has a relational dimension. It shapes how individuals experience governance. A system that allows for participation and explanation is more likely to be perceived as legitimate, even in the absence of perfect outcomes. Conversely, a system that produces decisions without explanation or opportunity for engagement risks generating distrust, even if its outcomes are broadly accurate.

These functions are closely tied to the assumptions outlined earlier: that decisions are made by actors who can engage with information, that processes can be articulated, and that individuals can respond. When these assumptions are altered, the role of procedural fairness becomes less clear. Automated systems, which operate through data processing and algorithmic optimisation, do not fit easily within this framework. They do not "hear" in the conventional sense, nor do they "reason" in a way that can be readily communicated.

This does not mean that procedural fairness is irrelevant in such contexts. On the contrary, it suggests that its application must be reconsidered. If the conditions under which decisions are made have changed, then the mechanisms through which fairness is ensured may also need to evolve. Understanding this shift requires first examining the nature of automated governance itself.

## 3. The Rise of Automated Governance

### 3.1 What is Automated Decision-Making?

Automated decision-making in governance refers to the use of computational systems to assist, structure, or in some cases determine administrative outcomes. These systems range from relatively simple rule-based processes to more complex models that rely on statistical inference and machine learning. What unites them is not a single technological form, but a shift in how decisions are produced: from deliberation centred on individual cases to processes that rely on data aggregation, pattern recognition, and standardised criteria.

It is important to distinguish between different degrees of automation. In some contexts, automated systems function as decision-support tools, providing recommendations that are reviewed by human officials. In others, the role of human intervention is minimal, with systems generating outputs that are implemented directly or subject to only cursory verification. The boundary between these categories is not always clear in practice. Systems described as "assistive" may exert significant influence over outcomes, particularly where administrative capacity constraints limit the scope for independent review.

Automated systems operate through a combination of data inputs, model structures, and decision thresholds. Data is collected from various sources, processed into a usable format, and analysed to produce outputs that are treated as indicators of eligibility, risk, or compliance. These outputs are not decisions in themselves, but they function as decisive elements within administrative processes. Once incorporated into institutional workflows, they shape how decisions are made and how discretion is exercised.

What distinguishes these systems from traditional administrative tools is not simply their scale or speed, but the nature of their reasoning. Human decision-making typically involves the interpretation of rules in relation to specific circumstances, with the possibility of deviation where justified. Automated systems, by contrast, rely on predefined parameters and statistical relationships. They process cases as instances within a dataset, rather than as discrete situations requiring contextual judgement. This does not necessarily render them inaccurate, but it alters the basis on which decisions are reached.

### 3.2 Domains of Application

The use of automated systems in governance is not uniform, but it is expanding across multiple domains

where large volumes of data must be processed and administrative decisions must be made at scale.

One prominent area is welfare administration. Governments increasingly rely on digital systems to determine eligibility, detect irregularities, and manage distribution mechanisms. These systems are designed to standardise decision-making and reduce administrative burden. However, their operation often depends on the integration of multiple datasets and the application of criteria that may not fully capture the variability of individual circumstances.

Digital identity systems provide another important example. In India, the Aadhaar system has been integrated into a wide range of administrative processes, including welfare delivery, banking, and public service access. Authentication through biometric or demographic data has become a precondition for accessing certain entitlements. While the system has been presented as a means of improving efficiency and reducing duplication, its reliance on automated verification processes introduces new forms of dependency on system performance.

Automated systems are also used in regulatory and enforcement contexts. These include systems designed to flag suspicious transactions, identify potential fraud, or prioritise cases for investigation. In such settings, automation is often justified on the basis that it allows authorities to focus limited resources on higher-risk cases. However, the criteria used to define risk and the thresholds for intervention are embedded within system design, rather than being determined through case-by-case assessment.

Across these domains, a common feature is the integration of automated outputs into decision-making processes that affect rights, entitlements, and obligations. The extent to which these outputs are treated as authoritative varies, but in many cases they function as default indicators that shape administrative action.

