

RESEARCH REPORT

India's digital payments and the small-enterprise productivity test

India's payment expansion creates records and can reduce transaction costs, but adoption, registration, credit access and productivity remain separate outcomes. Official enterprise evidence supports a staged test of the mechanisms connecting them.

Information cutoff: 2026-10-27

<https://aesenate.org/research/india-digital-payments-and-small-enterprise-productivity/>

Four outcomes that must be measured separately

Digital payments can make a transaction faster and easier to document. They do not, by themselves, establish that a business has registered with the relevant authority, obtained affordable credit or produced more value with its resources. Those outcomes involve different decisions, institutions and constraints. A serious assessment of India's digital infrastructure should test their connections rather than treat them as successive names for the same event.

The distinction matters because payment adoption can be commercially useful even when no measured productivity gain follows immediately. A merchant may reduce cash-handling effort, serve customers who prefer digital payment or improve reconciliation. Another may accept digital transfers while keeping incomplete business accounts. A lender may obtain more information but still reject an application because repayment capacity is inadequate. Each result provides information about a different part of the mechanism.

This report studies the chain from payment use to enterprise records, financing and productive performance. Official payment statistics establish system activity; enterprise-survey evidence identifies business characteristics; RBI material describes lending-data infrastructure. None of those sources alone estimates a causal national productivity effect. The report develops an evaluation framework, statistical bounds and transparent numerical accounts to distinguish the plausible channels.

The central finding is that digital infrastructure can lower specific information and transaction costs, but its productive effect depends on complementary business capabilities, lender behavior and demand. Payment adoption is an intermediate outcome. Registration is a separate administrative outcome. Better financing is a

contractual and economic outcome. Productivity requires a consistent measure of real value added and inputs, together with an explanation of what changed.

The enterprise evidence does not collapse those margins

MoSPI’s ASUSE release for October 2023-September 2024 reports internet use for entrepreneurial purposes at 26.7 percent of unincorporated nonagricultural establishments, up from 21.1 percent in the preceding survey. Registration with at least one act or authority was 37.2 percent, compared with 36.8 percent. These are distinct survey indicators, not a count of UPI merchants or a comprehensive definition of formal employment. [MoSPI, ASUSE 2023-24 results](#).

ASUSE indicator	2022-23	2023-24	Change
Internet used for entrepreneurial purposes	21.1%	26.7%	5.6 percentage points
Registered with at least one act or authority	36.8%	37.2%	0.4 percentage points

The different movements demonstrate why a digital-use indicator cannot be substituted for registration. They do not establish that internet adoption failed to cause formalization, because the table does not link the same establishments or identify their counterfactual behavior. Entry, exit, composition and survey measurement can also affect the aggregate shares.

The survey’s coverage matters. It concerns unincorporated nonagricultural manufacturing, trade and other services, with specified ownership forms. Its release describes area-based sampling and substantial use of oral monthly-reference inquiries. Those features should be considered when connecting the results to administrative payment records. A digital account and a sampled establishment are not automatically the same statistical unit.

MoSPI’s variable metadata asks whether the establishment used the internet during the last 365 days. That question is wider than accepting a particular payment instrument and different from a measure of transaction intensity. It should not be relabeled as an enterprise’s digital-payment share. [MoSPI microdata metadata, internet-use variable](#).

Payment-system scale answers a narrower question

NPCI’s published table records UPI transaction volume of 24,508.96 million and value of 2,982,355.95 crore rupees for August 2026. The table identifies monthly transaction units and participating banks. Those observations establish the scale of transactions processed in that period; they do not count unique merchants or measure enterprise value added. [NPCI, UPI product statistics](#).

A single customer can make many payments, and a single business can use several accounts or interfaces. Person-to-person transfers and merchant-related transfers also answer different questions. An aggregate system

total should not be divided by the number of establishments from another survey and labeled average sales without identifying the transactions and matching the population.

Gross payment value is particularly unsuitable as a direct GDP measure. Funds can move more than once while financing one chain of activity. A payment can settle an intermediate purchase, transfer money between people or repay an obligation. National accounts require a production and expenditure classification, not the total of financial settlements.

