International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 11 | Nov 2026 www.irjet.net p-ISSN: 2395-0072

The Dark Side of Digital India: An Investigation into UPI Fraud, Financial Cybercrime, and Payment Security

Pooja Mehta[1], Dhruv Khatri[2], Abdullatif Dhopawkar[2], Arnav Naidu[2]

[1]*Assistant Professor, D. Y. Patil Deemed to be University, School of Humanities & Sciences, Belapur, Maharashtra, India.*

[2]*Students, Department of MCA, D. Y. Patil Deemed to be University, School of Humanities & Sciences, Belapur, Maharashtra, India.*

Abstract - *Unified Payments Interface (UPI) activity has grown rapidly in India, while reported fraud presents challenges for consumers, merchants, payment providers and investigators. This focused narrative review condenses a 62-page research and awareness report using parliamentary statistics, government and regulatory publications, provider guidance and selected academic literature. It distinguishes payment throughput, recorded incidents, amounts involved and final consumer loss, and separates account compromise from payments authorised after deception. The source reports 24,161.69 crore UPI transactions in FY2025-26 and 13.42 lakh domestic UPI fraud incidents in FY2023-24 under a specified parliamentary series. The analysis develops a three-stage prevention framework covering contact verification, payment decisions and response; it also reviews merchant controls, reporting, accessibility, monitoring and a proposed awareness evaluation. No new survey, experiment, original transaction-level investigation or independent verification of the underlying source documents was conducted for this manuscript. The findings support transparent reporting, usable security interfaces and testable public education, but do not establish individual fraud probability or the effectiveness of an unevaluated intervention.*

Key Words: UPI, digital payment fraud, social engineering, financial cybercrime, payment security, consumer protection, cybercrime reporting, digital literacy

1. INTRODUCTION

India's digital payments are routinely used for retail purchases, transfers and merchant receipts. The expansion of Unified Payments Interface (UPI) makes the circumstances of a payment decision an important security concern. An apparently familiar application can be involved in very different events: someone may misuse account access, or a legitimate user may authorise an outgoing transfer after accepting a false explanation. The two events can produce comparable financial harm, but their evidence and preventive controls differ. [1, 2, 6]

This manuscript summarizes the supplied extended research report rather than presenting a new national fraud survey. Its title highlights a security challenge, not a claim that digital payments are inherently unsafe. The analysis is intentionally limited to source-defined payment activity, reported domestic UPI fraud, selected user research and documented reporting guidance. The period of the underlying evidence review ends on 14 September 2026.

1.1 Problem statement and objectives

The research problem is to interpret reported fraud without conflating payment growth, complaints, successful criminal transfers, final unrecovered losses and the number of affected people. It asks what published data can show; where false authority, payment authorisation and account access produce different risks; how an effective response should involve institutions as well as users; and how awareness could be evaluated rather than simply advertised as effective.

The resulting objectives are to (i) describe UPI activity and recorded fraud with explicit periods and units; (ii) explain commonly discussed deception and account-access mechanisms; (iii) translate the evidence into household, merchant and institutional safeguards; (iv) set out reporting and follow-up considerations; and (v) propose a safe, measurable education programme. These are descriptive and programme-design objectives, not causal claims.

1.2 Scope and definitions

Digital payment fraud can include cards, internet banking, wallets and other payment channels. The quantitative analysis here centres on UPI and does not merge other channels into a UPI fraud total. “Reported incident” means the unit counted by a cited registry: it does not necessarily equal a unique victim, criminal case or successful transfer. “Amount involved” preserves the source's classification; it is not automatically the consumer's final unrecovered loss. An illustrative scenario is a teaching example, not a verified case. [1, 14]

2. MATERIALS AND METHODS

2.1 Focused narrative review

The source report used purposive searching and reference chasing to assemble parliamentary responses, government payment statistics, central-bank and provider guidance, cybercrime-reporting information, a peer-reviewed qualitative study and a model-evaluation preprint. Its search was not registered as a systematic review and did not screen every academic database. The condensed manuscript inherits that method and the source's stated evidence limits. [1-14]