**3.3 Drivers of Adoption**

The expansion of automated governance is driven by a combination of administrative, economic, and political factors. One of the primary drivers is scale. Modern administrative systems are required to process large volumes of applications, transactions, and interactions. Manual processing is often resource-intensive and time-consuming, creating incentives to adopt systems that can operate more efficiently.

Cost considerations also play a role. Automated systems are frequently presented as a means of reducing administrative expenditure, particularly in contexts where public resources are constrained. By standardising processes and reducing reliance on human labour, such systems are seen as a way to achieve greater output with fewer inputs.

Another factor is the perceived objectivity of data-driven systems. Automated decision-making is sometimes viewed as a means of reducing human bias and discretion, thereby producing more consistent outcomes. This perception contributes to the legitimacy of such systems, even where their operation is not fully understood.

At the same time, the adoption of these systems reflects broader shifts in governance practices. There is an increasing emphasis on data-driven policy, performance metrics, and measurable outcomes. Automated systems align with this orientation by providing quantifiable outputs that can be integrated into administrative frameworks.

However, these drivers do not fully account for the implications of automation for procedural structures. Efficiency, cost reduction, and consistency are legitimate objectives, but they do not address how decisions are experienced by those affected. The adoption of automated systems is often framed as a technical improvement, yet it involves changes to the way authority is exercised and how individuals interact with administrative processes.

### 3.4 Structural Implications

The integration of automated systems into governance processes produces structural changes that extend beyond individual applications. One such change is the standardisation of decision-making. By applying uniform criteria across cases, automated systems reduce variability in outcomes. While this can enhance consistency, it may also limit the capacity to account for exceptional circumstances.

A second implication is the reconfiguration of discretion. In traditional administrative settings, discretion is exercised by officials who interpret rules in relation to specific cases. In automated systems, discretion is embedded within system design—through the selection of variables, the definition of thresholds, and the calibration of models. Once implemented, these design choices shape outcomes in ways that are not always visible at the point of decision.

A third implication concerns the temporal dimension of decision-making. Automated systems operate continuously, processing inputs and generating outputs in real time or near real time. This can reduce delays, but it also compresses the space for reflection, review, and intervention. Decisions may be implemented before they can be meaningfully examined.

These structural changes are not inherently problematic, but they alter the conditions under which procedural fairness operates. The assumptions that underpin procedural safeguards—regarding reasoning, participation, and accountability—are challenged when decision-making is mediated through systems that function differently from human actors. Understanding these tensions requires a closer examination of how automated systems interact with the principles outlined in the previous section.

## 4. Where the Tension Emerges

### 4.1 From Reasoned Decisions to Statistical Outputs

A central premise of procedural fairness is that decisions are grounded in reasons that can be articulated and evaluated. The obligation to provide reasons is not merely a formal requirement; it serves multiple functions. It allows affected individuals to understand the basis of a decision, enables meaningful challenge, and provides a record against which the decision can be reviewed. In this sense, reasoning operates as the connective tissue between authority and accountability.

Automated systems complicate this premise by altering the nature of reasoning itself. Decisions influenced by such systems are often derived from statistical relationships within data rather than from a process of deliberation that can be reconstructed in narrative form. A model may identify correlations between variables and generate outputs based on patterns observed in historical data, but these outputs do not necessarily translate into reasons in the conventional sense. They indicate that a particular outcome is more or less likely, but they do not explain why that outcome should be preferred in a normative or contextual sense.

This distinction has practical consequences. When an administrative decision is challenged, the expectation is that the decision-maker can justify it by referring to relevant factors and demonstrating how those factors were weighed. In the case of automated systems, the justification may be limited to the fact that the system produced a particular output based on its internal parameters. This shifts the locus of explanation from the decision itself to the design and functioning of the system, which may not be readily accessible or intelligible to those affected.