A higher transaction count can also arise from splitting a purchase into smaller payments or substituting digital transactions for cash. Those changes can improve convenience without increasing physical output. Average transaction value can change because of the mix of users and purposes. It is not a price index for goods or services purchased through the system.

For enterprise research, the useful next step is a matched merchant account: verified business transactions, refunds, internal transfers and the corresponding sales or invoices. The account should preserve customer and merchant privacy while enabling aggregate analysis. System statistics provide the environment in which adoption occurs; the business account identifies the mechanism relevant to financing and performance.

The marginal shares do not reveal the joint distribution

The ASUSE table gives the share using the internet and the share registered, but not their overlap in the presented exhibit. From those two marginals alone, the proportion both internet-using and registered can lie anywhere between zero and 26.7 percent. The lower bound is the greater of zero and $26.7 + 37.2 - 100$; the upper bound is the smaller share.

These are logical identification bounds, not confidence intervals. They show how little can be concluded about the relationship from two aggregate percentages. A joint table could narrow the descriptive question immediately. Establishing a causal effect would still require a design addressing why some establishments adopt and others do not.

A hypothetical population makes the distinction tangible. Suppose sixty of one hundred businesses accept digital payments, forty are registered and thirty belong to both groups. Registration among digital adopters is fifty percent, while registration among nonadopters is twenty-five percent. That association can reflect a payment effect, prior business capability, size, customer demand or another common influence.

The example is compatible with several causal stories. Registration may make digital adoption easier; adoption may encourage registration; or larger, better-managed firms may do both. A cross-sectional association cannot distinguish the directions. Treating the conditional rate as an impact estimate would convert a descriptive comparison into an unsupported causal claim.

A panel with documented adoption and registration dates would improve the evidence, but it would not eliminate selection automatically. Firms anticipating growth can adopt before sales rise. A registration campaign can coincide with wider payment access. The analysis should retain pre-existing trends and simultaneous changes rather than assume that temporal ordering alone identifies the effect.

Formalization is a set of institutional relationships

A business can hold a payment account, register an enterprise, file tax returns, maintain accounts and employ workers under particular arrangements. These relationships can overlap without being identical. The ASUSE indicator of registration with at least one authority should retain that specific meaning. It does not establish full compliance with every applicable requirement or a change in the employment status of all workers.

The economic benefits and costs also differ. Registration may improve access to contracts, finance or services while requiring administration and reporting. A digital record can lower bookkeeping effort but expose previously unrecorded transactions. The firm's response depends on incentives, capability and the rules relevant to its activity. No single transaction count can identify those tradeoffs.

An evaluation should therefore choose its outcome precisely. If the question concerns enterprise registration, observe registration and persistence. If it concerns tax reporting, examine the appropriate filings and reconciliations. If it concerns worker protections or benefits, use employment and benefit evidence. Calling every change formalization makes the claim harder to test and can conceal a mixed result.

Administrative registries also differ from surveys. A registration can remain active after operations cease, and one economic enterprise can have several registrations or establishments. Conversely, an active establishment may be represented only indirectly in a registry. Changes in registry totals should be reconciled with active production rather than interpreted automatically as new businesses or jobs.

For public policy, this distinction supports a more useful target: reduce avoidable administrative and information costs while preserving the purpose of the relevant institution. A successful registration initiative should be assessed through actual participation and outcomes, not merely through the number of payment accounts associated with the same population.

Better records can reduce one component of the lending cost

A lender assessing a small firm needs to establish identity, income, obligations, repayment capacity and the reliability of the underlying information. Some costs are largely fixed per application. They can make a small loan expensive to originate even when the borrower's operating activity is viable. Standardized data access can address part of that cost, but it does not remove credit risk or the need for funding and capital.

RBI's 2024-25 annual report describes ULI as connecting lenders and data providers through a standardized interface. It records 44 lenders, more than sixty data services and twelve loan journeys as of March 31, 2025. Those are infrastructure participation measures, not estimates of loans caused by the interface or gains in borrower productivity. [RBI, annual report, paragraph VI.46](#).

