EXCLUSION BY DESIGN: AI, IDENTIFICATION FAILURES, AND THE MARGINS OF CITIZENSHIP IN INDIA

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

ABSTRACT

This paper explores how AI-mediated identification systems are reshaping the contours of citizenship in India by embedding forms of exclusion within infrastructures designed to secure inclusion. As digital identification becomes increasingly central to welfare distribution, governance, and access to public services, citizenship is no longer experienced solely as a stable legal status but as an ongoing process of verification, authentication, and recognition. In this context, the paper asks: what happens to citizenship when its enactment depends on successful interaction with technological systems?

Drawing on the expansion of biometric and data-driven identification within India's governance architecture, the analysis argues that these systems do not merely operationalise citizenship but actively reconstitute its boundaries by defining who is legible to the state. Focusing on points of breakdown—such as biometric mismatches, data inconsistencies, and infrastructural constraints—the paper demonstrates that exclusion is not simply an accidental failure of implementation but a structural feature of standardised, machine-readable governance. These exclusions are further intensified by forms of administrative and algorithmic opacity, where decision-making processes remain difficult to contest and responsibility is dispersed across technical and institutional actors.

The paper's central contribution lies in conceptualising these dynamics as **"**exclusion by design**,"** a condition in which the technical logics of efficiency, standardisation, and scalability produce new margins of citizenship that are neither fully visible nor easily addressed within existing legal frameworks. By situating the Indian case within broader debates on surveillance and governance, the paper also considers how such models of digitally mediated citizenship may circulate globally, raising critical questions about accountability, recognition, and the future of political membership in an increasingly data-driven system of governance.

Keywords: Algorithmic governance, Citizenship, Surveillance, Digital identification, Exclusion by design, Datafication, India, AI and governance,

Accountability, Global South

1. Introduction: Reconfiguring Citizenship Through Recognition

Citizenship is conventionally understood as a stable legal status, one that defines an individual's formal relationship with the state and secures access to a range of rights and entitlements. It is typically treated as a categorical distinction—one either possesses citizenship or does not—and is assumed to provide a relatively consistent basis for participation in political and social life. However, in contemporary governance systems, this understanding is increasingly under strain. Access to rights and services is no longer determined solely by legal status, but is mediated through processes of identification, verification, and recognition that are embedded within digital infrastructures.

In India, this shift is particularly visible in the expansion of large-scale digital identification systems that are now central to governance. These systems are widely associated with efforts to improve administrative efficiency, reduce leakages in welfare distribution, and extend access to public services. Their scale and ambition reflect a broader transformation in how states manage populations, moving toward data-driven forms of governance that rely on continuous processes of authentication. In this context, citizenship is not simply recorded; it is operationalised through technological systems that determine whether individuals can be recognised as eligible recipients of rights and services.

While these developments have generated considerable discussion in terms of their benefits, including efficiency gains and expanded inclusion, they also introduce a less visible but significant shift. Citizenship is increasingly enacted through successful interaction with technological systems. This creates a situation in which formal legal status may not be sufficient on its own.[1] Instead, access to rights becomes contingent on the ability to be accurately identified and verified within system-based processes. The implications of this shift are particularly important in contexts where identification systems are deeply integrated into everyday governance.

This paper examines how AI-mediated and automated identification systems contribute to this transformation by reshaping the boundaries of citizenship. In doing so, it adopts a broad understanding of AI, not limited to advanced machine learning models, but encompassing

[1] T.H. Marshall, Citizenship and Social Class, in *Citizenship and Social Class and Other Essays* 1 (1950).

automated decision-making processes that structure identification, verification, and access within governance systems. These systems operate through standardisation and data-driven recognition, enabling large-scale administration while also introducing new forms of dependency on technological processes.

The central argument of the paper is that such systems do not merely facilitate the administration of citizenship. Rather, they actively reconstitute its boundaries by embedding forms of exclusion within their design. Exclusion, in this sense, is not only the result of implementation failures or administrative inefficiencies. It is also a structural outcome of the way in which these systems are designed to function, particularly their reliance on standardised, machine-readable forms of identification and their limited capacity to accommodate variability and error.

To develop this argument, the paper focuses on the concept of "exclusion by design," which captures how the logics of efficiency, standardisation, and scalability inherent in digital and AI-enabled systems can produce conditions under which certain individuals become difficult to recognise within governance processes. These forms of exclusion are often not immediately visible, as they arise within technical and administrative systems that operate with limited transparency. As a result, they challenge conventional understandings of citizenship as a stable and universally accessible status.[2]

The analysis proceeds by examining the mechanisms through which exclusion emerges, including technical limitations in identification processes, infrastructural dependencies, and forms of administrative and algorithmic opacity. It then considers how these dynamics contribute to a broader shift in the nature of citizenship, from a primarily legal status to a more conditional and procedurally mediated form of recognition. While the primary focus is on India, the paper situates these developments within wider debates on surveillance, governance, and international relations, particularly in light of the increasing global circulation of digital governance models.