The result is a gap between the form of reasoning expected by procedural frameworks and the form of reasoning produced by automated systems. This gap does not necessarily imply that automated decisions are arbitrary, but it does mean that their justification cannot always be expressed in terms that align with existing procedural norms.

**4.2 Opacity and the Limits of Explainability**

The difficulty of translating statistical outputs into reasons is compounded by the limited visibility of how automated systems operate. While not all systems are equally opaque, many involve processes that are not easily understood without specialised knowledge. Even where technical documentation is available, it may not provide a form of explanation that is meaningful to those subject to the system's decisions. It is important to distinguish between different forms of opacity. Some aspects of automated systems are opaque because they are technically complex; others are opaque because they are institutionally shielded, for example through claims of proprietary information or security concerns. In both cases, the effect is similar: the basis on which decisions are made is not fully accessible.

Efforts to address this issue have often focused on improving explainability. However, the notion of explainability itself is not straightforward. An explanation that satisfies a technical audience may not be sufficient for an individual seeking to understand why a particular decision was made in their case. Conversely, simplified explanations may obscure the factors that actually influenced the outcome.

This creates a tension between the demand for transparency and the practical limits of explanation. Procedural fairness does not require full disclosure of every aspect of a decision-making process, but it does require that individuals have enough information to understand and, where necessary, contest decisions that affect them. Where automated systems cannot provide such information in a usable form, the effectiveness of procedural safeguards is diminished.

**4.3 Diffusion of Responsibility**

Procedural fairness also relies on the ability to identify who is responsible for a decision. This identification is essential for both accountability and remedy. If a decision is flawed, there must be a clear locus at which responsibility can be assigned and from which corrective action can be sought.

In automated governance, responsibility is often distributed across multiple actors. System design may involve private vendors, data may be sourced from different agencies, and implementation may be carried out by administrative bodies that rely on system outputs. No single actor may have full control over the entire process, yet the outcome is experienced by the individual as a unified decision.

This diffusion complicates the application of accountability mechanisms. When a decision is challenged, it may not be clear whether the issue lies in the data, the model, the implementation, or the interpretation of the output. Each of these elements may fall under the responsibility of different actors, none of whom can be said to be solely responsible for the outcome.

From a procedural perspective, this fragmentation weakens the link between decision and responsibility. The ability to challenge a decision depends not only on understanding how it was made, but also on identifying who has the authority to review or alter it. Where responsibility is dispersed, the process of seeking redress becomes more complex and, in some cases, less effective.

**4.4 Standardisation and the Limits of Individual Assessment**

Another point of tension arises from the way automated systems standardise decision-making. By applying uniform criteria across large datasets, these systems promote consistency and reduce variability. However, this standardisation can come into conflict with the expectation that decisions should take into account the specific circumstances of individual cases. Procedural fairness does not require that every case be treated differently, but it does require that relevant differences be considered where they exist. Human decision-makers are able, at least in principle, to recognise when a case does not fit standard categories and to adjust their reasoning accordingly. Automated systems, by contrast, operate within predefined parameters that may not easily accommodate such deviations.

This limitation is not always apparent at the level of system design. Models may incorporate a wide range of variables and may be capable of capturing complex patterns. However, they still process cases as instances within a dataset, rather than as narratives with contextual meaning. As a result, factors that are difficult to quantify or that fall outside the scope of the data may not be adequately reflected in the outcome.

The consequence is a form of decision-making that is consistent but not necessarily responsive to individual circumstances. From the perspective of procedural fairness, this raises concerns about whether affected individuals are being treated as participants in the process or as data points within a system. The distinction is subtle, but it has implications for how fairness is understood and experienced.

**5. Case-Based Discussion**

**5.1 Aadhaar and Welfare Administration in India**

The Aadhaar system provides a useful point of reference for examining how automated processes intersect with procedural fairness in a governance context. As a digital identity infrastructure based on biometric and demographic data, Aadhaar has been integrated into multiple areas of public administration, including welfare delivery, banking, and service access. Its stated objectives include improving efficiency, reducing duplication, and ensuring that benefits reach intended recipients.

