

From Digital Public Infrastructure to Intelligent Finance: What the World Can Learn from India

Sameer Gaikwad¹, Nivedita Chaskar²

¹Independent Director (IICA-certified, DIN 10727192); Mumbai, India

²Independent Director (IICA-certified); Mumbai, India

Abstract

India has built the world's largest population-scale digital public infrastructure for finance, and is now layering artificial intelligence onto it. This paper examines that transition and asks which parts of the Indian experience are transferable to other jurisdictions. Six potential failure modes accompany India: authentication exclusion, predatory digital lending, fraud. In the paper, the Indian stack is mapped across three layers of consumer identity, payments and consented data, and in parallel it is shown how a fourth layer of standardised credit origination and machine-learning underwriting emerged. Key findings include that the Indian payment layer reached 241.6 billion transactions in the financial year ended March 2026, but so-called 'last-mile' infrastructure problems mean the rate of year-on-year growth has decelerated from 82 per cent to 30 per cent in the previous four years. The third, consent layer has enabled more than 3.82 trillion rupees (\$44.95 trillion) of credit and, by September 2026, had already surpassed 500 million fulfilled consent requests. The paper shows that the decisive variable was architectural, not technological and the state has supplied thin, open, interoperable protocols, leaving product competition to market forces. Four potential failure modes – authentication exclusion, predatory digital lending, transaction fraud enabled by lowfriction onboarding and concentration at the application layer – are also identified. Six design principles are identified which are transferable to other jurisdictions and which policymakers can use to manoeuvre around scale economics and specific institutional forms. The paper ends by providing a ten-point sequencing framework for policymakers.

Keywords: Digital Public Infrastructure, Financial Inclusion, Unified Payments Interface, Account Aggregator, Artificial Intelligence In Credit, Open Finance, Regulatory Technology, India Stack

1. Introduction

For two decades, the global financial inclusion policy agenda has been driven by one indicator: whether an adult holds a transaction account. On that indicator, there has been much to celebrate. On these other indicators, whether the account is used, whether usage produces a credit history and whether the history translates into affordable credit, there has been little to celebrate.

Three structural obstacles remain across most developing markets and some developed ones. It remains costly for small merchants to accept payments electronically, which forces them to transact in cash and corresponding to invisibility to lenders. Credit bureau files remain thin for informal and self-employed workers, who account for the majority of the workforce in most emerging economies. Financial data is locked within incumbent institutions and open banking mandates tend to cover only banks and to lack portable consent standards.

India has addressed all three obstacles at the protocol layer rather than the product layer. It has done this from starting from a position of low per capita income, low formal employment and weak state administrative capacity. That starting position appears to more closely mirror the conditions in much of Sub-Saharan Africa, South Asia and Latin America, rather than the United Kingdom, Singapore and the European Union, which have been the sites of most open banking scholarship.

1.1 Research problem

The large body of descriptive literature on the extent of Indian digital payments is well known. Much less literature exists on the architecture behind the Indian digital payments success, on how to differentiate between architectural elements and those that are uniquely Indian, and the follow-on transition currently underway from payments to credit and decision infrastructures. The question facing policymakers considering an Indian-style program is not the extent of success in India, but rather which elements can be replicated in other countries.

1.2 Contribution

The paper makes four contributions. First, it gives an architectural view of the Indian stack through the four layers instead of looking at payments alone. Second, it identifies the empirical markers of success and failure for each of the four layers as of September 2026 through primary regulatory and industry sources. Third, it examines failure modes with the same rigor as success stories, on the grounds that the failure stories make better lessons for other jurisdictions. Fourth, it develops a transferability framework which differentiates design principles from institutions and economics of scale.

1.3 Structure of the paper

Section 2 covers the literature review. Section 3 outlines the research goals and methods. Section 4 provides the theoretical framework. Sections 5 through 8 cover the empirical analysis regarding payment layer, global positioning, consent and data layer, and credit and machine learning layer respectively. Section 9 analyzes the governance framework. Section 10 provides a critical evaluation of the failure modes. Section 11 reviews international diffusion. Section 12 develops the transferability framework. Sections 13 and 14 cover the limitations and conclusions.

2. Review Of Literature

The relevant scholarship falls into four streams, which have developed largely in parallel and are seldom integrated.

2.1 Digital public infrastructure as a category

The idea of digital public infrastructure became official in the multilateral world when the nations comprising the G20 agreed to adopt a framework in support of systems that are by nature open, interoperable, secure, and inclusive. The International Monetary Fund has described the principle as “digital rails for inclusive growth” rather than proprietary technology and has acknowledged the Indian payments system as the biggest real-time payments system in the world in terms of transaction volume. It is estimated that around 2.9 billion people do not have any form of digital identity to conduct online business by the World Bank, based on its 2026 digital public infrastructure program..

2.2 Financial inclusion and the usage gap

Inclusion research has increasingly moved away from looking at ownership and towards the use of accounts, and then from use to access to credit. The constant pattern in the findings is that ownership gains do not lead automatically to use, and use gains do not lead automatically to credit. Accounts are inactive and informal credit is prevalent for micro-enterprises in many markets, including India.