Parliamentary annexures are used for historical activity and domestic UPI fraud reporting; an April 2026 government release provides a later activity total; the August 2026 observation is identified as secondary Reuters reporting. A December 2025 news report supplies a separately labelled update to the domestic-fraud series, which was not independently reconciled with the underlying primary response in the supplied review. The sources were selected for what each measures, not treated as interchangeable. [1-4]

2.2 Evidence categories and interpretation

Payment statistics measure numbers of transfers and money transferred. Fraud registers count events that enter a defined reporting process. User research explores the interpretation, constraints and experiences of recruited participants. Provider instructions establish suggested conduct, not the effectiveness of a warning. A technical classifier's test accuracy describes performance on a specified dataset, not the proportion of national fraud prevented. These distinctions shape all tables and figures in this summary. [5, 7, 8, 13]

Table -1: Measurement categories in the reviewed evidence (conceptual synthesis).

Evidence	What it measures	Key limit
UPI activity	Transfers and value	Not unique users or losses
Fraud reports	Incidents; amount involved	Not net loss or victim count
User interviews	Qualitative experiences	Not national prevalence
Model tests	Dataset classification	Not real-world harm reduction

For a comparison to be valid, units, period lengths, source vintages and reporting scopes must be consistent. Monthly gaps remain missing rather than zero; unfinished years are not projected into annual totals. An apparent rise in reports can reflect offending, improved detection, wider exposure, a new definition or better complaint access. A decline can likewise have multiple explanations. Without comparable groups or further case-level data, neither change demonstrates the effect of a single intervention.

2.3 Research limitations

The source did not collect original interviews, surveys, experiments, transaction records or case files. Its aggregate data cannot measure unique-victim prevalence, disaggregate the distribution of loss or identify the true cause of a time-series change. The smaller studies mentioned are not assumed to represent India's entire UPI population. This paper preserves those limitations: its conclusions concern measurement, practical decisions and research design, not estimated causal effects or a nationwide security ranking.

3. UPI ACTIVITY AND DESCRIPTIVE FRAUD FINDINGS

3.1 Growth in UPI transactions

The Ministry of Finance's April 2026 publication reports 24,161.69 crore UPI transactions in FY2025-26 and approximately INR 314 lakh crore in aggregate transferred value. The rounded annual chart shows 2,233 crore transactions in FY2020-21, 4,597 crore in FY2021-22, 8,375 crore in FY2022-23, 13,116 crore in FY2023-24, 18,587 crore in FY2024-25 and 24,162 crore in FY2025-26. The volume is a count of payments, not a count of distinct individuals or a fraud measure. [2]

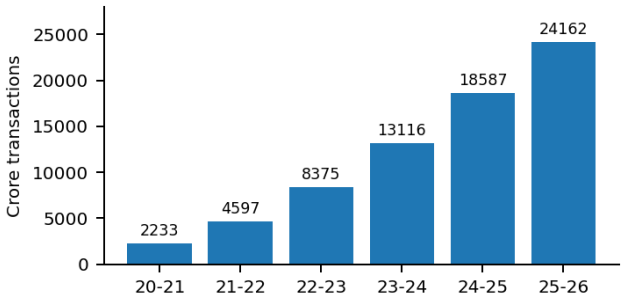


Chart -1: Annual UPI activity, rounded crore transactions, FY2020-21 to FY2025-26 [2].

The historical parliamentary annexure records 8,371.43 crore transfers for FY2022-23 and 13,112.96 crore for FY2023-24. These values differ slightly from the later rounded publication; the review retains the reporting vintages rather than silently combining them. The historical monthly annexure supplies FY2024-25 only through September, so its absent later months cannot be read as an activity decline. [1, 2]

The review also records March 2026 activity of 2,264 crore transfers with INR 29.53 lakh crore in value from the government release, while Reuters reported an August 2026 snapshot of 2,451 crore transfers and INR 29.82 lakh crore in value. These are two dated activity observations, not a continuous monthly series and not contemporaneous measurements of fraud. The August figures remain secondary reporting in this manuscript. [2, 4]

3.2 Reported domestic UPI fraud

In the cited CPFIR domestic UPI series, FY2022-23 had 7.25 lakh reported incidents involving INR 573 crore; FY2023-24 had 13.42 lakh incidents involving INR 1,087 crore. A separately marked December 2025 news account of newer parliamentary data reports 12.64 lakh incidents and INR 981 crore for FY2024-25. The later figure has a different source provenance and is not represented as an independently reconciled original annexure. [1, 3]