The paper does not argue against the use of digital or AI-enabled systems in governance. Instead, it seeks to provide a more nuanced understanding of their consequences, particularly the ways in which they reshape access to rights and redefine the practical meaning of

[2] Engin F. Isin & Bryan S. Turner, *Handbook of Citizenship Studies* (2002).

citizenship. By foregrounding the structural dimensions of exclusion, it aims to contribute to ongoing discussions on how technological systems can be designed and governed in ways that support, rather than inadvertently constrain, political and social inclusion.

2. Conceptual Framing: Citizenship, Surveillance, and Automated Governance

To understand how digital and AI-mediated identification systems reshape citizenship, it is necessary to clarify three interrelated concepts: citizenship as practice, surveillance as a mode of governance, and the role of data-driven and automated systems in mediating recognition. Rather than treating these as abstract theoretical categories, this section develops them as analytical tools for examining how access to rights is structured in contemporary governance systems.

Citizenship is often conceptualised as a formal legal status that defines membership within a political community. However, this formal definition does not fully capture how citizenship is experienced in practice. In many contexts, the exercise of rights depends on interactions with administrative systems that verify identity, eligibility, and entitlement. Citizenship, therefore, can be understood not only as a status but also as a set of practices through which individuals are recognised by the state. This distinction is important because it allows for the possibility that individuals who possess legal citizenship may nonetheless face barriers in accessing rights if they are not successfully recognised within these systems.

Surveillance plays a central role in shaping these practices of recognition. While surveillance is often associated with monitoring or control, it can also be understood more broadly as a mechanism through which states organise and manage populations. From this perspective, surveillance is not only about observing individuals, but about producing categories, assigning identities, and enabling administrative decision-making. Identification systems, particularly those that rely on digital infrastructures, form a key part of this process. They allow states to render populations legible in ways that can be acted upon, linking individuals to databases, records, and entitlements.

The increasing reliance on digital and automated systems intensifies this dynamic. Processes of identification and verification are now frequently mediated through technologies that convert personal attributes into standardised data formats. Biometric identifiers, database entries, and authentication protocols translate individuals into machine-readable forms that can

be processed at scale.[3] In this context, recognition becomes dependent on whether an individual can be successfully matched to their corresponding data within a system. This introduces a shift from recognition as an administrative judgment to recognition as a technical process.

AI and automated decision-making systems further extend this shift by structuring how recognition takes place. Even where advanced predictive models are not directly involved, automation shapes the conditions under which individuals are identified, verified, and granted access to services. These systems operate through predefined rules, probabilistic matching, and data-driven processes that can function without direct human intervention. As a result, decisions that affect access to rights may be made within systems that are not always transparent or easily contestable.

A key feature of these systems is their reliance on standardisation. In order to function at scale, identification systems require consistent and uniform data inputs. This enables efficient processing but also reduces the capacity to accommodate variation. Individuals whose data does not align with system requirements, whether due to technical error, bodily variation, or inconsistencies in records, may encounter difficulties in being recognised. This highlights an important tension between scalability and flexibility. Systems designed for large-scale governance often prioritise efficiency, but in doing so, they may produce rigid forms of recognition that do not account for the complexity of lived realities.

This tension is closely linked to the issue of opacity. As identification and decision-making processes become more automated, the pathways through which outcomes are produced can become less visible. Individuals interacting with these systems may not have access to information about how decisions are made or why certain outcomes occur. This lack of transparency can make it difficult to identify points of failure or to assign responsibility when errors occur. In this sense, opacity is not simply a technical feature, but a governance issue that affects accountability and access to redress.

Taken together, these dynamics suggest that citizenship in digitally mediated contexts is increasingly structured through processes of recognition that are embedded within surveillance-oriented and data-driven systems. Legal status remains important, but it is no longer sufficient on its own to guarantee access to rights. Instead, citizenship is enacted through

[3] Rogers Brubaker, Citizenship and Nationhood in France and Germany (1992).

successful interaction with systems that translate individuals into data and process that data according to predefined rules.