2.3 Open banking and consent architecture

Research done on open banking regimes in comparison has found that mandate design is what makes or breaks their adoption. Those regimes which have put into place access requirements on banks only, without the consent mechanism being portable and without including non-banking financial data, have been found to be lagging behind expectations in terms of adoption. This is not the case with the Indian account aggregator regime, which differs in two ways from the aforementioned trend.

2.4 Artificial intelligence in credit risk assessment

The risk assessment literature divides stock underwriting, which depends on past borrowing and repayment experiences, from flow underwriting, which depends on observed cash flows. For borrowers who do not have a bureau file, the two are not substitutes since only the latter type of information is available. Machine learning techniques used on transaction histories, tax returns, and merchant settlements have been found capable of uncovering repayment capacity information that is not accessible to score-based algorithms. Algorithmic fairness literature warns that machine learning techniques, when deployed on biased data sets, can reinforce exclusion and escalate exclusion from simple refusal into unfavorable pricing that the borrower cannot see or challenge.

2.5 Identified research gap

The four layers are inadequately linked to each other. Payment studies do not often delve into consent-related issues. Open-banking studies do not often study how payments can provide data that may be used as inputs for underwriting purposes. Artificial Intelligence studies also do not often focus on the impact of infrastructure on the availability of input data. This paper fills this gap by looking at the layers as an interdependent system.

3. Research Objectives And Methodology

3.1 Objectives

1. To create a map of India's layered architecture for its digital public infrastructure for finance and understand the design principles behind it.
2. To compile and analyze the present empirical indicators of usage at the payment, data, and credit layers.
3. To analyze the move from transactional infrastructure to decision-making infrastructure and the use of machine learning for credit decisioning.
4. To study the governance and regulation architecture behind this move.
5. To study failure modes and what aspects of the model can be replicated elsewhere.

3.2 Research questions

1. What are the architectural features responsible for the scale of adoption, aside from features particular to India?
2. Is there any evidence that data infrastructure and machine learning underwriting result in access expansion rather than pricing alone?
3. Which parts of the model are portable, institution-dependent, and population-size dependent?

3.3 Method

This study is essentially descriptive and analytical, relying solely on secondary data. There was neither any primary survey nor any institutional data used. There are four sources, which are described in Table 1 below.

Table 1. Data Sources And Their Use In This Study

Source category	Specific sources	Data used
Government and regulatory	Press Information Bureau, Reserve Bank of India publications, Lok Sabha replies	Transaction volumes and values, bank participation, regulatory frameworks
Industry self-regulatory	Sahamati ecosystem dashboard and impact reports, National Payments Corporation of India statistics	Consent counts, participant counts, credit enabled
Multilateral	International Monetary Fund, World Bank digital public infrastructure programme, Group of Twenty framework	Comparative positioning, global identity gap
Independent media and research	Business press, policy research institutes, cyber-intelligence reports	Failure mode evidence, adoption of exported components

3.4 Period and scope

The research time period is from April 2016 (which marks the launch of the unified payments interface) to September 2026. The years in financial statements refer to the Indian accounting practice of year ending on 31st March, which means that 2026 is the year from April 2025 to March 2026. Monetary values are given in Indian rupees in the Indian numbering system, where one crore is equivalent to ten million.

3.5 Analytical approach

The indicators have been abstracted from sources and organized by layer. The growth rates given in this document have been calculated by the author himself from the underlying annual values provided, and are designated as such. The architectural and governance research is based on the documentary study of the regulations. The transferability framework presented in Section 11 has been arrived at inductively by categorizing each individual design characteristic into one of the three categories depending on the criteria mentioned above.

4. Conceptual Architecture Of The Indian Stack

In the Indian approach, the system comprises different layers that operate separately from each other, are supervised using different mechanisms, and fail independent of one another. This layering itself is the fundamental design principle and not just an accidental one.




Figure 1. Layered Architecture Of India's Digital Public Infrastructure For Finance.

Source: Author's own construction based on regulatory documentation.

Layer upon layer strips away the ability of the lower layer. The digital identity layer lowered the marginal cost of account opening, making it economical to have low balance accounts. These accounts created the payments data. Once made portable by consent, the payments data became underwriting data. Underwriting data once standardized through a common gateway became the foundation of automated credit decisions.

4.1 Layer descriptions

Layer 1, identity. It is the novel identification number based on biometrics and online verification through electronic know-your-customer application programming interfaces. What matters economically is not the identity credential itself but the reduction in onboarding costs from a branch visit and paperwork to something programmable.

Layer 2, payments. Unified Payments Interface

UPI is the interoperable messaging protocol for transferring funds from one bank account to another and is managed by National Payments Corporation of India, which is a non-profit organization promoted by banks, operating under the regulation of the central bank. UPI does not require any prior commercial arrangement between the two entities for executing payments.

Layer 3, data. The model of an account aggregator entails that the financial information will be passed between regulated bodies but on the basis of a signed consent artefact which clearly states the reason, type of information, frequency and expiry, and the consent can be revoked by the consumer. The aggregator company is registered as a non-banking financial company and technologically incapable of reading the data payload it sends out.