In welfare administration, Aadhaar is commonly used as a means of authentication. Individuals seeking access to subsidised goods or services may be required to verify their identity through biometric or one-time password-based mechanisms. This process is designed to confirm eligibility and prevent leakage within distribution systems. In practice, however, the functioning of authentication mechanisms introduces a layer of system dependency into what was previously an administrative interaction.

Instances of authentication failure have been documented in public reports and litigation. These failures may arise from mismatches in biometric data, connectivity issues, or errors in data integration across systems. When such failures occur, the consequence is often denial or delay of access to benefits. From an administrative perspective, the system is operating as designed—it is enforcing a verification requirement. From the perspective of the individual, however, the outcome is experienced as a denial without a clear avenue for immediate explanation or remedy.

The procedural implications of this dynamic are not straightforward. The requirement of authentication is not, in itself, incompatible with procedural fairness. However, when access to entitlements is contingent on system outputs that may fail for reasons not attributable to the individual, the basis for decision-making becomes less transparent. The individual is required to comply with a process whose functioning they cannot fully understand or influence.

Judicial engagement with Aadhaar, particularly in the context of welfare delivery, has acknowledged the need to ensure that technological systems do not result in exclusion. At the same time, the reliance on automated authentication mechanisms continues in practice, reflecting the administrative advantages they offer. The resulting situation is one in which procedural safeguards coexist with system-based decision points that are not always easily reconciled with those safeguards.

**5.2 Algorithmic Welfare Fraud Detection: The SyRI Case**

A different, though related, set of issues can be observed in the use of algorithmic systems for welfare fraud detection. The System Risk Indication (SyRI), implemented in the Netherlands, was designed to identify individuals at higher risk of committing welfare fraud by analysing data from multiple sources. The system combined information from various government databases to generate risk scores, which could then be used to trigger further investigation.

The operation of SyRI raised concerns about both transparency and fairness. Individuals were not informed that they were being assessed by the system,

nor were they provided with meaningful information about how risk scores were generated. This limited their ability to understand or challenge the basis on which they were identified as potential risks.

In 2020, a Dutch court held that the use of SyRI violated the right to privacy under the European Convention on Human Rights. While the judgment focused primarily on privacy considerations, it also highlighted broader concerns about the opacity of the system and the imbalance between the state's data-processing capabilities and the individual's ability to contest its outputs.

From the perspective of procedural fairness, the SyRI case illustrates the difficulties that arise when decision-making processes are not visible to those affected. The identification of individuals as "high risk" had practical consequences, yet the reasoning underlying this classification was not accessible. This undermined the possibility of meaningful participation in the process, as individuals could not engage with a decision whose basis was effectively concealed.

The case also demonstrates how automated systems can shift the locus of decision-making. Rather than making determinations directly, the system generated indicators that shaped subsequent administrative action. This indirect influence complicates the application of procedural safeguards, which are typically designed to operate at the point of decision rather than at earlier stages of data processing and risk assessment.

**5.3 Lessons from Practice**

Taken together, these examples illustrate that the challenges associated with automated governance do not arise solely from errors or malfunctions. Even when systems operate as intended, they can produce outcomes that are difficult to reconcile with established procedural expectations.

One key observation is that automation does not eliminate discretion; it redistributes it. Decisions about which variables to include, how to define thresholds, and how to integrate system outputs into administrative processes involve forms of judgment that are embedded within system design. These judgments are made prior to the point at which individuals encounter the system, and they are not always visible within the decision-making process itself.

A second observation is that the point at which procedural safeguards are applied may not correspond to the point at which decisions are effectively shaped. In traditional administrative processes, procedural fairness is often associated with the moment of decision. In automated systems, however, significant aspects of decision-making occur earlier, during data processing and model development. By the time an outcome is produced, the scope for meaningful engagement may already be limited.