This conceptual framing provides the foundation for the analysis that follows. By understanding citizenship as a practice of recognition, surveillance as a mode of governance, and AI-enabled systems as mediators of access, it becomes possible to examine how exclusion emerges not only from failures in implementation, but from the structural features of the systems themselves. The next section builds on this framework to examine how these dynamics operate in the context of India's identification infrastructure, where digital systems play an increasingly central role in governance.

3. India's Identification Infrastructure as a Governance Model

The transformation of citizenship through digital and automated systems is particularly visible in India, where large-scale identification infrastructure has become central to governance. The Aadhaar system, administered by the Unique Identification Authority of India, represents one of the most expansive biometric identification projects globally. It assigns a unique identity number linked to biometric and demographic data, and has been progressively integrated into a wide range of governance functions, including welfare distribution, financial inclusion, and service delivery.

Aadhaar is often framed as a tool for improving efficiency and transparency in governance. By linking individuals to a verifiable identity, it is intended to reduce duplication, prevent fraud, and streamline administrative processes[4]. Its integration into welfare schemes such as the Public Distribution System (PDS), which provides subsidised food grains, has been a key example of this approach. The broader logic is one of targeted delivery, where benefits are linked to authenticated identities rather than distributed through more diffuse administrative channels.[5]

At the same time, the expansion of Aadhaar has not been without controversy. The legal and constitutional dimensions of its use were addressed in the landmark judgment of Justice K. S. Puttaswamy (Retd.) vs Union of India, where the Supreme Court of India upheld the

[4] Michel Foucault, *Discipline and Punish: The Birth of the Prison* (Alan Sheridan trans., Vintage Books 2d ed. 1995) (1975).
[5] James C. Scott, Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed (1998).

constitutional validity of Aadhaar while placing limits on its mandatory use, particularly in relation to privacy and proportionality. The judgment recognised both the potential benefits of the system and the risks associated with its expansive deployment, especially when linked to essential services.

Beyond the legal framework, a range of documented instances has highlighted the practical challenges associated with Aadhaar-based authentication in welfare delivery.[6] Reports by organisations such as Right to Food Campaign have pointed to cases where individuals were denied access to food rations due to biometric authentication failures. These failures can arise from a variety of factors, including worn fingerprints among manual labourers, connectivity issues in rural areas, and discrepancies between stored data and real-world conditions.

One widely cited example involves exclusion from the PDS in states such as Jharkhand, where Aadhaar-based authentication was linked to the denial of rations in certain cases. While the extent and causes of such exclusions remain contested, they have been sufficiently documented to raise concerns about the reliability of biometric verification in contexts where access to essential services is at stake. Importantly, these cases do not necessarily reflect systemic breakdowns across the entire system, but they do illustrate how reliance on authentication can create points of vulnerability.

These dynamics are further shaped by the infrastructural context in which Aadhaar operates. The effectiveness of biometric authentication depends not only on the accuracy of data but also on the availability of supporting infrastructure, including internet connectivity and functioning point-of-sale devices[7]. In areas where such infrastructure is uneven, authentication processes may fail even when individuals are correctly enrolled in the system. This highlights the extent to which access to rights can become dependent on the functioning of technological systems, rather than solely on legal entitlement.

In addition to welfare distribution, Aadhaar has also been linked to financial inclusion through initiatives such as the Jan Dhan–Aadhaar–Mobile (JAM) trinity, which aims to connect identity, banking, and mobile services. This integration reflects a broader shift toward digital governance, where multiple aspects of citizenship are mediated through interconnected

[6] Viktor Mayer-Schönberger & Kenneth Cukier, *Big Data: A Revolution That Will Transform How We Live, Work, and Think* (2013).
[7] David Lyon, *Surveillance Studies: An Overview* (2007).

systems. While such integration can enhance administrative efficiency and enable direct benefit transfers, it also increases the degree to which individuals' access to services is contingent on successful interaction with digital and automated systems.

From an analytical perspective, Aadhaar can be understood not only as a technological system but as a governance model. It embodies a particular approach to managing populations, one that prioritises standardised identification, data integration, and automated verification. This model enables large-scale administration but also introduces new dependencies on system performance and data accuracy. In doing so, it reshapes the conditions under which individuals are recognised and granted access to rights.

Crucially, the issues that arise in this context are not limited to isolated instances of failure. They reflect broader tensions inherent in the design of large-scale identification systems. The need for standardisation can conflict with the variability of lived realities, while the reliance on automation can reduce opportunities for discretionary intervention. These tensions do not necessarily invalidate the system as a whole, but they do highlight the importance of examining how its design shapes outcomes.