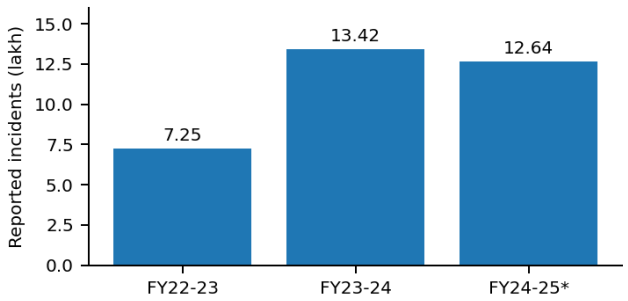


Chart -2: Reported domestic UPI incidents; the FY2024-25 value (*) is secondary-reported [1, 3].

Table -2: Reported UPI fraud. Asterisks indicate secondary reporting; partial periods are not comparable as annual totals [1, 3].

Reporting period	Incidents	Amount involved
FY2022-23, full	7.25 lakh	INR 573 cr
FY2023-24, full	13.42 lakh	INR 1,087 cr
FY2024-25, full*	12.64 lakh	INR 981 cr
Apr-Sep 2024	6.32 lakh	INR 485 cr
Apr-Nov 2025*	10.64 lakh	INR 805 cr

The last two table entries cover six and eight months respectively. Neither is a full financial year, and annualisation would require assumptions about seasonality, reporting delays and stable definitions. Combining partial-period fraud counts with a full-year activity denominator would create a mismatched ratio. The values also describe “amount involved”, not verified total unrecovered household loss.

3.3 Ratios and average amounts

For two complete matched financial years within the same parliamentary vintage, the original report derives about 8.66 reported incidents per million UPI transactions in FY2022-23 and 10.23 in FY2023-24. The corresponding amount-involved ratios are approximately INR 0.41 and INR 0.54 per INR million transferred. Mean amounts involved per reported incident are about INR 7,903 and INR 8,100. These are aggregate calculations, not individual victimisation probabilities or estimates of a typical person's final loss. [1]

A small ratio to an enormous payment denominator can coexist with serious harm to individual households. Conversely, an incident register may include several payments or repeat records related to the same individual. Without linked case-level data, a researcher cannot calculate a median loss, the proportion affected, the geographic distribution of personal risk or the probability that a future payment will be fraudulent.

3.4 Incompatible administrative series

A different December 2025 parliamentary response lists 457 UPI cases involving INR 2.13 crore for FY2021-22 to FY2025-26, through September 2025, in a bank/financial-institution digital-fraud table. Its coverage has not been shown to match the much larger CPFIR domestic UPI incident series. These counts must not be added, averaged or substituted for one another without an authoritative bridge of definitions, thresholds and reporting instructions. [14]

Geographic and demographic interpretations also require corresponding denominators. The cited historical parliamentary activity data do not record a usable transaction location, so they cannot support a state-by-state UPI fraud-rate map. The review contains no matching classification dataset with which to assign each national incident to a named scam mechanism. Accordingly, this manuscript does not invent maps, scam-type percentages or real-time fraud totals. [1]

4. HOW PAYMENT-RELATED FRAUD DEVELOPS

4.1 Three-stage prevention framework

The source proposes a three-stage sequence: (1) contact and claim; (2) consequential payment or account-access decision; and (3) detection, reporting and follow-up. At the first stage, the question is who appears to be asking, what they claim and how that claim can be verified independently. At the second, it is what action the screen actually authorises, where money would go and whether the action matches the independently confirmed purpose. At the third, it is how to limit further exposure, contact

authentic support and preserve a usable record. This is a conceptual teaching framework, not a statistical transition model.

Table -3: Three opportunities to interrupt fraud, conceptual framework from the supplied report.