Finally, these cases highlight the asymmetry between institutional and individual capacities. Administrative bodies have access to data, technical expertise, and system-level information that individuals do not. This asymmetry affects not only the substance of decisions, but also the ability of individuals to understand and respond to them. Procedural fairness, which relies on the possibility of participation and challenge, is strained under such conditions.

These observations do not suggest that automated systems are inherently incompatible with fair governance. Rather, they indicate that the application of procedural safeguards must take into account the structural characteristics of these systems. The challenge is not limited to correcting individual decisions, but extends to understanding how decisions are produced and where within that process fairness can be meaningfully ensured.

**6. Implications for Governance**

**6.1 Accountability Under System-Mediated Decision-Making**

The integration of automated systems into administrative processes complicates established models of accountability. In conventional

frameworks, accountability is tied to identifiable decision-makers who can be required to justify their actions and, where necessary, provide remedies. This structure presumes a relatively direct relationship between authority and outcome.

In system-mediated environments, this relationship becomes less direct. Decisions are often the product of multiple interacting components, including data sources, model architectures, and institutional practices. While the state remains formally responsible for administrative action, the pathway through which a particular outcome is produced may not be easily attributable to a single actor or decision point.

This does not eliminate accountability, but it alters its structure. Responsibility becomes distributed across different stages of the decision-making process, from system design to implementation. As a result, mechanisms that focus exclusively on the final decision may not capture the factors that most significantly influenced the outcome. Effective accountability in such contexts requires attention to how systems are constructed and integrated, rather than only how individual decisions are justified after the fact.

At the same time, there is a risk that the presence of automated systems may create a form of institutional distance. Administrative bodies may rely on system outputs as objective indicators, thereby reducing the scope for independent assessment. When this occurs, the ability to question decisions may be constrained not by formal barriers, but by the perceived authority of the system itself.

**6.2 Transformation of Citizen–State Interaction**

Procedural fairness is not only a set of legal requirements; it also shapes how individuals experience governance. Traditionally, administrative processes involve some degree of interaction between the individual and the state, whether through applications, hearings, or appeals. These interactions provide opportunities for explanation, clarification, and, where necessary, contestation.

Automated systems modify this interaction by introducing an interface between the individual and the administrative process. Decisions are increasingly communicated through digital platforms, often without direct engagement with a human decision-maker. While such interfaces can improve accessibility and efficiency, they also change the nature of participation.

In particular, the capacity to engage with a decision becomes dependent on the design of the system. If the system does not provide clear information about how an outcome was reached, or if it limits the avenues through which a decision can be challenged, the individual's role in the process is reduced. Participation becomes procedural in a narrow sense—limited to inputting data or following prescribed steps—rather than substantive.

This shift does not necessarily preclude meaningful interaction, but it does require that such interaction be intentionally designed. Without this, there is a risk that individuals experience governance as a series of outputs rather than as a process in which they can participate.

**6.3 Legitimacy and Acceptance**

The legitimacy of administrative decisions is closely linked to the perception that they are made through fair processes. Even where outcomes are unfavourable, individuals are more likely to accept decisions that they perceive as having been reached through procedures that are transparent, reasoned, and responsive.

Automated systems introduce a different basis for acceptance. Compliance may be driven less by an understanding of the decision and more by the absence of alternatives or the difficulty of contestation. Where systems are perceived as technical or objective, their outputs may be accepted without scrutiny, even if the underlying processes are not fully understood.

At the same time, the absence of visible reasoning can undermine trust. When individuals are unable to obtain explanations or to engage with the basis of a decision, the legitimacy of the process may be called into question. This is particularly significant in contexts where decisions affect access to essential services or rights.

The relationship between automation and legitimacy is therefore complex. Automated systems may enhance efficiency and consistency, but these benefits do not automatically translate into perceived fairness. Legitimacy depends not only on outcomes, but on the processes through which those outcomes are produced and communicated.