India's experience is significant not only because of the scale of its identification infrastructure, but also because of its influence beyond national boundaries[8]. Elements of India's digital governance model, particularly in relation to digital public infrastructure, have been promoted in international forums and adapted in other contexts. This makes it an important case for understanding how digitally mediated forms of citizenship may evolve and circulate more broadly.

The Indian case therefore provides a useful empirical grounding for the conceptual framework outlined earlier. It illustrates how citizenship, when mediated through digital and automated systems, becomes dependent on processes of recognition that are structured by technological design. The next section builds on this foundation by examining more closely how exclusion emerges within such systems, not only as a result of failure, but as a consequence of their underlying logic.

[8] Rob Kitchin, The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences (2014).

4. Exclusion by Design: Mechanisms of Structured Exclusion

The preceding sections have outlined how citizenship is increasingly mediated through digital and automated systems, particularly in the context of large-scale identification infrastructures. Building on this, the present section develops the paper's central analytical claim: that exclusion within such systems is not only the result of isolated failures, but is often embedded within their design. This is conceptualised here as **exclusion by design**, referring to the ways in which the technical, infrastructural, and administrative logics of these systems produce conditions under which certain individuals are more likely to be excluded from recognition.

This argument does not suggest that exclusion is intentional or that systems are designed with the purpose of denying access. Rather, it highlights how the core features that make these systems effective at scale—standardisation, automation, and efficiency—also create structural constraints that limit their ability to accommodate variability[9]. Exclusion, in this sense, emerges not as an anomaly, but as a byproduct of design choices that prioritise uniformity and scalability.

To unpack this, the section examines three interrelated mechanisms through which exclusion is produced: technical standardisation, infrastructural dependency, and administrative and algorithmic opacity.

Digital identification systems operate through the conversion of individuals into machine-readable forms. This involves the collection and storage of biometric and demographic data, which is then used for authentication through matching processes. In order to function effectively, these systems rely on standardised data inputs and consistent matching criteria.[10]

However, this reliance on standardisation introduces limitations. Human bodies and lived conditions do not always conform to the assumptions embedded within these systems. Biometric identifiers such as fingerprints and iris scans are treated as stable and reliable markers of identity, yet in practice they can be affected by age, occupation, environmental factors, and health conditions. Similarly, demographic data may contain inconsistencies due to errors in data entry, changes over time, or mismatches across records.

[9] Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (2015).
[10] Tarleton Gillespie, The Relevance of Algorithms, in *Media Technologies: Essays on Communication, Materiality, and Society* 167 (Tarleton Gillespie et al. eds., 2014).

When such discrepancies occur, the system may fail to authenticate an individual, even if that individual is correctly enrolled. In contexts where authentication is a prerequisite for accessing services, such failures can result in exclusion. Importantly, these outcomes do not necessarily indicate a malfunction in the system. Rather, they reflect the limits of a design that depends on precise and consistent data matching.

The key issue, therefore, is not simply the presence of error, but the way in which systems are structured to respond to it. In many cases, authentication systems operate on binary outcomes—either a match is successful, or it is not. This leaves limited room for accommodating uncertainty or variability, and can result in the denial of access even in cases where identity is not genuinely in doubt.

In addition to technical constraints, digital identification systems are embedded within broader infrastructural environments that shape how they function in practice. Authentication processes depend on a range of supporting elements, including reliable internet connectivity, functional hardware, updated databases, and accessible service points.

In contexts where infrastructure is uneven, these dependencies become significant. Even when individuals are correctly enrolled and their data is accurate, failures in connectivity or system availability can prevent successful authentication[11]. This is particularly relevant in rural or resource-constrained settings, where interruptions in network access or equipment malfunction may be more frequent.

The consequence is that access to rights becomes contingent not only on legal status or eligibility, but also on the functioning of technological systems at a given moment. Citizenship, in this sense, becomes conditional on infrastructure. This introduces a new layer of uncertainty, where individuals may be excluded not because they lack entitlement, but because the systems required to verify that entitlement are temporarily or persistently unavailable.

This form of exclusion is difficult to categorise within traditional administrative frameworks, as it does not arise from a formal decision to deny access. Instead, it emerges from the interaction between system design and infrastructural conditions, highlighting the extent to which governance outcomes are shaped by technological environments.

[11] Nick Seaver, Knowing Algorithms, 27 Media, Culture & Society 1 (2017).

A further mechanism through which exclusion is produced relates to the opacity of decision-making processes within digital and automated systems. As identification and verification become increasingly mediated through technology, the pathways through which outcomes are generated can become less visible to those affected.