A controlled loan account shows the potential cost channel. Assume a one-year loan of one hundred thousand rupees, funding cost of eight thousand, expected credit loss of two thousand and assessment and servicing

cost of three thousand. Ignore capital charges, taxes and profit for this exercise. Revenue of thirteen thousand is required to cover those stipulated costs.

If better data and processing reduce assessment and servicing cost to one thousand with all other inputs unchanged, required revenue falls to eleven thousand. The two-thousand reduction equals two percentage points of principal over the year. Whether the saving reaches the borrower depends on competition, pricing and other costs. This is an assumed cost reduction, not an observed Indian rate effect.

The expected-loss amount is separate from the information-processing cost. More detailed records can improve risk assessment, but they can also reveal that a borrower is less creditworthy than previously thought. Some applicants may receive better terms and others may be rejected. A lower rejection rate is not the only possible sign of improved underwriting; repayment outcomes, access and appropriate pricing should be considered together.

A transaction history is not a complete cash-flow statement

Incoming transfers can include business sales, household transfers, loan proceeds, refunds and movements among accounts. A lender needs to distinguish those sources before treating receipts as recurring operating revenue. Outgoing transfers likewise need classification. A payment history containing many observations can still omit cash transactions, supplier credit or liabilities held at another institution.

Suppose a hypothetical account receives one hundred thousand during a month. Sixty thousand represents customer receipts, twenty thousand is a transfer from the owner's other account, ten thousand is a new loan and ten thousand is a refund. Treating the full amount as sales would overstate customer receipts by forty thousand. More digital data would not solve the error if the classification remained wrong.

Even correctly identified sales do not establish repayment capacity. Intermediate costs, household withdrawals, taxes, investment and existing debt service affect available cash. A merchant with high turnover and a low margin can have less capacity than a smaller business with stable income. The appropriate underwriting account connects flows, balances and obligations rather than relying on transaction volume alone.

Seasonality requires another adjustment. A festival, harvest cycle or temporary contract can produce a strong month followed by weaker receipts. Extrapolating that month across a year can overstate sustainable income. A useful digital record should cover the cycle relevant to the business and preserve exceptional events rather than smooth them away without explanation.

Verification and contestability matter as well. Duplicate entries, mistaken counterparty labels or a temporary outage can affect a score. The economic value of data access depends on reliable information and a way to correct material errors. A highly automated decision can be fast while still being wrong; speed and decision quality are separate performance measures.

Consent and interoperability are part of the economic mechanism

The government's May 2026 explanation distinguishes payment infrastructure, lending-data interfaces and the Account Aggregator framework. It describes the latter as consent-based sharing of financial information across institutions. This establishes a stated design principle and institutional distinction; it does not demonstrate that every user experiences equally informed consent or that every lender uses data identically. [Government of India, financial-inclusion backgrounder, Account Aggregator section](#).

The distinction is economically important. A payment network moves funds. A data-sharing arrangement can help an authorized service provider retrieve information. A lender makes a credit decision and bears or transfers the resulting exposure. Interoperability can lower the cost of connecting those functions while leaving their responsibilities separate.

A borrower benefits from data portability when it enables competing offers rather than exclusive dependence on one provider. The assessment should examine whether comparable information can reach willing lenders, how total prices are presented and whether the borrower can understand and refuse particular uses. More connected accounts alone does not establish effective competition.

Data use also has costs beyond processing. Security failures, misuse or inaccurate classification can impose losses and reduce willingness to participate. Those risks should be assessed through actual incidents, controls and remedies, without assuming that all digital use is unsafe or that an interoperable system removes every vulnerability. Trust is an operating input to sustained adoption.

For research, data minimization supports both privacy and interpretability. Collect the information necessary to identify the mechanism, document authorized access and use appropriately aggregated outputs. The report does not prescribe a transaction-specific legal compliance procedure. Its methodological point is that the ability to observe a relationship does not justify unrestricted collection or an assumption that every available variable improves the decision.

Payment speed can release working capital without creating sales

Faster availability of receipts can reduce the funds tied up between delivery and settlement. That is a balance-sheet improvement even if the business sells the same quantity. Its productive effect depends on whether the released cash relaxes a binding constraint, reduces borrowing or is held as liquidity.