Layer 4, credit. The unified lending interface acts as a single application programming interface gateway throughout the life cycle of the loan, eliminating individual integrations between every lender and every data source. The Open Credit Enablement Network segregates customer-side distribution from balance sheet funding, enabling a digital platform to originate a loan which will be booked by a regulated lender.

4.2 Design principles

Table 2. Design Principles Of The Indian Model And Their Functional Purpose


Principle	Implementation	Functional purpose
Protocol not product	The state defines message formats and settlement rules; private firms build interfaces	Avoids a state monopoly over customer experience while retaining standard-setting control
Mandatory interoperability	Every participant must accept instructions from every other participant	Prevents the largest network from capturing the market through exclusivity
Layer unbundling	Identity, payments and data are separately licensed and separately supervised	Contains the operational and political consequences of any single failure
Architectural minimalism	Each layer performs one function and exposes it as an interface	Preserves composability for use cases not anticipated at design time
Utility economics	Base layer operated at or near cost by a non-profit or regulated entity	Removes rent extraction from infrastructure on which all competition depends
Consent as artefact	Signed, machine-readable, time-bound, revocable consent objects	Converts data portability from a legal aspiration into an auditable technical fact

4.3 Institutional separation

This structure parallels that of the technical stack. Payments are handled via a utility run by banks. Credit infrastructure is developed by an innovation subsidiary of the central bank. Aggregators of accounts operate as licensed third parties. The industry association for the data layer received recognition as a self-regulatory organization from the central bank in June 2026. There is no monopolist in charge of the entire stack.

5. Findings: The Payment Layer

Payment Layer carried out 241.62 billion transactions amounting to about 314 trillion rupees in FY 2026 while the volume of transactions was 20 million in FY 2017. The number of banks increased from 44 to 703 during the same period.




Bars show annual transaction volume in billions. Line shows year-on-year growth on the right axis.

Figure 2. Annual Transaction Volume And Year-On-Year Growth Rate, Financial Years 2022 To 2026.

Volumes sourced from the Press Information Bureau release on the tenth anniversary of the payment system and from a written reply in the Lok Sabha. Growth rates computed by the author.

The most important analytical trend is deceleration. The growth dropped from 82 percent in the financial year 2023 to 30 percent in the financial year 2026. The volume keeps increasing, but the growth is slowing down, and it is becoming closer to the growth experienced by a mature utility company. The above statement proves that the hypothesis put forward in section 11 was right about the issue that the future strategy will be based on what comes after the payments.

5.1 Composition of transactions



Person-to-merchant transactions dominate volume while person-to-person transactions dominate value.

Figure 3. Composition Of Transactions By Payment Type, Volume Share Against Value Share.

Source: NPCI transaction segmentation reported by the Press Information Bureau. Person-to-merchant value share computed by the author as the residual of the reported person-to-person value share.

This difference between the two panels provides valuable insights into the system. Person to Merchant accounts for 63% of the volumes but a small percentage of the total value, whereas the Person to Person segment accounts for 71% of the total value of transactions. As shown by the ticket size distribution, the percentage of merchant transactions that are below 500 rupees is 86%.

5.2 Adoption indicators

Table 3. Payment Layer Adoption Indicators

Indicator	Value	Period
Annual transaction volume	241.62 billion	Financial year 2026
Annual transaction value	314 trillion rupees	Financial year 2026
Daily average transactions	660 million	Financial year 2026
Banks live on the system	703	March 2026
Share of domestic digital payments	85 percent	Financial year 2026
Share of global real-time payment volume	Approximately 49 percent	2025
Countries accepting the system	8	2026
Record monthly volume	24.51 billion	August 2026

Source: Press Information Bureau and NPCI data as reported in the business press.

5.3 Economic mechanisms

The following four factors account for the observed adoption. The cost of onboarding was reduced drastically due to the process becoming programmatic, which made low balance accounts economically viable. There is no discount rate for payments to merchants from persons, which means that there are no costs of terminal rent and fees for each transaction which limits card acceptance to organized retail in similar markets. Merchants now have a settled account record which leaves a cash flow trail visible to lenders.

6. Findings: Global Positioning

India held about 49 percent share in the global volume of real-time payment transactions, as reported by the International Monetary Fund in its analysis in June 2025, and reconfirmed by the Press Information Bureau in its April 2026 report.

There are two important caveats to be made regarding this measure, in relation to the comparative discussion in this paper.

First, share in volume is partially dependent on population size, and also on the extremely low average transaction size discussed in Section 5. A payment system in which 86 percent of transactions made to merchants are under 500 rupees would produce a large number of transactions regardless of the amount being transacted. This would result in overstating the economic difference between India and other systems. Secondly, share of the aggregate globally is not an indicator of the quality of design. There is no reason why the Brazilian immediate payment system cannot boast of comparable domestic penetration based on a different institutional structure, where the system was run directly by the central bank, instead of having been established by the banks via a utility.