In traditional administrative settings, decisions are often made by identifiable actors, and there are established procedures for review and appeal. In contrast, digital systems distribute decision-making across multiple layers, including data inputs, software processes, and institutional practices. When an authentication attempt fails or a service is denied, it may not be immediately clear whether the issue lies in the data, the technology, or the administrative process[12].

The incorporation of automated and AI-enabled processes further complicates this dynamic. These systems can operate without direct human intervention, applying predefined rules or probabilistic matching techniques to determine outcomes. While this can increase efficiency, it can also reduce transparency, particularly where the logic of decision-making is not easily accessible or understandable.

This opacity has significant implications for accountability. If individuals cannot identify how a decision was made, it becomes difficult to challenge or correct it. Responsibility may be diffused across technical and institutional actors, making it unclear who is accountable for addressing errors or providing redress. As a result, exclusion can occur in ways that are not only difficult to detect, but also difficult to resolve[13].

Taken together, these mechanisms illustrate how exclusion emerges not only from discrete failures, but from the structural features of digital and automated systems. Technical standardisation limits the capacity for flexible recognition, infrastructural dependency conditions access on system functionality, and administrative and algorithmic opacity complicates accountability.

The concept of exclusion by design captures this shift from viewing exclusion as an exception to understanding it as a patterned outcome. It draws attention to the ways in which design choices shape the distribution of access, even in the absence of intentional discrimination or

[12] Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (2018).
[13] Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (2018).

overt policy decisions.

This does not imply that such systems are inherently exclusionary or that their benefits are negligible. Rather, it suggests that their operation produces a dual effect: enabling inclusion for many while simultaneously creating new forms of exclusion for others. These exclusions are often less visible than traditional forms of administrative denial, but they can be equally consequential in terms of access to rights and services.

Understanding exclusion in this way provides a basis for rethinking how citizenship is enacted within digitally mediated governance systems. The next section builds on this analysis by examining how these dynamics contribute to a broader transformation in the nature of citizenship itself, shifting it from a stable legal status to a more conditional and procedurally mediated form of recognition.

5. Rethinking Citizenship: From Legal Status to Conditional Recognition

The preceding analysis has shown how exclusion can emerge from the structural features of digital and automated identification systems. This has implications that extend beyond questions of implementation or governance efficiency. It points to a more fundamental shift in how citizenship itself is experienced and enacted. While legal status remains formally intact, the practical exercise of citizenship is increasingly mediated through processes of recognition that are contingent, procedural, and technologically structured.

Traditionally, citizenship is understood as a binary and relatively stable category. An individual either possesses citizenship or does not, and this status is assumed to provide a secure foundation for accessing rights, services, and forms of political belonging. This understanding is rooted in legal frameworks that define citizenship in terms of membership within a political community. However, as access to rights becomes tied to system-based verification, this binary conception becomes less adequate for describing how citizenship operates in practice.

What emerges instead is a distinction between formal citizenship and recognised citizenship. Formal citizenship refers to the legal status conferred by the state, while recognised citizenship refers to the ability to successfully demonstrate that status within administrative and technological systems. In digitally mediated governance contexts, it is the latter that often determines whether individuals can access entitlements. This creates a situation in which legal

recognition alone may not be sufficient if it cannot be translated into system-based verification[14].

This shift introduces a more conditional dimension to citizenship. Access to rights becomes dependent on successful interaction with identification systems, which, as the previous section has shown, are subject to technical limitations, infrastructural constraints, and forms of opacity. Citizenship, therefore, is no longer only a status that one holds, but a condition that must be repeatedly validated through procedural processes. This does not formally alter the legal definition of citizenship, but it reshapes how it is experienced on a day-to-day basis.

The implications of this shift are particularly significant in contexts where digital systems are deeply integrated into essential services. When access to food distribution, financial transfers, or other welfare provisions depends on successful authentication, the consequences of recognition failure become more immediate and material. In such cases, the distinction between legal status and practical access becomes highly consequential, as individuals may find themselves unable to exercise rights despite formally possessing them.

This transformation can also be understood in terms of performativity. Citizenship becomes something that must be enacted through compliance with system requirements. Individuals are required to present themselves in forms that are legible to technological systems, whether through biometric authentication, data consistency, or adherence to procedural protocols. Those who are unable to meet these requirements, for reasons that may be technical, infrastructural, or situational, risk being positioned at the margins of recognition.

At the same time, this shift does not occur uniformly across all populations. The ability to interact successfully with digital systems is shaped by a range of factors, including access to infrastructure, familiarity with technology, and the stability of personal data. As a result, the move toward system-mediated recognition can reinforce existing inequalities, even where the intention is to promote inclusion. Individuals who are already vulnerable may face greater challenges in navigating these systems, increasing their risk of exclusion.