Consider a hypothetical merchant with daily eligible sales of two thousand rupees and an average ten-day interval before receipts become usable. Under a deliberately simple steady-flow assumption, twenty thousand is tied up in that interval. Reducing the interval to one day leaves two thousand tied up and releases eighteen thousand. This is an illustrative cash-conversion calculation, not a description of UPI's actual settlement rules or a claim about all merchant receivables.

If the merchant had borrowed the eighteen thousand at an annual fifteen-percent rate, replacing that finance could save two thousand seven hundred over a full year under simple interest. If no borrowing was used, the benefit could instead be greater liquidity or earlier purchases. The same released amount need not generate the same interest saving or spending response across firms.

Faster payment does not resolve a customer's refusal or inability to pay an invoice. It can improve settlement after authorization while leaving commercial credit terms and disputes unchanged. An evaluation should separate transaction-processing time, invoice-credit time and overdue debt. Combining them can attribute a change in customer behavior to the payment rail incorrectly.

Nor should the release of working capital be counted as the same amount of new GDP. It is a stock adjustment. Additional production requires a subsequent use of resources and demand for the output. The cash account identifies a potential enabling condition; the production account measures whether that condition was converted into more or better activity.

Productivity requires output and input evidence

A firm's digital receipts can rise while total sales remain unchanged because customers substitute away from cash. Its recorded turnover can rise because previously unrecorded transactions become visible. Prices can rise while quantities remain unchanged. Each mechanism differs from producing more real value added per hour or per unit of combined input.

Take a hypothetical establishment with sales of ten thousand, intermediate inputs of six thousand and eighty hours of work. Value added is four thousand and value added per hour is fifty at the stipulated prices. If the digital share of receipts rises from fifty to eighty percent with the underlying activity unchanged, digital receipts rise from five thousand to eight thousand. Productivity remains fifty.

Now assume administrative improvements reduce hours from eighty to seventy-six while output and intermediate inputs remain unchanged. Value added per hour rises to approximately 52.63, a gain of about 5.26 percent. The gain concerns measured labor time. It should include the owner's administrative hours and any new outsourced service costs; otherwise the apparent saving could simply move work outside the measured boundary.

A broader productivity assessment would separate capital services, skills, operating scale and intermediate inputs. Buying a device or subscribing to software can improve capability while adding costs. The correct comparison is net of the resources used and adjusted for output quality. Revenue per worker is useful descriptive evidence but cannot isolate a technological contribution by itself.

The relevant horizon can be longer than the adoption month. Learning, customer uptake and bookkeeping changes can take time, while initial disruption can absorb resources. A credible study should follow outcomes over an adequate interval without assuming that every weak early result will become a later gain. The timing is an empirical question, not a reason to exempt the intervention from evaluation.

Selection can produce a convincing but misleading adoption premium

Larger or better-managed firms may adopt digital tools earlier because they have more customers, better connectivity and stronger administrative capability. They may also be more productive before adoption. Comparing their current productivity with nonadopters can therefore attribute a pre-existing difference to the payment instrument.

Customer demand creates another selection channel. A merchant serving buyers who prefer digital payments has an incentive to adopt and may already face a growing market. Sales growth after adoption could reflect the customer base rather than a reduction in transaction costs. Conversely, a struggling business may adopt specifically to find new customers, making adopters initially weaker than comparison firms.

Credit introduces further feedback. A lender can require particular records before approving a loan, while the loan enables investment in digital operations. The observed sequence of records, finance and growth can run in several directions. An analysis should document dates and conditions rather than assume a one-way chain from payments to productivity.

A useful design would combine establishment characteristics before adoption, changes in local access and the actual timing of use. A documented phased introduction can offer a comparison, but locations receiving earlier access may differ in infrastructure or demand. Those differences need testing. An instrument or rollout variable is not valid merely because it is geographical or determined by an institution.

The outcome should also include nonadoption and exit. Studying only successful continuing users omits firms for which the tool was unsuitable, inaccessible or insufficient. A policy evaluation needs a defined eligible population and follow-up of the initial cohort, including businesses that close or switch providers. Otherwise survivorship can create an apparent performance gain unrelated to the average effect.