6.1 Cross-border extension

These have included eight countries, which include the United Arab Emirates, Singapore, France, Bhut-

an, Nepal, Sri Lanka, Mauritius, and Qatar. The Singapore connection is bilateral real-time transfer network connecting Singapore's domestic instant payment scheme, but several of them act as merchant acceptance for the Indian diaspora or tourists traveling from India.

The difference is key. Merchant acceptance by foreign merchants is an extension of the domestic payment scheme to foreign shores. Bilateral real-time transfer networks connecting two domestic instant payment systems are infrastructure interoperability. Only a few of the eight connections are of this nature, implying that cross-border interoperability is more underdeveloped than what is reported.

7. Findings: The Consent And Data Layer

The consent layer exceeded 500 million consent requests being fulfilled in September 2026, five years after the implementation of the framework. This is the part of the Indian model that has the least number of equivalents around the world and, as suggested in this article, the most transferable one.




Figure 4. Cumulative Fulfilled Consent Requests On The Account Aggregator Network, August 2024 To September 2026.

Source: Milestone announcements and the ecosystem dashboard published by the self-regulatory organisation for the framework. Observation intervals are unequal and the series should be read as a trajectory rather than a uniform time series.

7.1 Participation and throughput

Table 4. Consent Layer Participation And Throughput Indicators

Indicator	Value	Period
Financial entities live on the network	1,117	September 2026
Financial information users	1,076	September 2026
Financial information providers	176	September 2026
Licensed account aggregators	17	September 2026

Indicator	Value	Period
Cumulative accounts enabled	2,880 million	September 2026
Cumulative consents raised	1,970 million	September 2026
Monthly data shares	295 million	June 2026
Daily consents processed	Above 700,000	Financial year 2026
Cumulative data fetches	Above 5 billion	Financial year 2026
Financial products and services facilitated	38 million	Financial year 2026

Source: Sahamati ecosystem dashboard and financial year 2026 ecosystem report.

7.2 Credit outcomes attributable to the layer

Total cumulative credit made available via the platform stood at 3.82 trillion rupees by August 2026. As per financial year 2025, around 1.67 trillion rupees have been disbursed via 18.9 million loan accounts using the network. Growth in reported usage has since extended beyond unsecured personal credit to include secured and business credit.

As another sign of deep penetration beyond credit, there have been 6.77 million accounts that have used the platform for derivative trading in financial year 2026, with top brokers reportedly conducting between 70 and 80 percent of income verifications through the platform. The number of personal finance management accounts stands at 59.6 million, growing at a CAGR of 164 percent from financial year 2023 onwards.

7.3 Structural distinction from open banking

There are three aspects that differentiate this model from the open banking models that have been put into practice in the UK and EU.

Firstly, the scope of the coverage not only includes banks but extends to include mutual funds, depositories, insurance firms, pension firms, and tax filings, and this is done under the supervision of four financial sector regulators. Coverage solely based on bank data would inadequately represent the financial condition of individuals in countries with informal employment ratios.

Secondly, the consent manager is separate from the data requester. The consent manager has been delegated to the requesting entity in many international frameworks, creating an apparent conflict of interest.

Thirdly, the consent mechanism is a signed technical object instead of being a contractual one.

7.4 Documented coverage gaps

However, the framework lacks completeness to an extent that renders qualification to the stated percentages. First, jointly held accounts are not included. Second, fixed and recurrent deposits have been enabled by less than half of participating banks. Third, certain types of instruments have not been covered up to the proposal stage only.

8. Findings: Credit Infrastructure And Machine Learning

The fourth level sets the standards for inputs into a credit decision process. This is clearly a much more difficult challenge than setting the standards for a payment process, since the former requires an outcome, a decision, a price, and a long-standing relationship between parties while the latter has on clear outcome.



Figure 5. Credit Decision Flow Across The Layers Of The Indian Stack.

Source: Author's own construction based on regulatory and platform documentation.

8.1 The unified lending interface

The interface acts as a universal entry point for all identity verification, eligibility assessment, application processing, and disbursement by collecting data in real time and doing verification through standardised schema. As of January 2026, there are 64 lenders and 136 data services on the platform. There are 12 different loan journeys which have been made live ranging from agricultural credit, unsecured loans to micro and small enterprises, housing, vehicle, personal, pension, and dairy sector credit.

The data sources become the unique analytical aspect. They consist of digitised land records, satellite images for assessing crop status, goods and service tax filings, milk cooperative payment records, property search record, and credit bureaus. They allow assessing customers who do not have any previous record of borrowing.

8.2 The shift from stock-based to flow-based assessment

Bureau-based underwriting involves the stock of past borrowing and repayments. Cash-flow underwriting involves movement in an account, with seasonal, volatile and counterparty concentrations taken into consideration. In the case of salary-earning clients, the two methods become one. In the case of street vendors, freelance earners and microenterprises, only the latter category of information exists.

This results in the reconstitution of loan portfolios to comprise lower amounts, short-term durations, more frequent origination and pricing which allows adjustment to actual cash flows instead of being static from origination. Management of a leading non-banking financial institution have described

structured consented access to bank statements as underwriting correctly, pricing correctly, greater exposures and portfolio early warning monitoring.