**6.4 Uneven Effects and Structural Asymmetry**

The implications of automated governance are not evenly distributed. Differences in access to digital infrastructure, levels of digital literacy, and familiarity with administrative processes shape how individuals interact with system-based decision-making. Those with greater resources and knowledge may be better positioned to navigate these systems, while others may encounter barriers that are not immediately visible.

This unevenness is compounded by asymmetries between institutions and individuals. Administrative bodies have access to technical expertise, system-level information, and the capacity to interpret outputs within a broader context. Individuals, by contrast, encounter the system at the point of outcome, often without insight into how that outcome was produced. Such asymmetries affect the practical operation of procedural safeguards. The ability to challenge a decision depends not only on formal rights, but also on the capacity to understand and engage with the process. Where this capacity is unevenly distributed, the effectiveness of procedural fairness is correspondingly uneven.

In addition, system design choices may have differential impacts across groups. Variables used in automated decision-making may correlate with socio-economic factors, leading to outcomes that disproportionately affect certain populations. While such effects are not always the result of deliberate discrimination, they raise questions about how fairness is assessed in contexts where decision-making is mediated through data.

**7. Rethinking Procedural Fairness in Automated Governance**

The preceding analysis indicates that the challenges posed by automated decision-making cannot be addressed solely through incremental adjustments to existing procedural frameworks. While current principles of procedural fairness remain relevant, their application presupposes conditions that are not fully present in system-mediated environments. The issue is not that these principles have become obsolete, but that they are being applied to forms of decision-making for which they were not originally designed.

A central limitation lies in the focus on the moment of decision. Procedural fairness is traditionally concerned with how a decision is reached in a specific instance: whether relevant factors were considered, whether the individual was heard, and whether reasons were provided. In automated systems, however, significant aspects of decision-making occur prior to this moment. The selection of data, the design of models, and the definition of thresholds shape outcomes in ways that are not visible at the point where a decision is communicated.

This temporal shift suggests that procedural fairness must be understood not only as a property of individual decisions, but also as a property of the systems through which those decisions are produced. Ensuring fairness in such contexts requires attention to upstream processes that influence outcomes before they take on a recognisable administrative form. This does not replace traditional safeguards, but it extends their scope.

A related limitation concerns the expectation of reasoned justification. As discussed earlier, automated systems generate outputs based on statistical

relationships rather than deliberative reasoning. While efforts to enhance explainability may improve transparency, they do not fully resolve the mismatch between statistical outputs and normative justification. A system may be able to indicate which variables influenced an outcome, but this does not necessarily provide a reason that can be meaningfully evaluated in relation to the individual case.

This suggests that the emphasis on explanation, while important, may need to be supplemented by other forms of accountability. Rather than focusing exclusively on the ability to explain individual decisions, attention may need to shift towards the conditions under which systems are designed and deployed. This includes questions about the appropriateness of the data used, the objectives encoded within the system, and the mechanisms available for oversight and correction.

Another point of reconsideration concerns the role of participation. Procedural fairness assumes that individuals have the opportunity to engage with the decision-making process, whether by presenting information, responding to allegations, or challenging outcomes. In automated systems, the scope for such engagement may be limited, particularly where decisions are produced rapidly and without prior notice.

Reframing participation in this context requires moving beyond the idea of direct interaction with a decision-maker. It may involve ensuring that individuals have access to meaningful avenues for review, that they are informed about how systems operate in general terms, and that there are mechanisms through which systemic issues can be raised and addressed. Participation thus becomes less about influencing a single decision and more about engaging with the structures that produce decisions over time.

The issue of responsibility also requires reconsideration. As noted earlier, automated systems distribute decision-making across multiple actors and stages. Traditional models of accountability, which focus on the identification of a single decision-maker, are not well suited to this structure. This does not mean that responsibility cannot be assigned, but it does suggest that it may need to be understood in a more layered manner.

