

Disintermediation by Function: Digital Public Infrastructure and the Selective Persistence of Welfare Mediation in India.

Abstract:

India's Digital Public Infrastructure — Aadhaar, Direct Benefit Transfer, and the payment rails built around them — promised to move welfare directly from the state to the poor, removing the officials and agents who stood between. The evidence is contradictory: some studies find that digital systems cut leakage and improved delivery; others find that intermediaries endured and new exclusions opened. This article treats the disagreement as a problem of aggregation. Welfare delivery bundles distinct operations — establishing eligibility, verifying identity, disbursing benefits, cashing them out, and resolving failures — that respond to digitalisation differently. Using MGNREGA wage payments and Aadhaar authentication in the Public Distribution System, it develops a functional account built on two properties: a function's *administrative computability*, which governs whether it can move from human hands to a system, and the *design of its failure-handling*, which shapes what automation does once it breaks down. The same authentication technology improved delivery in one state yet cut 1.5 to 2 million people off from their rations in another — a divergence explained by failure design, not technology. Digitalisation, then, does not abolish mediation but redistributes it: across functions, and toward those who design the rules.

Keywords: Digital public infrastructure; digital governance; administrative discretion; welfare delivery; disintermediation; administrative burden; algorithmic governance; Aadhaar; MGNREGA; India

1. Introduction

The case for routing India's welfare payments through digital systems was made in the language of removal. When the central government expanded Direct Benefit Transfer (DBT) from 2013, the promise was that money would move from the treasury to the beneficiary without passing through the officials whose handling of it had produced delay, diversion, and arbitrary gatekeeping (Ministry of Finance, 2013). Aadhaar carried a related promise. A unique biometric identifier would let the state recognise a resident directly, dispensing with the intermediary vouching that had governed access to public goods (Aiyar, 2017; Scott, 1998). Built on these foundations, and packaged with the payment interfaces layered on top of them, the resulting Digital Public Infrastructure (DPI) was described as substituting an unmediated relationship between state and citizen for a mediated one.

By *Digital Public Infrastructure* this article means the specific Indian assemblage sometimes branded the "India Stack" and popularised through the "JAM trinity" of Jan Dhan bank accounts, Aadhaar identity, and mobile connectivity: a set of interoperable, population-scale digital rails for identity, payments, and data on which welfare and other services are built (International Monetary Fund, 2023). The term is used here in this concrete, architectural sense rather than as a generic category of e-government, and the analysis concerns only those elements of the stack that touch welfare delivery — Aadhaar authentication, Aadhaar-seeded bank accounts, and the DBT payment rails — not the whole of the ecosystem.

Whether this infrastructure has delivered directness is contested in a way that is initially hard to reconcile. A large experimental literature reports that biometric and electronic payment systems reduced the leakage of funds and reached beneficiaries more reliably than the arrangements they replaced (Muralidharan et al., 2016). Official assessments go further, attributing cumulative fiscal savings on the order of ₹3.5 lakh crore over 2014–2024 to the deletion of duplicate and "ghost" beneficiaries enabled by Aadhaar seeding (Government of India, 2025). A large body of survey and ethnographic work reports the opposite: that local agents remained indispensable, that authentication failed for the people least able to withstand it, and that points of exclusion opened where old ones were meant to close (Drèze et al., 2017; Chaudhuri, 2019; Masiero & Das, 2019). This same literature disputes the savings figures, arguing that headline numbers conflate genuine leakage reduction with the exclusion of eligible beneficiaries and with unrelated price effects, and noting that a national audit found the widely cited savings on liquefied petroleum gas subsidies substantially overstated (Drèze & Khera, 2018). Both literatures are careful, and both describe the same infrastructure. That two well-executed research programmes should reach opposing verdicts is worth pausing over, because the difficulty may lie less in the evidence than in the unit to which it is applied.

Welfare delivery is not a single act that is either mediated or not. It is a sequence of operations that happen to be bundled together: establishing who qualifies, confirming that a claimant is the person entitled, computing and moving the benefit, converting that benefit into cash or grain at

the point of collection, and dealing with the failures that arise along the way. These operations differ in kind. Some turn on interpreting ambiguous evidence; others amount to applying a fixed rule to a verifiable fact. Because they differ in this respect, a digital system built to serve all of them touches each differently, displacing an intermediary here, leaving one in place there, and generating a new one somewhere else. Accounts that speak of disintermediation at the level of the programme rather than the operation flatten this variation, and end up describing different parts of the same process while appearing to contradict one another.

This article therefore shifts the unit of analysis from the programme or the technology to the administrative function. Where existing scholarship has largely asked whether digitalisation, taken as a whole, reduces leakage, improves access, or reproduces exclusion, this article asks a more disaggregated question: which operations within welfare delivery does digitalisation displace, which does it leave intact, and what governs the difference? Two Indian cases are used to develop and probe an answer. Wage payment under MGNREGA layered electronic transfer onto a scheme whose allocation and verification stayed in official hands. Aadhaar-based authentication in the PDS introduced a biometric check at the point where households collect subsidised food. In each, one infrastructure meets several operations, and its effect is uneven. It removes mediation from some, leaves it untouched in others, and, in a third, converts a visible and negotiable human intermediary into an opaque technical one. The cases also show that whether a computable operation is disintermediated to a beneficiary's advantage or turned into a mechanism of exclusion is not settled by the technology; it depends on how the system is built to behave when it fails.

The article makes three contributions, framed as analytical propositions rather than descriptions. First, it proposes the administrative function as the appropriate unit of analysis for studying welfare disintermediation, and shows that disaggregating delivery into distinct operations resolves much of the apparent conflict in the literature. Second, it distinguishes two properties that jointly condition the effect of digitalising any given operation: its *administrative computability*, which conditions whether the operation can move from human hands to the system level, and the *design of its failure-handling*, which conditions what happens once it does. Third, it uses this distinction to explain a fact that neither the disintermediation nor the persistence reading accounts for on its own — that the same authentication technology improved delivery under one implementation and excluded roughly two million people under another. These claims are set out formally as four propositions in Section 3 and assessed against the two cases in Sections 5 through 8.

The argument deliberately limits the matter. It does not hold that eligibility, or any function, is permanently beyond automation; that is an open empirical question the present evidence cannot settle, and the argument turns on how functions behave in the cases examined, not on claims about the ultimate limits of computation. It does not extend to the electoral or representational effects of digital welfare, for which the required evidence does not exist. What the two cases support is a more exact account of a familiar phenomenon: an explanation of why the

disintermediation reading and the persistence reading are both, within their domains, correct, and an identification of the conditions — a function's dependence on judgment, and the design of its failure-handling — that help decide which reading applies where.