8.3 Applications of machine learning

Table 5. Machine Learning Applications In Indian Financial Services And Their Infrastructure Dependency

Application	Function	Infrastructure dependency
Cash-flow underwriting	Repayment capacity assessment for thin-file borrowers	Consent layer for statement and tax data
Fraud and mule detection	Network-level identification of laundering structures	Transaction graph across institutions
Portfolio monitoring	Early warning from deterioration in cash flow	Continuing rather than one-time consent
Vernacular servicing	Conversational servicing in regional languages	Not infrastructure dependent
Document processing	Automated verification and extraction	Standardised schemas on the credit gateway

It is also noteworthy that the innovation unit within the central bank created a machine learning solution for detecting mule accounts and made it available to all regulated institutions rather than having them develop similar capabilities individually. Deployment reportedly covers 26 banks. Fraud detection is considered an infrastructure service in the same vein as the payment switch.

8.4 Assessment of the evidence base

There is not much independently verified information that machine learning underwriting on such infrastructure results in both increased risk-adjusted returns and greater access. Most claims come from vendors or from lenders talking about their own experience. It is clear how the mechanism works and there is some information about adoption, but the counterfactual has not been established yet. This paper assumes that the access effect is true and unsubstantiated, and describes it as the primary limitation of the current study in Section 12.

9. Governance And Regulatory Architecture

The governing structure may be described as techno-legal. The rules are built into the protocol itself in addition to being specified in the law, thus creating a form of consent which is not only an item on the terms-of-service page but a signed digital document.

Table 6. Governance Instruments Applicable To The Indian Stack

Instrument	Scope of governance	Status as at September 2026
Digital Personal Data Protection Act 2023	Lawful processing, consent, obligations of data fiduciaries, consent managers	Enacted, with rules and phased implementation in progress
Guidelines on Digital Lending	Direct disbursement to borrower, cost disclosure, conduct of lending service providers	In force

Instrument	Scope of governance	Status as at September 2026
Master Direction on Account Aggregators	Licensing, consent artefact specification, data-blind design	In force, with self-regulatory organisation recognised June 2026
Framework for Responsible and Ethical Enablement of Artificial Intelligence	Adoption of artificial intelligence by regulated entities	Committee report issued August 2025 with 26 recommendations
Discussion paper on digital payment fraud	Transaction friction, mule account controls, customer safeguards	Issued April 2026, consultation closed May 2026

9.1 The artificial intelligence framework

The framework has been formulated by a task force formed in December 2024 and headed by Professor Pushpak Bhattacharyya from Indian Institute of Technology Bombay, and delivered its report in August 2025. It consists of seven guiding principles and six strategic pillars on infrastructure, policy, capacity, governance, protection and assurance and contains 26 recommendations and annexed documents, comprising proposed modifications to master directions, sample board policy template and incident reporting guideline.

It can be distinguished from similar international tools by three characteristics.

First, it regards responsible innovation as a higher priority for regulation than precaution, reversing the default attitude of artificial intelligence governance frameworks.

Second, it suggests active involvement of the regulator in developing shared infrastructure, including curated data sets, innovation sandbox and domestic development of artificial intelligence solutions for the financial sector. The rationale is that keeping the capability in the hands of the largest players increases systemic risk.

Third, the disclosure of artificial intelligence applications and the final decision making power of humans as basic principles have been proposed.

This structure incorporates the principle of proportionality; whereby less onerous measures will apply to low-risk cases while supervision will apply in high-risk cases. The surveys done by the committee reveal that less than one-third of respondents, particularly large banks, had oversight at the board level with respect to artificial intelligence, which means that there was no policy framework.

9.2 An unresolved regulatory boundary

The data protection law creates a class of consent managers, whereas the central bank's regime creates a class of account aggregators. The two classes are similar but not exactly identical. Resolving a data protection law and the preexisting sectoral consent infrastructure within its scope continues to be an unresolved task, and the industry has made official submissions on this question.

This is not unique to the Indian experience. Any jurisdiction that develops sectoral consent infrastructure prior to adopting general data protection laws can expect the boundary issue, and should ideally plan for such an outcome rather than facing it once both measures are in place.

10. Critical Assessment Of Failure Modes

The infrastructure that makes the transaction cheaper for everyone will also make it cheaper for all parties involved regardless of whether the designers intended to accommodate them. Every weakness

outlined here is an inevitable result of an advantage of the infrastructure design; this is why the following section is analyzed as a critical one.

10.1 Fraud enabled by low-friction design

Immediate, irreversible and low-friction transfers are advantageous for authorised push payments fraud wherein the victim is manipulated to make the transfer. There is no method of authentication that will detect an authorised payment made by the account holder.




Figure 6. Amounts Involved In Reported Bank Fraud Cases, Financial Years 2023 To 2025.

Source: Reserve Bank of India annual reports as reported in the business press. The series covers all reported bank fraud rather than digital payment fraud alone, and reporting thresholds changed during the period, so the series should be read as indicative of direction rather than as a consistent measure.