It is also important to note that this transformation does not necessarily involve a withdrawal of state authority. Rather, it reflects a reconfiguration of how that authority is exercised. Decision-making is increasingly mediated through technological systems, which structure the

[14] Reetika Khera, Dissent on Aadhaar: Big Data Meets Big Brother, 54 Econ. & Pol. Wkly. 32 (2019).

terms under which individuals are recognised and granted access. In this sense, the state's role is not diminished, but redistributed across institutional and technological frameworks that shape governance outcomes[15].

The concept of recognised citizenship therefore provides a way of understanding how legal status and technological mediation interact in contemporary governance. It highlights the gap that can emerge between formal entitlement and practical access, and draws attention to the conditions under which that gap is widened or narrowed. This perspective does not deny the continued importance of legal frameworks, but it emphasises that their effectiveness is increasingly dependent on the systems through which they are operationalised.

This analysis also has broader implications for how citizenship is conceptualised within political and legal theory. If citizenship is understood not only as a status but as a set of practices mediated through systems of recognition, then questions of access, accountability, and system design become central to its meaning. The boundaries of citizenship are no longer defined solely by legal criteria, but also by the technical and administrative processes through which individuals are rendered legible to the state.

In this context, the transformation of citizenship can be seen as part of a wider shift toward governance models that rely on data, automation, and continuous verification. While these models offer new possibilities for efficiency and inclusion, they also introduce new forms of dependency and vulnerability. Understanding these dynamics is essential for assessing how citizenship is being reshaped in practice, and for identifying the conditions under which it can be exercised effectively.[16]

The next section extends this analysis beyond the national context by considering how digitally mediated forms of citizenship may travel and influence governance practices in other settings. By situating the Indian case within broader debates in international relations, it becomes possible to examine how these transformations may contribute to emerging global patterns of political membership and governance.

Nice, this is where your paper steps out of the national frame and proves it belongs in an IR

[15] Usha Ramanathan, The Aadhaar Project: A Biometric Database of Citizens, 48 Econ. & Pol. Wkly. 77 (2014).
[16] Jean Drèze et al., Aadhaar and Food Security in Jharkhand: Pain Without Gain?, 52 Econ. & Pol. Wkly. 50 (2017).

conversation.

6. International Relevance: AI, Governance, and the Global Circulation of Citizenship Models

While the preceding analysis has focused on India, the dynamics it highlights are not confined to a single national context. The increasing reliance on digital identification, automated verification, and AI-enabled governance reflects a broader transformation in how states manage populations and deliver services. As these systems expand and circulate across borders, they carry with them particular assumptions about identification, recognition, and access, with implications for how citizenship is structured in different contexts.

India occupies a significant position within this landscape, not only because of the scale of its identification infrastructure, but also because of its role in promoting digital public infrastructure as a model for governance. Initiatives associated with identity systems, digital payments, and data integration have been presented as scalable solutions for improving service delivery, particularly in the Global South. As elements of these systems are studied, adapted, and in some cases exported, they contribute to the diffusion of governance models that rely on standardised, data-driven forms of identification.

This process can be understood as a form of policy and technological diffusion within international relations. States often look to existing models when developing their own governance systems, particularly when those models are associated with efficiency gains or successful implementation at scale. In this context, digital identification systems are not merely technical tools, but part of a broader repertoire of governance practices that can travel across borders. The adoption of such systems is shaped by local conditions, but the underlying logics of standardisation, automation, and verification tend to persist.

The role of AI and automated systems in this process is particularly significant. As governance becomes increasingly mediated through data and automation, the criteria for recognition are embedded within technological systems that may not be easily adaptable to diverse contexts. This can lead to the standardisation of how individuals are identified and verified, potentially reducing the flexibility of governance systems to accommodate local variations. In effect, AI-enabled systems can contribute to the emergence of shared frameworks for operationalising

citizenship, even in the absence of formal international coordination[17].

This raises important questions about how citizenship is evolving within the international system. Traditionally, citizenship has been closely tied to the sovereignty of the state, with each state defining its own criteria for membership and rights. However, as digital governance systems become more interconnected and widely adopted, there is a possibility that the practical experience of citizenship may begin to converge across different contexts. This convergence does not necessarily imply uniformity in legal definitions, but it may shape how citizenship is enacted through similar processes of identification and verification.

At the same time, the global diffusion of these systems may also reproduce or reinforce existing inequalities. States differ significantly in terms of technological capacity, infrastructure, and regulatory frameworks. As a result, the implementation of digital identification and AI-enabled governance may produce uneven outcomes, both within and between countries. Individuals in contexts with limited infrastructure or weaker accountability mechanisms may face greater risks of exclusion, while those in more resourced settings may benefit from more reliable systems.