Stage	User question	Proposed response
Contact	Is the claim genuine?	Verify using a known channel
Decision	Who receives what, and why?	Read the actual authorisation
Response	Was there an unusual event?	Report promptly; save references

A legitimate-looking payment interface may correctly authenticate a transfer but cannot establish that the outside story motivating it is genuine. This is the difference between successful authentication and informed purpose. A request describing a refund can still involve an outgoing transfer; a caller claiming official authority can provide plausible personal details while lacking authority to ask for a credential or a payment. [6-8]

4.2 Credentials, impersonation and account access

The reviewed provider guidance recommends keeping UPI PINs, one-time passwords and other payment secrets private and checking unexpected activity through legitimate support. Educational examples should focus on the exact moment when someone asks a user to reveal a secret, not merely on memorising a list of scam names. Banking, merchant, employment or family identities can be impersonated in a message; a logo, profile image or title is presentation rather than proof of authority. [6-8]

The user should be able to end a suspicious interaction and use a previously known number, official app or established relationship for verification. The source treats such independent verification as a useful proposed practice, not as a guarantee that every account or communication channel is uncompromised. Institutions must make legitimate contact routes consistent and accessible so that their own staff practices do not normalise unnecessary disclosure.

4.3 QR codes, false refunds and receipt screenshots

A QR code encodes instructions for an application; its appearance cannot authenticate the underlying reason for making a payment. Users need to read the displayed recipient, amount and transaction direction. Provider guidance states that receiving money does not require disclosure of a UPI PIN. Awareness should not overgeneralise the precise on-screen approval flow across every product and future feature; the durable check is what the current action authorises. [7]

A customer's screenshot is not sufficient merchant evidence that a credit actually posted. Shops should reconcile their bank or payment-provider record and handle a delayed confirmation without automatically alleging dishonesty. Likewise, a seller's invoice and a supplier's newly supplied account details may require separate verification of the commercial relationship and payment destination. These are proposed controls rather than outcomes tested by this review.

4.4 Remote assistance and persuasive offers

Remote-control or screen-sharing requests create risk when permissions expose personal or transaction data or allow another party to guide an authorisation. The educational exercise proposed by the source uses fictional static screens: participants must not reveal real accounts, install remote-control applications or enter live credentials as part of training. The review also discusses employment deposits, investment promises and successive unlocking charges as illustrative contexts in which a payment may be technically valid while the promise is false. [7, 12]

A securities-payment verification method such as SEBI Check has a specific institutional scope; it cannot validate every digital-payment request or guarantee investment returns. Training should distinguish verification of a destination from independent verification of the underlying employer, contract, transaction or claim. Prior satisfactory small transactions do not by themselves prove that a later or larger demand is genuine. [12]

5. HOUSEHOLDS, MERCHANTS AND PRIVACY

5.1 Merchant payment controls

Merchants receive customer transfers, pay suppliers, issue refunds and sometimes allow employees to operate payment devices. The summary proposes a workflow centred on actual credits, unique transaction references and controlled changes of payee details. Receipts should be reconciled against the merchant or bank transaction record; refunds should be linked to the original sale and authorised destination; a displayed QR code should be checked for unexpected replacement; and supplier-account changes should be verified through an established channel.

Different exceptions should have explicit procedures. A delayed credit requires an appropriate pending-payment process, not an accusation based on insufficient evidence. An unexpectedly changed supplier account may require a temporary pause and independent confirmation. A dispute

should preserve references, timestamps and the staff member's response. The report does not estimate how frequently any particular merchant workflow is attacked, and does not claim that these controls have been trialled in a representative merchant sample.

5.2 Assisted payments and accessible design

Households may include shared devices, translated interfaces, relatives who help interpret prompts or users who depend on assistance. A categorical instruction never to accept help would ignore these circumstances. The safer design objective is to preserve the user's control over the final transaction while allowing assistance with reading screens, finding genuine support and understanding a stated purpose. The source gives no evidence for a national vulnerability ranking by age, language or education, so such rankings are not inferred.

Payment instructions should clarify terms including send, receive, pending, approve and refund in the relevant language and interface context. Readable type, high contrast, meaning conveyed by text rather than colour alone and transcripts for audio materials are proposed accessibility measures. Their actual usability requires testing with intended participants rather than assuming that a translated warning will be understood automatically.

5.3 Minimal and confidential incident records

A private incident record may include contact and payment times, visible recipient details, transaction reference, amount, message evidence, support-case number and subsequent action. It should not contain an actual PIN, OTP, password or unnecessary complete account information. Public teaching examples should be fictional or appropriately consented and redacted: removing a name may not protect a person if the accompanying circumstances remain uniquely identifying.