Mule accounts make up the laundering layer in this process. In one cyber-intelligence analysis, there were 524,121 suspected mule accounts and virtual payment accounts in just March 2026, where payments banks made up 41 percent of the suspicious activity while merchant accounts accounted for 11 percent. Onboarding friction is certainly something diagnostic here since the same reduction in onboarding friction that enabled financial inclusion to become commercially feasible reduced the cost of opening a mule account.

What has been outlined in the discussion paper published in April 2026 is the regulatory response to such practices, which includes the selective introduction of friction, where there will be a delay on certain transfers over 10,000 rupees, an annual credit ceiling of about 2.5 million rupees for accounts that do not have any enhanced due diligence, and a customer-managed facility to block digital payment channels.

10.2 Authentication exclusion

Biometric identification is especially prone to failure among blue-collar workers who have deteriorated fingerprints and older people and those suffering from chronic illnesses. When the process of identification is the door to a monetary benefit payment or a withdrawal of cash, then a failure to identify becomes a failure to deliver on the actual entitlement.

The fundamental problem is that while the digital process may become more resourced and streamlined, the analog channel will always be the least resourced and will continue to be utilized by people who are the least able to challenge it.

10.3 Predatory digital lending

Before the implementation of the digital lending guidelines, unlicensed applications had been using stolen contacts for debt collection through coercion, obscure pricing and hidden lending mechanisms. Cases of deaths have been recorded. The guidelines, which require direct disbursement, transparency on costs, and accountability for the lenders, were meant to address damages that had already occurred.

The transferable lesson is one of sequencing. The conduct regulations came after the infrastructure; not with it. Potential adopters should be aware that bad actors can get to the new rails before the conduct rulebook.

10.4 Credit quality and over-indebtedness

Frictionless origination is frictionless in both directions. Small ticket unsecured lending has been through an obvious growth cycle with stress, then supervision of risk weighting and concerns with concentration of lenders. Immediate decision removes the period during which the borrower could rethink.

10.5 Privacy, function creep and concentration

Four common themes appear in the literature and are captured here without judgment.

The first is the quality of consent. While the consent artifact can be legally valid, the borrower who is being asked to provide it does not fully comprehend what he or she is agreeing to.

The second is function creep. Systems designed for one function tend to accrue more functions, and legal constraints regarding secondary usage are more difficult to enforce than technical constraints.

Third is the capacity for state surveillance. Identity and transaction infrastructures on the scale of the entire population are long-lasting capabilities, which outlive the government that designs them.

Fourth is concentration at the application layer. For a very long time, two applications accounted for most of the payment volume, with the share of both falling below 80 percent only for the first time in May 2026. Several proposals for maximum market shares have been floated and then put off again. The rail is open, but the market above it is not evenly fought for.

10.6 The unresolved funding model

There is no rate for person-to-merchant payments, which means nobody benefits directly from the main purpose of the system. The expenses are covered through the banking system and subsidized in part by the government occasionally. Such a vital infrastructure running without a clear way of funding is a form of vulnerability and not a successful example of zero pricing.

11. International Diffusion

India has transformed its domestic infrastructure initiative into a tool of foreign policy. More than 23 to 25 countries have inked memoranda of understanding or cooperation on implementing variations of the model, while the G20 presidency for 2023 yielded a multilateral framework supporting the fundamental principles.

The exported initiative is purposefully designed to be a non-product format. Open-source blueprints are provided to partner governments, who assemble and own them.

Table 7. Channels Of International Diffusion

Channel	Component exported	Documented reach
Modular Open Source Identity Platform	Open-source foundational identity platform developed at a Bangalore institute	Deployment or pilot in Morocco, the Philippines, Ethiopia, Sri Lanka, Uganda, Sierra Leone, Guinea, Togo and Burkina Faso; approximately 11 countries with adoption and around 90 million registrations reported
NPCI International Payments Limited	Payment protocol adaptation, domestic card scheme construction, acceptance corridors	8 countries accepting the Indian payment system; domestic system builds for partner central banks
Policy and design assistance	Consent framework guidance, document wallet architecture, implementation roadmaps	Centre for Digital Public Infrastructure engaged with approximately 65 government departments internationally
Adjacent open-source programmes	Beneficiary and benefit-transfer platforms originating from the same institutional base	Adoption by five countries within two years

It is imperative to understand the importance of the demand context in this scenario. According to the 2026 programme of the World Bank, about 2.9 billion people who are mostly found in Sub-Saharan Africa and South Asia do not have digital identification.

11.1 The transfer of risk alongside architecture

Importing the architecture means importing its riskiness. The identification system of India is an object of litigation and criticism within the country on account of exclusion and surveillance, while being admired abroad, and this contrast some scholars have called a paradox.

This characterization is correct. The system decreased transaction costs significantly while producing exclusions at its margins. For a state introducing the system into practice, it is necessary to analyze the exception management system, sequencing of regulation of the conduct, and questions of funding at least as attentively as the statistics of adoption of the system. This is the message of the paper.