One possible approach is to distinguish between different forms of responsibility corresponding to different stages of the process. For example, responsibility for system design may rest with those who develop and configure models, while responsibility for implementation may lie with administrative bodies that integrate these systems into decision-making processes. Such distinctions do not eliminate complexity, but they provide a framework for analysis. How accountability might be structured. Finally, the relationship between consistency and individualisation must be reassessed. Automated systems are often valued for their ability to produce consistent outcomes across large datasets. However, procedural fairness requires that relevant differences between cases be taken into account. Reconciling these objectives involves recognising that consistency and individualisation are not mutually exclusive, but must be balanced within system design.

This may require incorporating mechanisms that allow for deviation from standardised outputs in appropriate cases, as well as ensuring that individuals have the opportunity to present information that may not be captured within the system. Such measures do not undermine the efficiency of automated systems, but they acknowledge that fairness cannot be fully achieved through standardisation alone.

In sum, rethinking procedural fairness in the context of automated governance involves a shift in perspective. It requires moving from a focus on discrete decisions to an understanding of systems, from an emphasis on explanation alone to a broader conception of accountability, and from direct participation to more complex forms of engagement. These shifts do not resolve the tensions identified in

earlier sections, but they provide a basis for addressing them in a more systematic manner.

**8. Policy Recommendations**

The analysis in this article does not suggest that automated systems should be removed from governance processes. Their adoption reflects legitimate administrative objectives, including efficiency, scalability, and consistency. The policy challenge, therefore, is not to reverse automation, but to ensure that its integration does not erode the procedural safeguards that underpin fair governance.

First, greater emphasis should be placed on system-level auditability. Current regulatory approaches tend to focus on individual decisions, requiring that outcomes be justified after they occur. In automated environments, this approach is insufficient because many of the relevant factors are embedded within system design. Independent audit mechanisms—conducted by technically competent bodies—can provide a means of examining how systems operate in practice, including the data they rely on, the assumptions they encode, and the patterns they produce. Such audits need not require full disclosure of proprietary information, but they should be structured in a way that allows for meaningful external scrutiny.

Second, there is a need to strengthen documentation and traceability requirements. Administrative bodies that deploy automated systems should be required to maintain records of how these systems are configured, how decisions are generated, and how outputs are incorporated into administrative processes. This is not merely a technical exercise; it creates a basis for accountability by ensuring that decision pathways can be reconstructed when necessary. Without such documentation, procedural safeguards that rely on explanation and review are difficult to operationalise.

Third, human oversight should be recalibrated rather than assumed. In many policy discussions, the presence of a human "in the loop" is presented as a safeguard. In practice, however, the effectiveness of such oversight depends on the capacity and authority of the human actor involved. If officials are expected to rely on system outputs without the resources or expertise to question them, oversight becomes formal rather than substantive. Policies should therefore specify the conditions under which human intervention is required, particularly in cases where decisions have significant consequences or where system outputs are uncertain.

Fourth, regulatory frameworks should address the issue of meaningful access to contestation mechanisms. Individuals affected by automated decisions must have avenues to challenge outcomes that are accessible in practice, not merely in principle. This includes providing clear information about how decisions are made in general terms, as well as ensuring that review processes are not overly complex or resource-intensive. Where automated systems operate at scale, mechanisms for collective or systemic review may also be necessary to address patterns that cannot be identified through individual complaints alone.

Fifth, attention should be given to data governance as a component of procedural fairness. The quality and relevance of data used in automated systems have a direct impact on the fairness of outcomes. Policies should therefore include standards for data accuracy, periodic updating, and the identification of potential biases. This does not require eliminating all imperfections in data, which may not be feasible, but it does require acknowledging that data is not neutral and that its limitations have procedural implications.