Organisers of awareness sessions should not gather bank statements or personal identifiers merely to measure comprehension. The source proposes participant codes and separation of any genuinely necessary follow-up contact details from exercise answers. A retention plan should specify who can access records, what will be reported in aggregate and when private material will be deleted. These are design proposals; no actual participant-level collection took place in the source project.

6. INSTITUTIONAL MONITORING AND MODEL EVALUATION

6.1 Layered technical and service controls

Payment security is not solely a matter of user attention. Device binding, authentication, transaction controls, fraud-monitoring systems, service design and incident handling each address a different opportunity for prevention or recovery. The source notes an NPCI fraud-monitoring solution available to banks, but its cited public aggregates do not measure how many specific incidents were prevented by any named technical control. [1]

A suspicious transaction signal may lead to a warning, a review, a temporary restriction or another action. An operational evaluation should count confirmed outcomes and not treat raw alert numbers as prevented losses. It should document false positives, time to resolution, user recourse and whether legitimate transfers are obstructed. The ability to contest a false alarm and identify authentic support is part of the proposed security response, not an unrelated customer-service feature.

6.2 Interpreting artificial-intelligence claims

A cited research preprint on large-language-model-based scam analysis reported 93.33% accuracy in a curated Google Pay-related evaluation. This is the preprint's dataset-specific classification result, not a tested national UPI fraud-prevention rate or a guarantee for other users or institutions. Before deployment, evaluation would need the tested unit, prevalence and class balance, held-out data, missed fraud, false alarms, shifts in messages over time and the real intervention triggered by an alert. [13]

A model can sound certain while offering an incorrect reason for a flag. Where an alert can stop or delay an important payment, a person or institution should have a documented mechanism to review errors and communicate the basis of the action. Privacy also matters: collecting broad message or transaction histories for model training can create exposure unless access, retention and purpose are controlled. These are implementation questions, not findings from an original machine-learning experiment in this manuscript.

7. REPORTING, LOSS AND LIABILITY

7.1 Response after suspected fraud

On recognising an unusual payment, a person should stop further interaction with the suspicious party and promptly contact the bank or payment provider through an authentic route. The supplied review also identifies India's national financial-cybercrime reporting helpline 1930 and cybercrime.gov.in. A bank/service-provider complaint and

a national portal report can serve related but different purposes; users should keep both references where applicable and follow official instructions for submission and follow-up. [9]

The source encourages concise factual reports: when contact occurred, which transfer is disputed, how the contact represented itself, what appeared on-screen and when the bank and reporting service were contacted. Preserve transaction references, acknowledgment numbers and relevant messages privately. Do not include real secrets in a shared awareness handout or public post. This response guidance does not guarantee a reversal, fixed refund time or positive outcome for every complaint.

7.2 Why a complaint is not a refund

A report can be acknowledged, awaiting information, under investigation, associated with temporarily held funds, closed without reimbursement or followed by an actual reversal. These are separate states. Money frozen in an account is not necessarily money returned to an affected consumer. A resolved complaint may still be rejected. The source's monetary fraud totals therefore remain “amount involved” where that is the underlying label; they are not automatically the net permanent loss of victims.

7.3 Policy chronology and case-specific liability

The original review cites RBI's 6 July 2017 framework concerning customer liability for unauthorised electronic banking transactions, and a 24 June 2026 RBI release announcing amendment directions stated to take effect on 1 January 2027. Its evidence snapshot is 14 September 2026, before that stated effective date. Case-specific rights depend on the applicable institution, operative directions, timing, circumstances, the classification of the transaction and the reporting record. [10, 11]

An action induced by deception and an unauthorised transaction cannot automatically be placed in the same liability category. A training guide should avoid promising a universal refund or declaring a complainant negligent from an unsuccessful outcome. A careful incident record preserves facts for the institution or competent authority to assess under the relevant rules.