12. Discussion: A Transferability Framework

The key analysis in this paper rests on the proposition that the success of India is due to the design choice rather than technological advantage. The government provided light, open and interoperable protocols and let products compete in the marketplace. Volume of transactions is the result of this choice rather than its accomplishment.

The attributes of the model should be categorized based on their dependencies for the purpose of policy discussion. There are three types of categories.

Table 8. Transferability Classification Of Model Features

Class	Features	Dependency	Transferability
Codified design	Protocol over	Rules and standards	High. These are

Class	Features	Dependency	Transferability
principles	product, mandatory interoperability, layer unbundling, minimalism, consent as artefact, utility economics	only	design decisions independent of national conditions
Institutional form	Bank-promoted non-profit operator, innovation arm within the central bank, recognised self-regulatory body for the data layer	Political durability across administrations and multi-year institution building	Moderate. Equivalents can be constructed but require sustained political cover
Scale-dependent economics	Zero-price merchant acceptance, cost recovery through volume, a domestic technology sector able to build immediately	Population size and market structure	Low. Smaller markets must design an explicit funding model from the outset

12.1 Sequencing

The sequence of the construction of the different layers is not random since each layer's economic structure depends on that of the layer below it.




Figure 7. Recommended Sequencing For Jurisdictions Adopting The Model.

Without developing a consent infrastructure before implementing the machine learning underwriting process, the jurisdiction will realize that all that the models are capable of observing are those data sets that the implementing organization already possesses. This only ensures further reproduction of current biases more rapidly and at reduced costs, and not the intended result.

12.2 Ten decisions for policymakers

1. Decide on ownership of the payment switch prior to the construction phase. A private incumbent maximizes the rail for its own distribution; a government department usually invests too little in reliability. The bank-led or regulator-led utility has demonstrated that this is the workable structure, and the decision is virtually irreversible.
2. Require interoperability while the market is still small. The political cost of doing so increases every month that incumbents mature.
3. Create, finance, staff, and publicly monitor the exception path. Disaggregated authentication failure rates by district, age, and occupation should be a reporting supervisory metric.
4. Bring about conduct regulation prior to scale of the rails in terms of cost disclosure, disbursal, recovery conduct, and grievance redress.
5. Decide on and publicize the funding model. Zero pricing is a policy decision, and comes at a price.
6. Turn consent into a machine-readable artefact with purpose, scope, expiration and revocation.
7. Grow the data layer from just banking to tax, securities, insurance and pension data.
8. Create the regulatory technology in parallel with market infrastructure, including shared fraud detection technology.
9. Publicize reasonable expectations for artificial intelligence early, including disclosure, human override, model documentation, bias testing and incident reporting.
10. Avoid combining the layers into one government platform. It's separation that prevents political capture, technical blast radius and vendor lock-in.

12.3 Implications for financial institutions

Public infrastructure takes away some competitive advantages and also adds new ones. Accepting payments, having exclusive data storage capabilities and proprietary onboarding process becomes ineffective in creating barriers once all participants are able to transact with each other, and customers can transfer their data whenever they want.

Those aspects which still work are the cost of funds, proprietary models using data regarding loan repayments which cannot be transferred, collecting loans in widely dispersed portfolio of smaller amounts, and trust of the institution while making important financial decisions. This partly explains why the infrastructure era in India benefited big deposit-taking institutions rather than those predicted by the commentary on financial technology disruption.

For the institutions entering the markets where this kind of infrastructure is being implemented, the period of gaining competitive advantage is the one which follows implementation and precedes market maturity. Business models earning margin from friction which the infrastructure aims at eliminating, such as data aggregation, screen scraping, payment routing and verification, have definite expiration date.

13. Limitations Of The Study

There are five factors limiting the results above.

First, the research is done using secondary information only. There has been no use of surveys of lenders, borrowers, or of those who were not included in the loaning process. Any claims related to the results of underwriting are based on the published statements of the institutions themselves on their performance.

Second, some of the metrics come from the self-regulatory organization for the consent layer of the industry, which clearly has its own interest in publishing positive adoption figures. These have been cited as such. Whenever there has been any minor difference between the government and industry figures of the same metric, both have been reported.

Third, the counterfactual to the central claim concerning access to credit is not set out. In this paper, evidence is provided that machine learning underwriting using consented data is being rolled out and how this could potentially be used to provide access to credit to borrowers who would not otherwise receive credit is outlined. Evidence is not provided to show that access has increased compared to the counterfactual, nor is there any evidence that new borrowers accessing credit are being charged appropriately.

Fourthly, the evidence on fraud and exclusion is inconsistent. The data on fraud series is biased by changes in reporting thresholds throughout the period covered, and the mule accounts estimate is based on commercial intelligence firms with undisclosed techniques. The data on authentication failure rates are unavailable, hence the analysis of the issue presented in Section 10 is based on cases.

Lastly, the transferability model presented in Section 12 was developed in an inductive fashion using the case study of a single country. This model has not yet been tested against the experiences of other countries that have incorporated elements of the model.

14. Conclusion

This paper reviewed the shift in India from digital public infrastructure to intelligent finance and considered how much of that experience can be replicated elsewhere.