The remainder of the article proceeds as follows. Section 2 sets out the two readings of India's welfare digitalisation and shows that the divide cannot be reduced to a contrast of methods. Section 3 develops the analytical framework — mediation and its neighbouring concepts, the migration of discretion, administrative computability, and the design of failure — and states four propositions. Section 4 describes the research design and the basis on which evidence was selected and organised. Section 5 reconstructs the pre-digital architecture of the two programmes. Sections 6 and 7 examine, respectively, the digitisation of the MGNREGA payment rail and Aadhaar authentication in the PDS. Section 8 compares the two cases function by function, presents the evidence systematically, and assesses the propositions. Section 9 addresses objections and states the limits of the claim. Section 10 concludes.

2. Two readings of a single infrastructure

The scholarship on India's welfare digitalisation has settled into two broad readings, and it helps to state what each establishes before asking why they diverge.

The first reading emphasises measurable gains. Its strongest statement is a randomised evaluation of biometric "Smartcards" in undivided Andhra Pradesh, which found that the reform delivered wage payments under the National Rural Employment Guarantee Scheme faster and with less corruption, cutting leakage in the scheme by roughly two-fifths without reducing access, at a cost the beneficiaries' own time savings repaid (Muralidharan et al., 2016). On this account, the intermediaries whose discretion had enabled diversion were disciplined by a payment system that recorded who was paid and made it harder to skim. Official evaluation extends the claim across programmes: the government reports that Aadhaar-linked authentication has removed large numbers of duplicate ration cards and ineligible beneficiaries and generated substantial fiscal savings (Government of India, 2025). The experimental evidence is real, and any honest treatment has to hold it in view.

The second reading emphasises endurance and exclusion. Household surveys of Aadhaar-based authentication in the PDS found that mandatory biometric verification excluded eligible cardholders and raised the costs of collection, while leaving untouched the quantity fraud that constituted much of the actual diversion (Drèze et al., 2017). Ethnographic and information-systems research reached compatible conclusions from a different angle: that Aadhaar does not dispense with human intermediation so much as depend on and reproduce it (Chaudhuri, 2019), and that datafying an anti-poverty scheme introduces its own legal, informational, and design-related injustices (Masiero & Das, 2019). This literature also contests the official savings figures directly, arguing that they have no solid basis and conflate distinct effects, a critique lent weight by an audit that found the flagship savings claim overstated by an order of magnitude (Drèze

&Khera, 2018). Here the infrastructure appears less as a solvent of mediation than as a new terrain on which mediation and exclusion are reorganised.

The obvious explanation for the divide — that it reflects methods, with experiments finding gains and fieldwork finding harms — does not survive scrutiny, because the divide runs through the experimental literature as well. The team responsible for the Andhra Pradesh results later evaluated Aadhaar-based authentication in the Jharkhand PDS and found that corruption fell there too, but that between 1.5 and 2 million legitimate beneficiaries lost access to their entitlements at some point during the reform (Muralidharan et al., 2025). The same researchers, using the same broad method, documented improvement in one setting and large-scale exclusion in another. That pattern is the point of entry for this article. If comparable technologies, competently studied, produce opposite outcomes, the explanation must lie in something other than the technology and other than the quality of the research — in what the systems were asked to do, and in how they were built to behave.

3. Mediation, computability, and the design of failure: an analytical framework

To ask whether digitalisation removes intermediaries, one needs a working definition of what an intermediary is, some discipline about neighbouring terms, and a criterion for deciding when mediation has genuinely been displaced rather than merely relocated. This section supplies these, situates the argument within the broader literature on digital and algorithmic governance, develops the two properties on which the analysis turns, and states four propositions.

3.1 Mediation, its neighbours, and a criterion for displacement

Mediation here means the institutional layer that stands between a citizen's claim and its satisfaction: the actors, offices, and procedures through which a person obtains a service and through which the state acts on its policies. The definition is deliberately neutral. An intermediary is not by definition an obstacle; whether mediation helps or harms depends on what it does in a given case. This matters in the Indian setting, where the record cuts both ways. Gupta's (2012) ethnography shows how routine dealings with an indifferent or predatory bureaucracy operate as a form of structural violence, so that the intermediary is often the barrier. Yet mediation is also how the poor reach the state at all (Corbridge et al., 2005), and removing it can strip away a channel of redress along with a source of extraction.

Because the article's central claim is that mediation persists, it must specify what would count as persistence and what as displacement; otherwise the claim risks being unfalsifiable, since some actor or system always stands between claim and satisfaction. Three forms of mediation are therefore distinguished. *Human intermediation* exists where a specific person stands between citizen and benefit and exercises judgment or control over access — the dealer, the panchayat official, the banking agent. *Institutional mediation* exists where an office or procedure, rather than a nameable person, governs the passage of a claim. *Technical mediation* exists where a

device, database, or authentication system performs a function that a human once performed. A fourth layer, *brokerage* — the informal intercession of party workers, local leaders, and fixers through whom the poor often navigate the state (Auerbach & Thachil, 2018) — is distinct again, and largely untouched by the systems examined here. With these distinctions, the article's criterion can be stated precisely: an operation is *disintermediated* when human intermediation over it is removed and not reconstituted as an equivalent human choke-point elsewhere; it exhibits *persistence* when human discretion over the operation survives in a new location or a new form; and it undergoes *displacement into technical mediation* when a human function is performed by a system but the capacity to exclude, delay, or err is retained by that system. On this criterion, the replacement of a dealer's identity check by a biometric device is a genuine displacement of human intermediation, whereas the movement of eligibility judgment from a counter to a set of database rules is a relocation, not a removal.

Several further distinctions keep the analysis from sliding. Mediation is not the same as *discretion*: mediation names the channel, discretion names the judgment sometimes exercised within it, and a system can close the channel while leaving the judgment intact, or automate the judgment while the channel remains. Mediation is not *representation* in the political-theory sense of acting for others through authorised institutions; the operations examined here are administrative, and no claim is made about representation. And mediation is broader than *street-level bureaucracy*, which refers to the specific frontline officials whom Lipsky (1980) placed at the centre of implementation. These are not pedantic separations. The article's claim concerns administrative mediation and its discretionary content; it is silent on representation and brokerage except where the evidence reaches them, which is rarely.

3.2 The migration of discretion

Lipsky's (1980) account explains why the discretionary element of mediation resists removal. Officials at the point of service cannot avoid judgment, because the rules they administer are incomplete and cannot anticipate every configuration of informal income, irregular household, or borderline asset a real applicant presents. Maynard-Moody and Musheno (2003) show how that judgment actually works: frontline workers act less as state agents mechanically applying rules than as citizen agents who form normative judgments about the deservingness of the person before them and respond accordingly. Discretion, in this tradition, is the medium through which general rules are fitted to particular cases, which is also why legislating it away is difficult.