8. PUBLIC AWARENESS PROGRAMME AND PROPOSED EVALUATION

8.1 Education focused on decisions

The source proposes teaching observable decisions rather than requiring participants to memorise a catalogue of scam names. A useful learner should explain the direction of a payment, identify the intended recipient and amount, recognise a request for a secret, verify an outside claim using an independent route and find an authentic reporting channel. An urgent request may allow more than one safe response, including consulting an established contact or postponing a transfer. [5-8]

A 60-minute proposed workshop begins with privacy rules and baseline scenarios (0-15 minutes), explores payment meaning and access protection (15-25), practises independent verification in fictional situations (25-40), rehearses reporting and record keeping (40-50), and ends with a new post-session decision assessment and questions (50-60). All screens, recipients and identities must be fictional or non-functional. It should not send a real transfer, solicit passwords, install remote software or make test complaints to live helplines.

8.2 Measurable outcomes and comparison

An awareness programme's effectiveness cannot be inferred from attendance, confidence or the distribution of a poster alone. The source proposes a governed evaluation using baseline, immediate and delayed non-identical scenarios. Outcomes might include correctly recognising an outgoing transfer, declining a credential request, independently confirming a claim and selecting a valid reporting route. Legitimate payment scenarios must also be included so the programme does not train indiscriminate refusal.

Where ethically feasible, participants could be assigned to an enhanced workshop or a suitable comparison activity, with a pre-specified primary decision outcome. Recruitment, sample frame, sample size justification, attrition, language adaptation, scoring and missing-data handling should be documented before assessment. Delayed tasks could test retention several weeks later, while follow-up should not pressure participants to disclose private financial history. Any measured effect would describe the actual recruited population and design, not automatically all UPI users.

8.3 Monitoring unintended consequences

An overly alarming programme could increase anxiety, impede legitimate payments, generate unnecessary support contacts or create excessive confidence after a simple quiz. A responsible evaluation should measure mistaken refusal, support-seeking, comprehension, accessibility and adverse experiences, rather than reporting only an increase in correct answers. Any real victimisation disclosed during a session should trigger an appropriate private support pathway, not become an unsolicited research record.

A publishable pilot would disclose the curriculum, fictional examples, primary and secondary outcomes, observed

implementation, negative as well as positive findings and departures from the planned method. Since no such pilot was performed in the supplied research, this section is a proposed protocol and does not claim a measured improvement in fraud resistance.

9. OPERATIONAL CHECKS AND EVALUATION SAFEGUARDS

9.1 Six fictional decision scenarios

The original report includes a draft, unvalidated assessment instrument intended for refinement before research use. Its questions address an alleged refund with an outgoing-payment screen, a caller requesting a secret code, an unexpected support number, a merchant receiving an uncertain screenshot, an unfamiliar transfer and changed payment details at a familiar organisation. These are fictional decision prompts, not documented crimes or observed case frequencies. For the instrument to measure reasoning, learners should be able to explain why they choose an action rather than merely label the scenario a scam.

Table -4: Scenario constructs adapted from the source's proposed assessment instrument; not validated cases.

Fictional prompt	Observable decision
Refund story; outgoing transfer	Explain direction and mismatch
Caller requests payment code	Withhold a secret; verify
New support number in message	Find independent contact
Buyer shows a receipt screenshot	Check actual merchant credit
Unfamiliar transaction appears	Use official reporting channels
Familiar supplier changes payee	Independently confirm changes

For the refund prompt, a correct explanation identifies that money would leave the account and that the on-screen action contradicts the stated purpose. For the credential prompt, the learner should identify what must remain private and distinguish a legitimate app flow from a person asking for the secret. For the support-number prompt, the response should describe an established route not supplied solely by the suspicious contact. The expected conduct need not be a single rigid sentence: stopping and checking, consulting a reliable contact or obtaining the institution's authentic support details can all be reasonable ways to handle uncertainty.

Merchant scenarios need balanced treatment. If the customer shows an apparent proof of payment but the seller cannot see a corresponding credit, the exercise should encourage checking reference, amount, status and the appropriate merchant or bank record. A delayed payment notification is not itself evidence of dishonesty. The learner should identify a proportionate pending-payment or dispute process, preserve the available details and avoid either releasing goods without confirmation or publicly accusing a customer solely because a screenshot is inconclusive.

Changed destination details are a further stress test. A known organisation may send a genuine invoice while a deceptive message claims that its bank account or UPI address has changed. The relevant question is not whether the new destination appears to have a familiar name, but whether the change was confirmed through a route independent of the new instruction. The exercise should treat the invoice, contract and destination as separate items of evidence. It should not request real supplier payment details from participants.