Existing credit research narrows one part of the question

A study by Shashwat Alok, Pulak Ghosh, Nirupama Kulkarni and Manju Puri examines UPI using consumer-loan data. Its September 2025 working-paper abstract reports deeper formal credit access, differences between bank and fintech lending, and complementarity with internet and bank-account access. The authors report no evidence of increased defaults in their analysis. These are attributed findings from a working paper, not an AES replication or an estimate of small-enterprise productivity. [Alok and coauthors, NBER Working Paper 33259](#).

This evidence is relevant because it directly studies a credit channel rather than inferring one from transaction growth. It also creates a clear boundary for the present report. Consumer borrowing is not the same object as an establishment's working-capital facility, and additional credit is not a direct measure of real productive output. Extending a credit finding to an enterprise-productivity claim requires another link and a suitable dataset.

The reported complementarity reinforces a staged interpretation. Infrastructure may be more effective where users also have the means to connect, hold accounts and supply usable information. A national platform can be technically available while the local economic conditions governing its use differ. That heterogeneity should be examined rather than converted into one average effect assumed to apply to every business.

A research synthesis should retain the paper's population, outcome and identification strategy before comparing its result with survey associations. This report inspected the primary abstract and does not claim to have reproduced the authors' proprietary data or audited their complete empirical design. The responsible inference is that serious evidence exists for a particular credit-access mechanism, while the enterprise-registration and productivity questions remain distinct tests.

Platform economics and the location of costs

A low charge at the point of payment does not establish that the system has no operating cost. Banks, infrastructure providers, application providers, merchants and the public sector can bear different parts of the cost. Revenue can come from related services or explicit support. Evaluating efficiency requires those resources to be identified rather than treating a zero visible charge as a zero social cost.

Interoperability can improve competition by allowing users of different providers to transact. It can also coexist with concentration in customer-facing applications or data access. A useful competition assessment should distinguish the shared payment rail, the bank account, the application and the lending relationship. Market shares at one layer do not automatically describe control at every layer.

The same distinction matters for innovation. A common interface can reduce the cost of entry for a new service, but distribution, trust and access to customers may remain difficult. A large existing provider can offer useful scale while also shaping which products users see. Observing the number of applications is therefore insufficient to measure competitive choice; switching costs, effective access and terms matter.

Support measures should be evaluated against that layered account. A subsidy may expand acceptance initially, but its durable effect depends on merchant use after support changes and on the funding of maintenance and resilience. The analysis should record costs borne by each participant and changes in related-service prices. Otherwise a cost reduction at one point can be mistaken for an equal reduction across the system.

Firm gains and aggregate productivity need another reconciliation

An adopting merchant can gain sales by attracting customers from a nearby nonadopter. That is an important private benefit, but aggregate output may rise by less if the gain is largely a reallocation of existing demand. A national productivity assessment should distinguish better matching and lower resource use from a redistribution of sales among firms.

Reallocation can itself improve productivity if more efficient firms expand. Establishing that channel requires comparable production and input measures across the affected firms, including those losing customers or exiting. Studying only adopters' revenue omits the offset. A local evaluation should therefore consider the market in which customers substitute among suppliers, rather than assume that each firm's gain is additional national demand.

Network effects can create benefits outside the treated firm as well. A customer may make fewer cash-withdrawal trips or a supplier may reconcile receipts more easily. Those are potential resource savings, but they require measurement and should not be assigned automatically as a fixed multiplier. The boundary of the evaluation determines which external effects are included.

A complete national estimate would combine direct firm effects, market reallocation, system operating costs and relevant external benefits. It should also account for price changes and the possibility that saved time becomes leisure rather than additional market output. The report's enterprise calculations isolate useful mechanisms; they do not replace that aggregate reconciliation. This is why a credible digital-economy claim must identify both the unit experiencing the gain and the unit over which the claimed total is measured.