What happens to that discretion when decisions are encoded in software is the subject of a now-substantial literature (Busch & Henriksen, 2018). Its foundational statement is Bovens and Zouridis's (2002) account of the movement from street-level to system-level bureaucracy. When an agency's decisions are written into code, they argue, discretion does not vanish; it moves upstream, from the official at the counter to the analysts and designers who fix what the system will and will not recognise. Implementation of the law is in one sense perfected, and in another sense the judgment that once governed hard cases becomes harder to see and harder to contest.

This insight has been generalised by two adjacent bodies of work. Scholarship on digital-era governance treats the reintegration of administrative functions around digital systems as the defining wave of public-sector reform succeeding new public management, reorganising the state around information infrastructures rather than around discrete agencies (Dunleavy et al., 2006). Scholarship on algorithmic governance examines how computational procedures of ranking, scoring, matching, and classification become instruments of social ordering, redistributing authority among officials, systems, and the private vendors who build them (Katzenbach & Ulbricht, 2019; Yeung, 2018; Andrews, 2019). A recurrent finding across this literature is that automation can erode the procedural safeguards that once attached to discretionary decisions, because a coded determination offers no obvious equivalent of the hearing, the explanation, or the appeal (Citron, 2008).

A second, more critical strand concerns what this migration does to the poor specifically. Peeters and Widlak (2018) give the argument a sharp edge with the notion of a “digital cage”: an information architecture can exclude a citizen whose circumstances the system was not built to accommodate, and can do so without anyone deciding to exclude them, because the exclusion is a property of the architecture rather than of a person’s choice. At the level of the welfare state as a whole, Eubanks (2018) documents how automated eligibility systems, risk scores, and data-matching regimes in the United States tend to profile and discipline the poor rather than serve them, and the United Nations Special Rapporteur on extreme poverty warned that a “digital welfare state” risks turning social protection into an apparatus of surveillance and control (Alston, 2019). These accounts are not confined to low-capacity settings: high-capacity administrations have produced their own architectural exclusions, from the Dutch civil registry’s rigidities (Peeters & Widlak, 2018) to automated fraud-detection regimes that wrongly penalised large numbers of legitimate claimants. The present article does not discover the migration of discretion or the architecture of exclusion; it takes both as established. What the Indian cases add is that the migration is *selective* — complete for some functions, stalled for others — and that whether it excludes depends on how the architecture treats the cases it cannot fit.

3.3 Administrative computability

From these strands the article develops its first organising property. A function is amenable to disintermediation to the degree that its output can be generated from specified and verifiable inputs without case-specific interpretation. It is important to be precise about the sense of “computable” at work here, because a purely technical reading would invite the objection that, with sufficiently rich data, almost anything can be automated. The concept used here is *administrative computability*: the extent to which a programme function can produce an administratively valid decision from specified and verifiable inputs without contextual judgment, under the data, rules, and institutional arrangements that actually obtain. This is a property of an operation within a given administrative order, not a claim about the ultimate reach of computation. It is also a matter of degree rather than a binary: operations range along a continuum from those whose outputs follow deterministically from verifiable facts to those that

turn substantially on the interpretation of ambiguous, informal, or contested evidence. Confirming that a fingerprint matches an enrolled record sits at the computable end; deciding whether a household's fluctuating informal income places it below an eligibility threshold sits far from it. Where an operation lies high on this continuum, it can move to the system level and human intermediation over it can be displaced; where it lies low, discretionary mediation persists, though it may migrate to those who design and populate the databases. Administrative computability thus conditions *whether* an operation can be disintermediated at all.

3.4 The design of failure

Computability, though, does not by itself fix the outcome, because even a highly computable operation fails part of the time — a fingerprint does not read, a record is missing, a network is down — and what the system does at the point of failure is a matter of design rather than of the technology. This is the article's second organising property. Two conceptual tools specify what is at stake when a computable operation fails.

Herd and Moynihan (2018) show that access to benefits is shaped by the administrative burden a programme places on claimants — the learning, compliance, and psychological costs of finding out about a benefit, meeting its requirements, and enduring the process — and that these costs fall hardest on those least able to absorb them. Much of what a digital system does to a beneficiary who must travel to authenticate, or re-enrol after a failed fingerprint, is legible here: the work of mediation reappears as a cost carried by the citizen rather than the state. Fox (2007) adds a distinction between transparency and accountability that the cases repeatedly bear out. Knowing that something has gone wrong yields no remedy unless some actor or channel can be made to correct it. A system may be entirely transparent about a denial while offering no way to contest it, which makes contestability a separate property from visibility, and one that digitalisation can erode even as it improves the record.

Because “failure” and “failure-handling” cover several distinct things, it helps to disaggregate both. Failures in these systems are of at least four kinds: *technical* failure, where a device, biometric, or network does not perform (a fingerprint that will not read, a point-of-sale terminal offline); *data* failure, where records are missing, mismatched, or wrongly mapped (an Aadhaar number incorrectly seeded, a name spelled inconsistently across databases); *institutional* failure, where the surrounding administrative machinery does not deliver (funds not released, a grievance not answered); and *procedural* failure, where no exception or alternative route exists for cases the system cannot process. The institutional response to these failures can likewise be arrayed along a single dimension that this article treats as decisive. At one pole is a *closed* failure system: when authentication or transfer fails, there is no override, no accessible alternative, and no effective recourse, so the failure resolves into exclusion and its cost falls on the beneficiary. At the other is a *recoverable* failure system: a failed match can be met with an alternative verification, a human override, an offline fallback, or a working grievance channel, so that access is preserved and the cost of failure is absorbed by the state rather than the citizen. The design of

failure-handling — where a system sits between these poles — conditions *what a computable operation actually does* once automated, independent of the technology it uses.

3.5 Propositions

The framework can be stated as four propositions, which Sections 5 through 8 assess against the two cases.

Proposition 1. The likelihood that a welfare function is disintermediated increases with its administrative computability — the degree to which its output can be generated from specified and verifiable inputs without contextual judgment.

Proposition 2. Where a function depends substantially on contextual judgment, digitalisation is more likely to redistribute mediation — relocating discretion to those who design and populate the decision systems — than to eliminate it.

Proposition 3. For functions that are administratively computable, whether digitalisation improves delivery or produces exclusion depends substantially on the design of failure-handling — the availability of exceptions, overrides, alternatives, and recourse — rather than on the technology itself.