The key point is that the determining factor was architecture. India did not use better technology. Rather, it made an architectural decision that other jurisdictions have not made, which is to maintain the base layer thin, open and interoperable and allow competition on top of it. The transaction numbers cited in Section 5 result from this decision, not the accomplishment itself.

Second, maturation is relevant to the case. Transaction growth is slowing down due to the base growth, going from 82 percent during financial year 2023 to 30 percent during financial year 2026. This makes the strategic importance of India's case during the next decade in the consent and credit layers rather than in transaction expansion.

The third finding is that the second transition is substantively more difficult than the first. The transportation of money is a well-understood engineering problem. The transportation of consented, explainable, and contestable decisions about who gets credit and on what terms is not. India is undertaking this experiment on a scale unparalleled in any other jurisdiction, under the oversight of a regulatory regime that has outlined an explicit framework for it. The results are still unknown. Machine learning algorithms applied to consented data may lead to lending where lending has never been available before, or they may result in a faster and more documented process of exclusion, in which price discrimination is less visible to the affected borrower.

The fourth finding is that failure modes are inherent properties of the architecture rather than accidents. Low-friction customer onboarding facilitated both financial inclusion and the development of mule

accounts. Instantaneous clearing led to both merchant acceptance and authorised push payment scams. Frictionless loan origination led to both thin-file loans and overindebtedness. When adopting this architecture, countries should anticipate the secondary consequences of each architectural choice when making it.

14.1 Recommendations

For policymakers, the ten decisions presented in Section 12.2 represent the actual recommendation of this paper, based on the sequence shown in Figure 7. The three most important decisions are ownership of the switch, interoperability required before incumbents become consolidated and conduct regulation introduced before scale is achieved.

For regulators, the Indian policy framework on responsible enablement of artificial intelligence serves as a model that recognises innovation as a responsibility, advocates for shared infrastructure to avoid the centralisation of capabilities, and considers human override as a basic principle, not a control.

For banks and other financial institutions, the key takeaway is that any competitive advantages built on infrastructure, exclusive data or onboarding technology are not going to survive the change, whereas competitive advantages based on funding costs, proprietary outcome data, collections and trust will.

14.2 Directions for further research

Three paths of inquiry emerge immediately. First is the issue of panel analysis at the borrower level to determine whether the use of consented data for underwriting expands or reprices access. Second is the assessment of adoption countries against the proposed transferability framework. Third is authentication exclusion at population level, which has no series whatsoever in published form.

The framework is the tractable portion of the problem. What is built on top of that remains a matter of policy.

References

1. Biometric Update. (2026, February 11). MOSIP Connect 2026 focuses on sustainability and practical uses of digital identity. Retrieved September 18, 2026
2. Biometric Update. (2026, August). India's CSM Technologies partners with MOSIP to bridge the gap between GovTech and Digital Identity. Retrieved September 18, 2026
3. Business Standard. (2026, September 1). UPI hits new high in August, logs 24.51 billion transactions; worth 29.82 trillion rupees. Retrieved September 18, 2026
4. Chambers and Partners. (2025). RBI Committee Report on Responsible AI in the Financial Sector: FREE-AI Framework. Retrieved September 18, 2026,
5. Dvara Research. (2025). Summary of the RBI FREE-AI committee report. Retrieved September 18, 2026, from <https://dvararesearch.com/summary-of-the-rbi-free-ai-committee-report/>
6. IANS. (2026, August 1). UPI transactions grow 22 per cent in July, value rises to
7. Impact and Policy Research Institute. IMPRI. (2026). Unified Lending Interface: The Digital Push By The RBI For Faster Credit Access. Retrieved on September 18, 2026, from <https://www.impriindia.com/insights/policy-update/unified-lending-interface-rbis-digital-push-for-faster-credit-access-2025/MediaNama>.
8. Multibagg. (2026, April 16). RBI digital fraud strategy: One-hour UPI delay in 2026. Retrieved September 18, 2026, from <https://www.multibagg.ai/market-pulse/articles/rbi-digital-fraud-upi-delay-cmo112jnnujxxny0jzznl90j9>

9. RBI Innovation Hub. (2026). Unified Lending Interface. Retrieved September 18, 2026, from <https://rbihub.in/projects/unified-lending-interface>
10. Reserve Bank of India. (2025, August 13). Report of the committee on the framework for responsible and ethical enablement of artificial intelligence. Mumbai: Reserve Bank of India.
11. Reserve Bank of India. (2023). Guidelines on digital lending. Mumbai: Reserve
12. Sahamati. (2024, August 20). India's Account Aggregator framework crosses 100 million consents in three years. Retrieved September 18, 2026, from <https://sahamati.org.in/media-article/indias-account-aggregator-framework-crosses-100-million-consents-in-three-years/>
13. Sahamati. (2025). The credit reimagined: Account Aggregator impact report, H2 FY25. Retrieved September 18, 2026, from <https://sahamati.org.in/the-credit-reimagined-account-aggregator-aa-impact-report-h2-fy25/>
14. Sahamati. (2026, August 27). Account Aggregator powers 3.82 lakh crore rupees in credit, with secured lending emerging