9.2 Valid scoring and retention claims

A proposed scoring rubric should identify concrete decision features such as noticing the direction of money, identifying an inconsistent claim, locating an independent verification route, protecting a secret and selecting appropriate reporting channels. The source does not provide a validated numerical threshold that categorises a person as safe or unsafe. Any rubric should first be tested for ambiguous wording and whether different trained assessors can score an explanation consistently. Where possible, scorers should not know whether a response came from a workshop or comparison group.

An immediate post-session improvement could reflect short-term recall of a recent example. A delayed assessment using new scenarios can test retention and the transfer of a principle to a different context. Even successful delayed scenario performance does not establish a reduction in actual monetary loss: people know that they are in an assessment and may behave differently during an unexpected real-world interaction. The eventual publication should therefore name the measured endpoint, its time point and its participant population precisely.

The source recommends including legitimate comparison requests in the assessment, because a learner could appear successful on a set of suspicious examples by declining every action. This approach can reveal false refusal, uncertainty that leads to prolonged delays and possible harm from an excessively alarming message. Confidence should be compared with demonstrated decision performance instead of accepted as proof that the person has learned. Open explanations can reveal a correct choice made for the wrong reason or an appropriate cautious choice that a simplified multiple-choice key would incorrectly penalise.

9.3 Sampling, statistical analysis and uncertainty

The target population needs to be named before recruitment. A convenience sample from a single college can describe those students, not every Indian consumer; a merchant-group sample may reflect that group's routines rather than occasional household transfers. Recruitment should specify eligibility, recruitment channel, language, assistance needs, refusals and exclusions. Any power or precision calculation must be justified by the primary outcome, plausible variation, intervention allocation and expected loss to follow-up, not selected after seeing the outcome.

If entire classes or community groups receive a workshop, responses may be correlated within the group. The design should therefore consider whether allocation and analysis need to account for such clustering. If individuals in one setting are placed in different conditions, sharing between them may affect the interpretation. Neither design is automatically suitable for every setting; the source proposes making the trade-offs explicit in advance. Research participants who disclose active financial harm should receive publicly available assistance independently of their study allocation or continued participation.

For a future binary decision outcome, the source proposes a transparent group comparison of the proportion making the prespecified safe decision, alongside uncertainty appropriate to the design. For a continuous score, an interpretable difference may be considered if the instrument supports that scale. The analysis plan should set the main time point, baseline adjustment, exclusions, missing-response rules and any limited subgroup analyses before viewing the main results. The descriptive national statistics used earlier in this paper are published aggregates, not a randomised experiment, and do not acquire causal certainty by adding p-values.

Attrition can distort results if participants who found the material confusing are less likely to return. An evaluation should report follow-up completion by condition and examine reasonable sensitivity analyses for missing responses. It should avoid treating unanswered items as successes and should not conceal unplanned changes in task wording or score definitions. The interpretation should attend to the size and practical relevance of any observed difference, not only to whether an arbitrary significance threshold was crossed.

9.4 Ethical safeguards and publication transparency

An awareness study could cause distress, embarrassment or unwanted disclosure of a real financial incident. A future project requires suitable institutional review, meaningful voluntary consent, private options for questions and a referral pathway for participants seeking actual reporting guidance. The supplied report used public materials only and does not claim that review approval, recruitment or consent occurred. No simulated exercise should ask a participant to send real money, provide bank credentials, install remote-control software or place a test call to a live cybercrime helpline.

Contact information for a delayed assessment should be limited to what is necessary, safeguarded and stored separately from scenario answers where practical. The plan should specify access permissions, retention, withdrawal arrangements and aggregate publication rules. Participants should not be required to share case details publicly or participate in research in order to receive general safety information. Any real case assistance should be separate from the scoring process, and facilitators should not promise a privileged recovery channel.

Authorship and source responsibility remain with the eventual named author. Before journal submission, the author should validate the bibliography and numerical calculations against the cited documents, confirm the manuscript complies with the journal's current instructions, supply real affiliations, disclose any AI-assisted preparation as required by the venue and make accurate ethics, funding and competing-interest statements. The material should be described as a condensed narrative review with descriptive secondary analysis and a proposed evaluation, not as newly conducted fieldwork.