Proposition 4. Where failure cannot be resolved through an accessible human or institutional override, the administrative burden of coping with failure is shifted from the state to beneficiaries, and falls hardest on those least able to bear it.

These propositions are analytical rather than statistical; the two cases are used to develop them and to probe whether they hold across dissimilar programmes, not to estimate their frequency across India's welfare architecture.

4. Research design and analytical approach

The article is a theory-building comparative study that develops an analytical framework and probes it against two cases, using secondary empirical evidence. It is useful to be explicit about what this design does and does not attempt, since the cases are asked to do analytical rather than representative work.

Case selection. MGNREGA wage payment and Aadhaar authentication in the PDS were selected for analytical contrast, not statistical representativeness. Both are large-scale, nationally significant welfare programmes; both historically delivered benefits through a concentrated local intermediary; both have been subjected to digital identity and payment infrastructure drawn from the same national stack; and both have attracted an unusually rich body of independent empirical evaluation, including randomised experiments, large household surveys, administrative-data studies, and ethnography. Crucially, they differ in the *mix* of administrative functions they foreground. MGNREGA is organised around employment and wage disbursement, so its digitalisation acts most visibly on the payment and last-mile operations while leaving work

allocation and verification in official hands. The PDS is organised around entitlement to subsidised food and point-of-service collection, so its digitalisation acts most visibly on identity authentication and, at one remove, on eligibility. The two programmes thus expose different combinations of computable and judgment-dependent operations, which is precisely what a functional account requires: if one taxonomy, applied to two dissimilar programmes, sorts their operations in the same way, that consistency is evidence that the taxonomy tracks a real property rather than an artefact of a single case. The design is, in this sense, a most-different-systems comparison at the level of the programme, holding the analytical framework constant. The two cases are not claimed to represent the whole of Indian welfare, and Section 9 treats this as a boundary of the argument.

Evidence base. The analysis draws on four kinds of secondary evidence, deliberately spanning methods and viewpoints so that the argument does not rest on any single tradition: randomised and natural experiments on payment and authentication reforms (Muralidharan et al., 2016, 2025); large household and representative surveys of PDS authentication and Aadhaar exclusion (Drèze et al., 2017; Somanchi et al., 2017; Abraham et al., 2018; Dalberg, 2019); administrative-data and transaction-level studies of payment delay and last-mile access (Narayanan et al., 2019; LibTech India, 2020, 2024); and ethnographic and information-systems research on intermediation and datafication (Chaudhuri, 2019; Masiero & Das, 2019). Official evaluations and their published critiques are included on both sides (Government of India, 2025; Drèze & Khera, 2018), so that the strongest claims of each camp are represented rather than filtered. The evidence was selected purposively to characterise each of the five operations in each programme; the article does not claim a systematic review of the entire empirical corpus, and where the evidence is thin or contested this is noted rather than smoothed over.

Analytical procedure. Welfare delivery in each programme is decomposed into five operations — eligibility determination, identity authentication, benefit computation and disbursement, last-mile access, and grievance-handling. Each operation is then assessed along the two properties developed in Section 3: first, its administrative computability, judged by whether its output can be generated from specified and verifiable inputs without contextual judgment under the arrangements that actually obtain; and second, the design of its failure-handling, judged by whether the system provides exceptions, overrides, alternatives, and recourse when the operation fails, and by who bears the resulting cost. The same two questions are put to each operation in each case, and the results are arrayed in a common matrix (Section 8, Tables 1 and 2). The propositions of Section 3.5 are treated as claims to be probed against this matrix: the design looks for whether computability tracks disintermediation across operations (Propositions 1 and 2), and whether, holding computability constant, variation in failure design tracks variation in outcomes (Propositions 3 and 4). The identity-authentication operation is analytically pivotal because it is highly computable yet produced opposite outcomes in the two implementations, which isolates the effect of failure design from that of computability.

5. The pre-digital architecture of welfare delivery

Before digitalisation, both programmes delivered benefits through a set of operations concentrated in local hands, and reconstructing that architecture allows the effect of digital systems to be located within it rather than described in the aggregate.

Under MGNREGA, enacted in 2005 to guarantee up to a hundred days of wage employment a year to rural households (Mahatma Gandhi National Rural Employment Guarantee Act, 2005), the sequence ran from registration to payment through the same apparatus. A worker registered with the gram panchayat; panchayat and block officials identified and sanctioned worksites; officials verified the work at the site; wages were computed from the record; and the wage was paid, historically through post offices, banks, or cash at the worksite. The operations were integrated, and this is what mattered politically: the local authority that allocated work also verified it and controlled its payment, so that power over the wage was continuous with power over access. Payment could be delayed to discourage further claims, inflated for favoured workers, or withheld from opponents. It functioned as an instrument of a local political economy of employment rather than as a discrete administrative step (Muralidharan et al., 2016).

The PDS delivered subsidised food through a parallel concentration. Eligible households, identified against state criteria, held a ration card; at the fair-price shop the dealer confirmed the cardholder's identity, recorded the transaction, and issued the monthly entitlement. Here too a single actor determined, in practice, whether a claimant would be recognised, how strictly quantities would be honoured, and whether the records matched what was actually given. That concentration produced two characteristic diversions: siphoning against the names of people who received nothing, and quantity cuts skimmed from those who received something (Drèze et al., 2017; Muralidharan et al., 2025).

Both programmes, then, bundled eligibility determination, identity verification, benefit computation, last-mile delivery, and grievance-handling into one mediating layer of concentrated human intermediation. Digitalisation did not replace that layer as a whole. It entered at particular points within it, and the task is to see which operations it reached, which it left in place, and what became of the people who had performed them.

6. MGNREGA: digitising the payment rail

The change that DBT brought to MGNREGA was, at bottom, a change to one operation. Once work is verified and a wage computed, an electronic Fund Transfer Order moves the amount from the treasury toward the worker's bank account; the Aadhaar-based payment routing later standardised this, directing the credit to whichever account is mapped to the worker's Aadhaar number. The verification of work and the allocation of it were untouched. The pre-determined wage now travels along a digital rail rather than through an official's hands.

Where that rail is implemented well, the gains are genuine, and it would misrepresent the evidence to minimise them. The Andhra Pradesh Smartcard evaluation remains the clearest demonstration. Across 157 sub-districts and some nineteen million people, the reform produced a faster, more predictable, and less corrupt payment process, cutting leakage in the scheme by about 41 percent relative to the control group without reducing access, and generating time savings to workers that alone matched the intervention's cost; beneficiaries preferred it (Muralidharan et al., 2016).