Finally, policy approaches must remain context-sensitive, particularly in jurisdictions such as India where administrative systems operate under conditions of scale and diversity. Solutions that assume high levels of digital literacy or access may not be appropriate in all contexts. Procedural safeguards must therefore be designed with attention to how systems are experienced by different groups, including those who may face barriers in interacting with digital interfaces.

These recommendations are not exhaustive, nor do they offer a complete solution to the challenges identified in this article. They are intended to indicate directions in which policy can move to better align automated governance with the principles of procedural fairness. The objective is not to impose a single model of regulation, but to ensure that the integration of automated systems into governance processes is accompanied by mechanisms that preserve the ability of individuals to understand, engage with, and, where necessary, challenge the decisions that affect them.

**9. Conclusion**

The expansion of automated decision-making in governance does not merely introduce new tools into existing administrative processes; it alters the conditions under which those processes operate. This article has examined how this shift affects procedural fairness, a principle that has traditionally relied on assumptions about identifiable decision-makers, reasoned justification, and the possibility of meaningful participation.

The analysis has shown that automated systems do not negate these principles, but they complicate their application. Decisions influenced by such systems are often produced through processes that are statistical, distributed, and embedded within system design. As a result, the connection between decision, reasoning, and responsibility becomes less direct. Procedural safeguards that operate at the level of individual decisions are therefore limited in their ability to address the factors that shape outcomes upstream.

The case-based discussion illustrates that these challenges are not confined to hypothetical scenarios. In contexts such as digital identity verification and algorithmic risk assessment, individuals encounter outcomes that are mediated by systems whose operation is not fully visible or easily contestable. These experiences highlight a shift in how authority is exercised: from interactions with identifiable officials to engagement with structured processes that operate through data and code.

The implications of this shift extend beyond questions of technical design. They affect how accountability is structured, how individuals interact with administrative systems, and how legitimacy is constructed. Automated systems may enhance efficiency and consistency, but these attributes do not, on their own, satisfy the requirements of procedural fairness. Where individuals are unable to understand or engage with the basis of decisions, the procedural foundations of governance are placed under strain.

Responding to this challenge does not require abandoning automated systems, nor does it involve treating them as inherently incompatible with fair governance. Instead, it requires recognising that procedural fairness must be adapted to the conditions created by automation. This involves extending attention beyond the moment of decision to the systems that produce decisions, reconsidering how responsibility is distributed, and ensuring that mechanisms for review and contestation remain meaningful in practice.

The broader significance of this analysis lies in its implications for how governance is conceptualised. As decision-making becomes increasingly mediated by automated systems, the focus of inquiry must shift from individual acts of administration to the structures within which those acts occur. Procedural fairness, in this context, cannot be confined to discrete interactions; it must be understood as a feature of the systems that organise those interactions.

The task, therefore, is not to preserve existing procedural frameworks unchanged, but to refine them in a way that reflects the realities of automated governance. This requires both conceptual clarity and institutional adaptation. Without such efforts, there is a risk that procedural safeguards will persist in form while diminishing in substance.

*****

**References**

Aadhaar (Targeted Delivery of Financial and Other Subsidies, Benefits and Services) Act, 2016 (India).

Court of The Hague. (2020). *NJCM c.s./De Staat der Nederlanden (SyRI case)*, ECLI:NL:RBDHA:2020:865.

Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society, 20*(1), 14–29. https://doi.org/10.1080/1369118X.2016.1154087

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

United Nations Development Programme. (2021). *Digital public goods for development: A guidebook*. https://www.undp.org

Ministry of Electronics and Information Technology. (2021). *Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021* (India).

Supreme Court of India. (2018). *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2019) 1 SCC 1.

Supreme Court of India. (1978). *Maneka Gandhi v. Union of India*, (1978) 1 SCC 248.

Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance, 12*(4), 505–523. https://doi.org/10.1111/rego.12158

Zuboff, S. (2019). *The age of surveillance capitalism*. PublicAffairs.

European Commission. (2022). *Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. https://eur-lex.europa.eu