A practical design that separates the channels

The first layer measures access: connectivity, account availability, device capability and ability to obtain assistance. The second measures adoption and intensity: active use, verified business receipts, transaction failures and substitution from other methods. The third measures records and finance: reconciliation quality, applications, offers, total borrowing cost and repayment. The fourth measures operating outcomes: real value added, hours, inventory, investment and survival.

These layers should be linked through a stable establishment identity with documented changes in ownership or operation. Matching solely by a payment account can miss multiple accounts or shared household use. Matching solely by registration can exclude precisely the unregistered firms central to the question. The research design needs a defensible crosswalk and a clear account of unmatched observations.

A randomized encouragement to obtain training or assistance could identify the effect of that encouragement under appropriate conditions. It would not automatically identify the effect of forced adoption or every feature of a national platform. Compliance, spillovers and the content of assistance should be recorded. The intervention actually assigned is the object whose effect can be estimated.

Where randomization is unavailable, a panel comparison should test pre-adoption trends and concurrent changes, use appropriate survey weights and report sensitivity to firm size, sector and location. Statistical adjustment can reduce observable differences but cannot guarantee removal of unobserved management quality or expected demand. The conclusion should match the strength of the design.

The channel tests should be pre-specified. If transaction time falls but loan terms do not change, the evidence supports a payment-efficiency channel before a credit channel. If credit expands but real output does not, funds may have replaced other borrowing or relieved financial stress. If output per hour rises, verify that

prices, labor coverage and outsourcing do not explain the result. A mixed outcome is analytically useful because it locates the remaining constraint.

Reliability and assisted access belong in the outcome

An enterprise may be registered and willing to accept digital payments but face unreliable connectivity or difficulty resolving failed transactions. The relevant service is not only successful authorization under normal conditions. It includes the ability to identify settlement, correct mistakes and continue trading when a channel is unavailable. Measuring failure rates and time to resolution would add information that an annual adoption indicator cannot supply.

Assisted access also changes the resource account. A proprietor may rely on an employee, relative or agent to operate the service. That can enable participation, but the helper's time and any charge should enter the cost assessment. Excluding it can overstate the labor saved by digital adoption, particularly where the proprietor's own recorded hours fall while another person's unpaid effort rises.

The appropriate comparison retains the alternatives available to the enterprise and its customers. A new channel can improve choice while some customers still need another means of payment. Research should measure whether access broadens across relevant groups and whether new frictions exclude others, without assuming that one payment method is feasible for every transaction.

Policy implications and the standard of success

The official evidence demonstrates substantial payment activity and rising business internet use, alongside a much smaller movement in the presented registration share. It also documents a distinct infrastructure effort to connect lenders and information providers. Those observations justify examining the links among transactions, records and finance. They do not establish a national productivity multiplier.

Policy should follow the specific friction. Connectivity and assistance address access; simpler bookkeeping and reconciliation address usable records; portable and reliable data can address information costs; competition and appropriate underwriting determine credit terms; skills, capital and demand determine productive use. Expanding a payment interface cannot be expected to solve every downstream constraint independently.

The standard of success should include durable use, lower relevant costs, suitable finance and improved productive or welfare outcomes, while preserving security and meaningful choice. A business that saves administrative time without borrowing can benefit. A borrower obtaining cheaper finance without expanding output can also benefit. Those results should be described accurately rather than promoted into evidence of a different outcome.

India's digital infrastructure offers a setting in which these mechanisms can be studied at scale. Its analytical value is greatest when adoption, administrative status, financing and productivity remain separate measurable objects. Connecting them through actual enterprise records and a credible comparison can establish where the

infrastructure changes economic possibilities, where complementary reforms are needed and where the evidence does not yet support a causal claim.

Source addresses

The links below retain the source addresses cited in the analysis. References should be read within their stated periods and scope.

<https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=2097316&lang=2®=48>

<https://microdata.gov.in/NADA/index.php/catalog/238/variable/V57>

<https://www.npci.org.in/product/upi/product-statistics>

<https://www.rbi.org.in/scripts/AnnualReportPublications.aspx?Id=1436>

<https://www.pib.gov.in/PressReleasePage.aspx?PRID=2260497&lang=1®=6>

<https://www.nber.org/papers/w33259>