One feature of that result complicates the vocabulary of removal, and it is worth elevating from an empirical footnote to an analytical point. The Andhra Pradesh gains came not from eliminating the human intermediary but from *redesigning* it: locally hired, bank-employed agents authenticated beneficiaries biometrically and paid them in cash in the village (Muralidharan et al., 2016). What improved payment was a more accountable mediator, not the absence of one. This suggests a general possibility that the disintermediation vocabulary obscures — that digitalisation may improve delivery less by removing human intermediation than by *reconstituting* it in a more accountable form, substituting one kind of human mediation for another while automating only the record and the transfer behind it. The national model that followed took the opposite tack, dropping the in-village agent and directing beneficiaries to banks and banking correspondents, and its record is correspondingly uneven. The contrast is itself instructive: the strongest single piece of evidence for “disintermediation” in fact involved a redesigned intermediary, and the move to a thinner human layer at the last mile did not straightforwardly improve on it.

At national scale, the claim that electronic transfer made payment prompt does not hold. Analysing roughly nine million transactions across ten states for 2016–17, Narayanan and colleagues (2019) found that only about one payment in five was completed within the mandated fifteen days once the full cycle was counted. Their contribution is to separate the two “legs” of the transfer — the state's generation of the payment order, and the central government's subsequent release of funds — and to show that official figures, by measuring only the first, understated delay substantially. On their account, delay is driven less by the payment technology than by the flow of funds behind it: where budgets are short, the rail cannot move money that has not been released. Both the Ministry of Finance and the Supreme Court, they note, have acknowledged the problem (Narayanan et al., 2019). This is an *institutional* failure in the terms of Section 3.4, and it is largely independent of any failure in the technology.

The rail delivers money to an account, not to a hand, and the account must still be reached. Across much of rural India this last mile runs through banking correspondents and common service centres, and the passage is neither quick nor dependable. A multi-state study found that roughly four to five percent of wage transactions were rejected outright, that workers in Jharkhand and Rajasthan spent on the order of four to five hours obtaining their wages against about an hour in Andhra Pradesh, and that the large majority of grievances went unresolved (LibTech India, 2020). The Aadhaar-based routing introduced its own failures: transfers rejected

because an Aadhaar number was not correctly seeded, and transfers misdirected to another person's account through mapping errors that are difficult to detect and harder to reverse (Narayanan et al., 2019) — *data* failures, in the terms used here. When Aadhaar-based payment was made effectively mandatory, culminating in the contested rollout of the Aadhaar-based payment system as the default for wage disbursement, the accompanying churn of account mapping coincided with the deletion of large numbers of workers from the rolls, a share of them wrongly, prompting protest and judicial attention (LibTech India, 2024).

Set against the pre-digital architecture, the pattern sorts by operation. The disbursement rail has been genuinely digitalised, and where funds are released it improves on the discretionary, delay-prone distribution it replaced. Work allocation and verification have not been automated at all and remain with panchayats and block officials, because deciding which projects to sanction and confirming that earthwork was actually done are judgments entangled with local interests rather than computations over verifiable inputs — operations low in administrative computability. Last-mile access has changed hands rather than disappeared, moving from post offices and worksite cash to banking correspondents and service centres — private actors with their own incentives and widely varying reliability, a re-intermediation rather than a disintermediation. Grievance-handling has become a more salient and poorly performing operation, as the failures thrown up by the rail and the routing generate problems the system is ill-equipped to resolve. The image of money passing directly from treasury to citizen describes one operation accurately and the operations around it not at all.

7. Aadhaar and the Public Distribution System: identity without eligibility

If MGNREGA shows the payment rail coming apart from the operations around it, Aadhaar in the PDS shows two operations that are habitually conflated — verifying identity and determining eligibility — behaving in opposite ways. Aadhaar, launched in 2009 and covering more than 1.2 billion residents within a decade (Abraham et al., 2018), assigns a unique identifier tied to biometric and demographic records; enrolment has since risen further, and the system now processes tens of billions of authentication transactions annually. It has entered the PDS at two points: the seeding of ration cards to Aadhaar numbers, and the authentication of the cardholder at the shop, where collecting the monthly ration now typically requires a fingerprint match on a point-of-sale device.

For the authentication operation, this is genuine disintermediation, and it should be named as such. Confirming that the person at the counter is the enrolled cardholder is a matter of matching a biometric to a record, a computation that needs no human judgment — an operation high in administrative computability. Where the machinery works, the dealer's former role in vouching for or disputing a claimant's identity is simply gone; the device performs it. The identity check that a frontline worker once exercised has, in Bovens and Zouridis's (2002) terms, left their hands. Among all the operations in either case, this is the clearest instance of a human function being displaced into technical mediation rather than merely relocated.

Eligibility is a different matter, and in the cases examined Aadhaar does not automate it. Deciding whether a household meets the income, residence, and composition criteria for subsidised food requires interpreting evidence that is often informal, irregular, or inconsistent across sources — an operation low in administrative computability. Aadhaar de-duplicates identity and blocks one person from holding two entitlements; it does not compute deservingness. This limit is not merely observed by critics but built into the architecture: the JAM infrastructure was not designed to verify the socio-economic data on which eligibility rests, and the assessment of eligibility criteria continues to be performed by scheme administrators against their own databases (International Monetary Fund, 2023). Masiero and Das (2019), studying the datafied PDS, draw the same line from the other side, distinguishing the system's capacity to recognise a beneficiary from its capacity to assign an entitlement, and locating the harder problems in the latter. What digitalisation changes is the manner of the eligibility judgment, not its necessity. Where an official once weighed paper documents face to face, the determination now runs through the cross-referencing of digital records, and the discretion involved shifts to those who decide which databases are consulted, how conflicts among them are read, and what is done when a record is missing. The judgment persists, relocated from the counter to the console — a redistribution of mediation rather than its removal.

The most instructive evidence, though, concerns the authentication operation whose computability seemed to promise a clean displacement, and it shows that computability conditions less than one might expect. The largest evaluation, covering some fifteen million beneficiaries in Jharkhand, found that mandatory biometric authentication in the PDS did reduce corruption, but that between 1.5 and 2 million legitimate beneficiaries lost access to their rations at some point during the reforms (Muralidharan et al., 2025). The authors' interpretation carries as much weight as the number. The adverse effects, they conclude, were driven primarily by decisions about how the transition was managed rather than by the technology, and the contrast with their Andhra Pradesh results reflected differences in programme design and emphasis — Andhra Pradesh's rules aimed to improve the beneficiary's experience, Jharkhand's to secure fiscal savings — rather than differences in state capacity (Muralidharan et al., 2025). The same authentication technology improved delivery under one design and excluded millions under another. In the terms of Section 3.4, one implementation approximated a recoverable failure system and the other a closed one.