These dynamics can also be situated within broader debates on digital sovereignty and governance. States are increasingly seeking to assert control over data, infrastructure, and technological systems as part of their governance strategies. Digital identification systems play a key role in this process, as they link individuals to state-managed databases and enable the administration of services. At the same time, the reliance on shared technological standards and external expertise can create dependencies that complicate claims of sovereignty. The governance of citizenship, therefore, becomes entangled with questions of technological control and international influence.

From an international relations perspective, this suggests a shift in how political membership is organised and experienced. Citizenship remains formally defined within national boundaries, but its practical operation is increasingly shaped by transnational flows of technology, policy models, and governance practices. The spread of AI-enabled identification systems contributes to this shift by embedding particular forms of recognition within the infrastructures that

[17] Susan Strange, *The Retreat of the State: The Diffusion of Power in the World Economy* (1996).

mediate access to rights[18].

Importantly, this does not imply that all states will adopt identical systems or that local variations will disappear. Rather, it points to the emergence of a shared set of challenges associated with digitally mediated governance, including issues of exclusion, accountability, and transparency. As more states integrate AI and automated systems into their governance frameworks, these challenges are likely to become more prominent within international discussions on development, governance, and rights[19].

The Indian case, in this context, serves as an important reference point. It illustrates both the potential and the limitations of large-scale digital identification systems, and highlights the need to critically examine how such systems shape access to citizenship in practice. By situating these developments within a broader international framework, it becomes possible to see how the dynamics of exclusion by design may extend beyond individual contexts and contribute to emerging patterns in the governance of political membership.

The final section of the paper turns to the question of how these challenges might be addressed. Rather than proposing comprehensive solutions, it outlines a set of policy considerations that aim to mitigate the risks associated with digitally mediated and AI-enabled systems, while preserving their potential benefits.

7. Policy Considerations: Designing for Recognition and Accountability

The analysis presented in this paper does not suggest that digital or AI-enabled identification systems should be abandoned. On the contrary, such systems can play a significant role in improving administrative efficiency, expanding access to services, and enabling large-scale governance. However, the preceding discussion highlights the need to reconsider how these systems are designed and governed, particularly in light of the forms of exclusion that may arise from their structural features.

The policy challenge, therefore, is not simply one of adoption, but of design and implementation. It requires attention to how technological systems mediate access to rights, and how their limitations can be addressed without undermining their benefits. The following

[18] Martha Finnemore & Kathryn Sikkink, International Norm Dynamics and Political Change, 52 Int'l Org. 887 (1998).
[19] Shoshana Zuboff, *The Age of Surveillance Capitalism* (2019).

considerations outline key areas where intervention may be necessary.

7.1 Designing for Failure Rather than Perfection

A central issue identified in this paper is the assumption of consistent and accurate recognition within identification systems. In practice, however, technical errors, data inconsistencies, and infrastructural disruptions are unavoidable. Systems that operate on the expectation of perfect authentication risk translating these errors directly into exclusion.

Designing for failure involves recognising that errors will occur and building mechanisms to accommodate them. This may include allowing multiple attempts at authentication, enabling alternative modes of verification, and ensuring that temporary failures do not result in immediate denial of access. In systems that incorporate AI or automated decision-making, this also requires safeguards to prevent automated outcomes from being treated as final without the possibility of review.

The aim is not to eliminate error entirely, which is unrealistic, but to ensure that errors do not have disproportionate consequences for individuals' access to essential services.

7.2 Strengthening Accountability in Automated Systems

As identification and verification processes become increasingly mediated through automated and AI-enabled systems, questions of accountability become more complex. When exclusion occurs, responsibility may be distributed across multiple actors, including system designers, implementing agencies, and administrative authorities.

Addressing this requires clearer lines of responsibility and more accessible mechanisms for grievance redress. Individuals should be able to identify how a decision affecting them was made, and who is responsible for addressing errors. This is particularly important in cases where automated systems play a role in decision-making, as the absence of human oversight can make it more difficult to challenge outcomes.

Strengthening accountability does not necessarily require full technical transparency, but it does require that the processes through which decisions are made are sufficiently clear to allow for meaningful review and correction.

7.3 Maintaining Non-Digital and Human Alternatives

A further implication of the analysis is the need to avoid making access to rights entirely dependent on digital or AI-mediated verification. While digital systems can enhance efficiency, their limitations mean that they should not function as the sole gateway to essential services.