Independent survey research points the same way while weighting the costs more heavily. Surveying households across villages in Jharkhand, Drèze and colleagues (2017) found that biometric authentication raised exclusion and transaction costs for the most vulnerable — the elderly, manual labourers, those whose fingerprints no longer read cleanly — while failing to reach the quantity cuts that made up much of the diversion, since skimming from a genuine transaction is invisible to an identity check. Exclusion of this kind appeared even in Hyderabad, where connectivity and administration were comparatively favourable (Somanchi et al., 2017). Representative estimates put the magnitude in perspective: a three-state survey attributed

between 0.8 and 2.2 percent of PDS exclusion to Aadhaar-related failures — a share that, in those three states, corresponds to roughly two million people, a separate estimate from the Jharkhand figure though of similar magnitude — even as overall satisfaction with Aadhaar remained high and the offline authentication that might have reduced such failures went underused (Abraham et al., 2018; Dalberg, 2019). The scholarly disagreement over the Jharkhand experiment — whether the leakage it reduced justified the exclusion it caused (Drèze et al., 2020) — is itself evidence for the larger argument, since the two sides are weighing the same real and offsetting effects, and the balance between them is conditioned by design.

What makes authentication failure consequential is not only its frequency but its effect on the citizen's position. A dealer who refused a claimant face to face could be argued with, appealed to, or reported to an elected representative; that mediation, whatever its faults, was visible and could be worked on. When the point-of-sale device declines to authenticate and the dealer can say only that the system has refused, the same denial becomes opaque and non-negotiable, and the citizen is left with a transparent outcome and no evident route of recourse (Fox, 2007). The exclusion that Peeters and Widlak (2018) attribute to information architecture is exactly this: a person whose biometric does not read, or whose seeding is wrong, is turned away by the design rather than by a decision, and unless the design provides an exception path, there is no one to persuade. The work of resolving the failure — the return visits, the re-enrolment, the travel — falls on the person least able to bear it (Herd & Moynihan, 2018). Authentication, the operation that looked most cleanly automatable, regenerates mediation precisely at the point of its failure, and how much failure it generates is conditioned by whoever writes the rules for exceptions.

8. Comparing functions across the two cases

Read together, the two cases yield a consistent pattern that neither the disintermediation reading nor the persistence reading captures alone. Table 1 sets out the five operations across both programmes along the dimensions developed in Section 3; Table 2 arrays the principal evidence by case, method, and outcome, so that the empirical basis for the pattern is visible in one place.

Table 1. How digitalisation affects distinct welfare functions

Function	Computability	What digitalisation did	Effect on mediation	Main mode	failure	Who bears the cost
Eligibility	Low	Paper checks replaced by record cross-checking	Persists; discretion moves upstream to record-holders	Data: missing or wrong records; ambiguous cases		Applicant, via delay or denial
Authentication	High	Biometric	Displaced	Technical/data:		Beneficiary

Function	Computability	What digitalisation did	Effect on mediation	Main mode	failure	Who bears the cost
n		match replaces human vouching	where it works; returns when it fails		mismatch, connectivity, seeding	, via return visits or exclusion
Payment	High	Computed wage routed via electronic transfer	Reduced where funds are released	Institutional/data : fund-release delay; misrouted transfers		Worker, via delayed or withheld wages
Last-mile access	Low–medium	Cash-out shifted to agents and centres	Re-intermediate d to private agents	Institutional: distance, agent liquidity		Beneficiary , via time, travel, cost
Grievance	Low	Digital records and formal complaint channels	Newly salient but weakly embedded	Procedural: no override; unresolved complaints		Beneficiary , via unremedied loss

Note. “Computability” denotes administrative computability (Section 3.3): the degree to which an operation yields an administratively valid decision from specified and verifiable inputs without contextual judgment. Failure modes follow the typology in Section 3.4 (technical, data, institutional, procedural).

Table 2. Principal evidence by case, method, and outcome

Case and function	Intervention	Method	Reported outcome	Resulting mediation
Andhra Pradesh NREGS payment	Biometric Smartcards with in-village agents	Randomised evaluation	Leakage down ~41%; faster payment; savings; beneficiaries preferred it	Redesigned, more accountable human agent
National MGNREGA payment	DBT and Aadhaar-based routing	Administrative-data / transaction study	~1 in 5 payments within 15 days once full cycle	Institutional and technical

Case function	and Intervention	Method	Reported outcome	Resulting mediation
			counted; delay driven by fund release	
National MGNREGA — last mile	Banking — correspondents and service centres	Multi-state field study	~4–5% transactions rejected; 4–5 hours to obtain wages in some states; most grievances unresolved	Re-intermediated to private agents
Jharkhand PDS — authentication	Aadhaar-based biometric authentication	Randomised and natural experiments	Corruption down; 1.5–2 million lost access at some point during reforms	Technical, closed failure design
Jharkhand PDS — authentication	Aadhaar-based biometric authentication	Household survey	Exclusion and transaction costs concentrated on the most vulnerable; quantity fraud untouched	Technical plus shifted citizen burden
Three-state PDS — authentication	Aadhaar seeding and authentication	Representative survey	0.8–2.2% of exclusion attributable to Aadhaar failures (~2 million people); offline mode underused	Technical, avoidable with fallback

Note. Sources as cited in Sections 6 and 7. The two “roughly two million” figures — the Jharkhand PDS exclusion (Muralidharan et al., 2025) and the three-state survey estimate (Abraham et al., 2018; Dalberg, 2019) — are distinct estimates of similar magnitude, not the same population counted twice.

Two operations anchor the extremes, and they map onto Propositions 1 and 2. Benefit disbursement — computing and moving a determined amount — is administratively computable, and where the conditions behind it hold, it is automated and human intermediation over it is displaced; this is the truth in the disintermediation reading. Eligibility determination, in the cases examined, is not; it rests on interpreting ambiguous evidence, so its mediation persists and

migrates to those who control the records; this is the truth in the persistence reading. The two readings are observations about different operations, and once the operations are separated the appearance of contradiction largely dissolves. Proponents who point to direct payment and critics who point to continuing gatekeeping are, for the most part, both right, because they are describing operations at opposite ends of the computability continuum.

The cases also identify what conditions the outcome where computability alone is silent, which is the substance of Propositions 3 and 4. A first condition is dependence on contextual judgment, which conditions whether an operation can move to the system level at all: disbursement can, eligibility cannot. A second condition is the design of failure-handling, which conditions what a computable operation actually does once automated. Authentication isolates the second condition, because it is computable in principle yet produced improved delivery in Andhra Pradesh and the exclusion of some two million people in Jharkhand, and the difference lay in the protocols governing the transition and the treatment of failures rather than in the technology or in state capacity (Muralidharan et al., 2025). The pattern is consistent with the closed-versus-recoverable distinction of Section 3.4: the Andhra Pradesh design, with in-village agents and rules oriented to the beneficiary's experience, approximated a recoverable failure system; the Jharkhand rollout, faster and oriented to fiscal savings, with the offline fallback underused, approximated a closed one. A computable operation is not thereby a benign one. Whether its automation serves beneficiaries or excludes them appears, in these cases, to turn on whether exceptions can be made, whether recourse exists, and who is made to carry the cost when the match fails — the conditions that Peeters and Widlak (2018) gather under the architecture of exclusion and that Herd and Moynihan (2018) track as administrative burden.

Failure-handling is one condition among several, not a master variable, and the cases are clear about the others. Fund flow conditions whether the disbursement rail delivers on time, independent of any failure design (Narayanan et al., 2019). The density of banking and the reliability of agents condition the last mile. What the cases establish is narrower and more useful than a general theory: that the effect of digitalising a computable operation cannot be read off its computability, and that the institutional handling of failure is a decisive part of the difference. The through-line across both programmes is that mediation is seldom removed cleanly. It is redesigned into a more accountable agent, as in Andhra Pradesh; passed to a private correspondent at the last mile; shifted upstream to those who write the decision logic, as with eligibility; or folded into the system itself, as with authentication failure. This is consistent with the migration of discretion that Bovens and Zouridis (2002) describe, with the broader accounts of algorithmic governance and the digital welfare state (Katzenbach & Ulbricht, 2019; Eubanks, 2018; Alston, 2019), and with Chaudhuri's (2019) account of intermediation reproduced within Aadhaar. The contribution here is to specify what those general claims leave open: which operations the migration completes, which it stalls, and the design conditions under which a displaced operation returns as exclusion.

9. Objections and limits

Because the pattern is consistent across two very different programmes, it is worth answering the strongest objections directly and then stating the boundaries of the claim.

The first objection is that this is ordinary electronic government — that relocating administrative work into systems is a general feature of computerised bureaucracy and nothing specific to Indian welfare. The mechanism is indeed general, and the article claims no more; it is the same migration of discretion that the digital-era-governance and algorithmic-governance literatures describe (Dunleavy et al., 2006; Katzenbach & Ulbricht, 2019). Its value lies in disaggregating that mechanism by operation, in adding failure-handling as a second condition, and in doing so where the stakes of a failed transaction are subsistence rather than inconvenience. That the mechanism is general is a reason to credit the account, not to set it aside.

A second objection holds that persistent mediation is a transitional problem of implementation that maturing systems will dissolve. The possibility cannot be excluded, and the leading experimental evidence cuts in a particular direction here: the Jharkhand exclusions are attributed to how the reform was managed rather than to the technology (Muralidharan et al., 2025). That finding supports rather than undermines the argument, because the claim is not that eligibility mediation is permanent but that outcomes across operations are conditioned by computability and by failure design — and if exclusion is an implementation artefact, then failure design is precisely the variable in question. What the article declines to assert is the stronger structural claim, that any operation is necessarily and forever beyond automation, which would require longitudinal comparison of high- and low-capacity systems over time that the present evidence does not provide. This is why the article speaks of administrative computability under the arrangements that obtain, and phrases eligibility's resistance as a feature of the cases examined rather than a law.

A third objection runs the other way, that digital systems eliminate more mediation than an emphasis on persistence allows. This should be granted without reservation. The reductions in leakage and in fictitious beneficiaries documented in both the Andhra Pradesh and Jharkhand experiments are real and substantial (Muralidharan et al., 2016, 2025), and even the contested official savings figures likely capture some genuine reduction in diversion. Any fair account registers these gains alongside the exclusions. The claim is not that digitalisation fails but that its effects are operation-specific and design-contingent, which is compatible with real gains in some operations and serious costs in others.

A fourth objection is that intermediaries persist because of weak state capacity rather than anything about the operations themselves. Capacity plainly matters, but the sharpest available evidence attributes the Andhra Pradesh–Jharkhand divergence to programme design and emphasis rather than to capacity (Muralidharan et al., 2025), and even under high capacity the eligibility judgment remains a judgment — as the architectural exclusions documented in high-

capacity administrations attest (Peeters & Widlak, 2018; Eubanks, 2018). Capacity and design both shape outcomes; the operational structure of delivery organises the space within which they do so.

Three boundaries of the argument should be equally plain. First, it rests on two programmes and on secondary evidence concentrated in a handful of states, and does not survey the whole of India's welfare architecture; the functional taxonomy is offered as a framework to be tested elsewhere, not as an established generalisation. Second, it is silent, by design, on questions that administrative evidence cannot answer — whether digital welfare has altered patronage, brokerage, or electoral competition — which would require political-economy research not attempted here and which the literature on brokerage suggests are not settled by administrative disintermediation alone (Auerbach & Thachil, 2018). Third, the beneficiary's own experience of these systems — how directness is perceived, whom people turn to when the device refuses — remains thin relative to the transaction-level metrics on which official evaluation concentrates, and the framework would be sharpened by qualitative work at that register. Marking these boundaries is part of keeping the claim defensible.

10. Conclusion

India's Digital Public Infrastructure has not replaced a mediated welfare state with an unmediated one, and it has not left the old arrangements in place. It has acted on welfare delivery operation by operation: automating the disbursement rail where funds are released, leaving the eligibility judgment intact while moving it from the counter to the console, and displacing the verification of identity while generating, at the point where verification fails, a newer and less answerable form of exclusion. Mediation proves durable across both cases. What digitalisation changes is *where* mediation sits and *who* exercises it — a redesigned agent, a private correspondent, the designers of the decision logic, or the system itself. The organising finding is therefore not that mediation persists but that it is *redistributed*: across administrative functions, according to their computability, and across institutional locations, according to the design of failure. And whether automating a computable operation helps beneficiaries or excludes them is conditioned less by the technology than by the way the system is built to handle its own failures, as the divergence between two implementations of one authentication technology makes plain.

The significance of this for the relationship between the state and the citizen is narrower than the language of digital transformation implies, and steadier for being narrow. What changes when mediation is redistributed in these ways is its visibility and its contestability. A denial that could once be argued with becomes one that cannot; discretion that could be located in a person moves to a system whose logic is set elsewhere; and the effort of coping with failure passes from the state to the citizen, weighing most on those with the least capacity to absorb it. These are consequences for administrative access, responsiveness, and recourse — for the texture of state–citizen relations at the point of delivery — not for representation or electoral politics, on which the evidence here is silent. They imply that the useful question for reform is not whether

mediation persists, since in some operations it must, but whether the mediation that persists stays visible to the people it governs and open to their challenge.