Maintaining non-digital alternatives and human oversight mechanisms is critical in ensuring that individuals are not excluded due to system failures or limitations. This may involve providing manual verification options, allowing for administrative discretion in cases of authentication failure, and ensuring that individuals have access to in-person support when needed.

Such measures recognise that technological systems, while powerful, cannot fully capture the complexity of individual circumstances, and that human judgment continues to play an important role in equitable governance.

7.4 Enhancing Transparency and System Legibility

Transparency is a key requirement for ensuring that digital and automated systems operate in a manner that is both accountable and understandable. Individuals interacting with these systems should have access to clear information about how identification and verification processes work, and what factors may affect outcomes.

In the context of AI-enabled systems, this includes making the role of automation visible, even if the underlying technical details remain complex. Transparency, in this sense, is less about full disclosure of algorithms and more about ensuring that the logic of decision-making is accessible to those affected.

Greater transparency also supports broader public oversight, enabling policymakers, researchers, and civil society organisations to assess how these systems function and where improvements may be needed.

7.5 Balancing Efficiency with Inclusion

Finally, there is a need to recognise the trade-offs inherent in the design of large-scale governance systems. The emphasis on efficiency, scalability, and standardisation has been a

key driver of digital identification initiatives. However, as this paper has shown, these same features can contribute to exclusion if not carefully managed. Policy approaches must therefore seek to balance efficiency with inclusion, ensuring that efforts to streamline governance do not come at the cost of equitable access. This may require accepting a degree of administrative complexity in order to accommodate variation and provide safeguards against exclusion.

Taken together, these considerations suggest that the effectiveness of digital and AI-enabled governance systems should not be evaluated solely in terms of efficiency or scale, but also in terms of their ability to support inclusive and accountable access to rights. Designing systems with these principles in mind is essential for ensuring that technological innovation enhances, rather than constrains, the practical experience of citizenship.

8. Conclusion: Citizenship in the Age of System-Mediated Recognition

This paper has examined how digital and AI-mediated identification systems are reshaping the contours of citizenship, with a particular focus on the Indian context. It has argued that these systems do not simply facilitate the administration of citizenship, but actively reconfigure its practical boundaries by mediating how individuals are recognised within governance processes. The analysis has shown that exclusion within such systems cannot be understood solely as the result of implementation failures or administrative inefficiencies. Instead, it often emerges from the structural features of system design, including reliance on standardised data, dependence on technological infrastructure, and the opacity of automated decision-making processes. These features, while enabling large-scale and efficient governance, also introduce constraints that limit the capacity of systems to accommodate variability and error. The concept of exclusion by design has been developed to capture this dynamic. It shifts the focus from isolated instances of failure to the patterned ways in which system architecture shapes access to recognition and, by extension, access to rights. This perspective highlights the need to consider not only whether systems function, but how their underlying logics distribute inclusion and exclusion across different populations. Building on this, the paper has proposed a shift in how citizenship is understood in practice. While legal status remains central, it is increasingly supplemented by what has been described as recognised citizenship—the ability to successfully demonstrate one's status within system-mediated processes of verification. This distinction draws attention to the gap that can emerge between formal entitlement and practical access, particularly in contexts where digital systems are deeply integrated into governance.

Although the empirical focus has been on India, the implications of this analysis extend beyond a single national setting. As digital public infrastructure and AI-enabled governance models are adopted and adapted across different contexts, the mechanisms through which citizenship is operationalised may begin to exhibit similar patterns. This raises broader questions for international relations, particularly regarding the diffusion of governance models and the evolving relationship between technological systems and political membership. At the same time, it is important to emphasise that the expansion of digital and AI-enabled systems does not inevitably lead to exclusionary outcomes. These systems have the potential to enhance efficiency, improve service delivery, and expand access in significant ways. The challenge lies in how they are designed and governed. As the policy considerations outlined in this paper suggest, attention to issues of flexibility, accountability, transparency, and alternative access pathways is essential in mitigating the risks associated with system-mediated governance. Ultimately, the transformation examined in this paper reflects a broader shift in the nature of governance itself. As states increasingly rely on data, automation, and continuous verification, the conditions under which individuals are recognised and granted access to rights are being redefined. Citizenship, in this context, is no longer solely a legal status secured by formal recognition, but a condition that is enacted through interaction with technological systems. Understanding this shift is critical, not only for analysing contemporary governance practices, but also for shaping future approaches to the design and regulation of digital and AI-enabled systems. If citizenship is increasingly mediated through processes of recognition, then the design of those processes becomes central to how inclusion, exclusion, and belonging are structured in practice.