The framework is offered as a tool rather than a verdict. Existing scholarship has shown both that digital systems can disintermediate welfare and that intermediaries endure, and the apparent conflict between these findings largely resolves once welfare delivery is disaggregated into its constituent operations, which differ in their administrative computability and in the design of their failure-handling. Seeing which operations digitalisation displaces, which it leaves intact, and why comparable instruments improve delivery in one place and exclude in another does not overturn the existing accounts. It explains how both can be true at once — and it suggests where the design of digital welfare, if it is to preserve access rather than merely secure savings, most needs attention: not in the choice to automate a computable operation, but in what the system is built to do when that operation fails.

UNDER PEER REVIEW

680 **References**

- 681 Abraham, R., Bennett, E. S., Sen, N., & Shah, N. B. (2018). *State of Aadhaar report 2017–18*.
682 IDinsight.
- 683 Aiyar, S. (2017). *Aadhaar: A biometric history of India's 12-digit revolution*. Westland.
- 684 Alston, P. (2019). *Report of the Special Rapporteur on extreme poverty and human rights:*
685 *Digital welfare states and human rights* (UN Doc. A/74/493). United Nations General Assembly.
- 686 Andrews, L. (2019). Public administration, public leadership and the construction of public value
687 in the age of the algorithm and 'big data'. *Public Administration*, 97(2), 296–310.
- 688 Auerbach, A. M., & Thachil, T. (2018). How clients select brokers: Competition and choice in
689 India's slums. *American Political Science Review*, 112(4), 775–791.
- 690 Bovens, M., & Zouridis, S. (2002). From street-level to system-level bureaucracies: How
691 information and communication technology is transforming administrative discretion and
692 constitutional control. *Public Administration Review*, 62(2), 174–184.
- 693 Busch, P. A., & Henriksen, H. Z. (2018). Digital discretion: A systematic literature review of
694 ICT and street-level discretion. *Information Polity*, 23(1), 3–28.
- 695 Chaudhuri, B. (2019). Paradoxes of intermediation in Aadhaar: Human making of a digital
696 infrastructure. *South Asia: Journal of South Asian Studies*, 42(3), 572–587.
- 697 Citron, D. K. (2008). Technological due process. *Washington University Law Review*, 85(6),
698 1249–1313.
- 699 Corbridge, S., Williams, G., Srivastava, M., & Véron, R. (2005). *Seeing the state: Governance*
700 *and governmentality in India*. Cambridge University Press.
- 701 Dalberg. (2019). *State of Aadhaar: A people's perspective*. Dalberg.
- 702 Drèze, J., Khalid, N., Khera, R., & Somanchi, A. (2017). Aadhaar and food security in Jharkhand:
703 Pain without gain? *Economic and Political Weekly*, 52(50).
- 704 Drèze, J., & Khera, R. (2018, February 8). Aadhaar's \$11-bn question: The numbers being
705 touted by government have no solid basis. *The Economic Times*.
- 706 Drèze, J., Khera, R., & Somanchi, A. (2020). *Balancing corruption and exclusion: A rejoinder*.
707 Ideas for India.
- 708 Dunleavy, P., Margetts, H., Bastow, S., & Tinkler, J. (2006). New public management is dead—
709 Long live digital-era governance. *Journal of Public Administration Research and Theory*, 16(3),
710 467–494.

711 Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the*
712 *poor*. St. Martin's Press.

713 Fox, J. (2007). The uncertain relationship between transparency and accountability. *Development*
714 *in Practice*, 17(4–5), 663–671.

715 Government of India. (2025). *A quantitative assessment of India's Direct Benefit Transfer system*
716 (BlueKraft Digital Foundation, cited by the Ministry of Finance). Government of India.

717 Gupta, A. (2012). *Red tape: Bureaucracy, structural violence, and poverty in India*. Duke
718 University Press.

719 Herd, P., & Moynihan, D. P. (2018). *Administrative burden: Policymaking by other means*.
720 Russell Sage Foundation.

721 International Monetary Fund. (2023). *Stacking up the benefits: Lessons from India's digital*
722 *journey* (IMF Working Paper No. 2023/078). International Monetary Fund.

723 Katzenbach, C., & Ulbricht, L. (2019). Algorithmic governance. *Internet Policy Review*, 8(4).

724 LibTech India. (2020). *Length of the last mile: Delays and hurdles in NREGA wage payments*.
725 Azim Premji University.

726 LibTech India. (2024). *MGNREGA implementation in India: Insights and trends, April–*
727 *September 2024*. LibTech India.

728 Lipsky, M. (1980). *Street-level bureaucracy: Dilemmas of the individual in public services*.
729 Russell Sage Foundation.

730 Mahatma Gandhi National Rural Employment Guarantee Act, No. 42 of 2005 (India).

731 Masiero, S., & Das, S. (2019). Datafying anti-poverty programmes: Implications for data justice.
732 *Information, Communication & Society*, 22(7), 916–933.

733 Maynard-Moody, S., & Musheno, M. (2003). *Cops, teachers, counselors: Stories from the front*
734 *lines of public service*. University of Michigan Press.

735 Ministry of Finance, Government of India. (2013). *Direct Benefit Transfer: Operational*
736 *guidelines*. Department of Economic Affairs.

737 Muralidharan, K., Niehaus, P., & Sukhtankar, S. (2016). Building state capacity: Evidence from
738 biometric smartcards in India. *American Economic Review*, 106(10), 2895–2929.

739 Muralidharan, K., Niehaus, P., & Sukhtankar, S. (2025). Identity verification standards in
740 welfare programs: Experimental evidence from India. *The Review of Economics and Statistics*,
741 107(2), 372–392.

- 742 Narayanan, R., Dhorajiwala, S., & Golani, R. (2019). Analysis of payment delays and delay
743 compensation in MGNREGA: Findings across ten states for financial year 2016–2017. *Indian*
744 *Journal of Labour Economics*, 62(1), 113–133.
- 745 Peeters, R., & Widlak, A. (2018). The digital cage: Administrative exclusion through
746 information architecture — The case of the Dutch civil registry's master data management
747 system. *Government Information Quarterly*, 35(2), 175–183.
- 748 Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition*
749 *have failed*. Yale University Press.
- 750 Somanchi, A., Bej, S., & Pandey, M. (2017). Well done ABBA? Aadhaar and the public
751 distribution system in Hyderabad. *Economic and Political Weekly*, 52(7).
- 752 Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*,
753 12(4), 505–523.