9.5 Implementation and revision cycle

The source envisages a sequence from source-checked educational material and preliminary cognitive testing to small pilot delivery, a prespecified main evaluation and appropriately labelled public distribution. Pilot sessions should identify confusing instructions, unrealistic scenarios and accessibility barriers before an efficacy claim is made. Delivery logs can record curriculum version, facilitator, language, group setting, session length and deviations without creating a large collection of unnecessary financial details. Attendance is a process measure and should remain distinct from participant consent and demonstrated learning.

A public awareness handout should also have an update owner. Support routes, applicable legal directions and app-specific screens can change after a report is written; a printed message should identify the date on which it was checked. If a substantive instruction changes, the prior material should be retired or marked outdated rather than silently edited and redistributed. Revisions should distinguish corrections of numeric source values from changes in the underlying measurement category. Maintaining a version history makes later research and public communication more reproducible.

Finally, transparent distribution requires naming what was actually delivered and assessed. An education programme can claim a reach number only after it records attendance, a learning improvement only after a suitable assessment, and a reduction in financial harm only after measuring that outcome with an appropriate comparison and observation period. Reporting null results, unintended barriers and unresolved limitations allows other educators and institutions to adapt the programme rather than merely repeating favourable promotional claims.

9.6 A traceable incident-response pathway

The report's reporting model begins when a person first notices an unfamiliar debit, a suspicious request for a secret or a mismatch between an expected receipt and a displayed outgoing payment. The first practical distinction is between an attempted deception and a transfer that has already occurred. A person who has not authorised a payment can stop the interaction and verify the claim without inventing a transaction reference. A person who sees an actual unexpected transfer may need the bank or provider to locate that transaction immediately. Both situations deserve a factual record, but they are not the same procedural event.

An incident timeline helps maintain clarity when several channels are involved. It can distinguish the time a suspicious call or message arrived, the moment an authorisation was shown, the time a transfer was submitted or confirmed, the time the discrepancy was discovered and the time each support service was contacted. When a timestamp is uncertain, the record should say so rather than supplying a false precision. An affected user may record whether money moved, whether the account or device behaved unusually and what message initially persuaded them to act. These observations are more useful than a speculative conclusion about who ultimately controlled the recipient account.

The source recommends using genuine bank or provider support and the national cybercrime reporting route without assuming the two systems have identical roles. The user should keep the acknowledgement or ticket number from each service, the method and time of submission and any requested follow-up. Copies of notices and payment evidence should be stored privately and provided only through appropriate channels. Where instructions from a suspicious caller conflict with an independently reached support service, the independent channel is the relevant verification route. This discussion is educational and does not replace case-specific directions issued by the responsible institution or authority. [9]

Table -5: Proposed incident record and follow-up stages. Records should never include secrets.

Incident stage	Record to preserve	Meaning of the stage
Suspected contact	Time, medium, claimed identity	Attempt or claim observed
Payment event	Amount, status, reference	Observed action or transfer
Bank/app report	Submission time, ticket ID	Institutional acknowledgement
Cybercrime report	Portal/1930 reference if used	National-report record
Follow-up	Requests, decision, fund status	Held, returned and rejected differ

A payment record and a complaint record should not be merged into a single ambiguous status. One suspicious message may lead to no transfer, one transfer or a sequence of payments. Multiple support contacts may be connected with a single underlying event, and a report may concern more than one transfer. Without linked records and a documented deduplication rule, counting all records as independent victims overstates what is known. That distinction matters both to an individual seeking a coherent case history and to a researcher attempting to interpret administrative totals.

A request for further evidence should be tracked separately from a decision. An acknowledgement confirms receipt of the complaint, not its substantiation. A temporary hold identifies an interim financial action, not completed reimbursement. A reversal, refund, rejection and still-open case are different final or current states. The incident record should update the status only when the institution supplies new information, noting the date and the communication channel. No entry should be labelled “recovered” merely because another institution has reported that money was located or restricted.

9.7 Making statistical figures reproducible

The manuscript's data treatment begins with a compact register for each figure: source identity, publication date, period covered, payment channel, measurement unit, coverage description, whether the number is rounded and whether the document is primary or secondary. For instance, a crore-transaction annual total must not be placed directly against a lakh-incident count on an unlabeled common axis. Similarly, an INR lakh-crore throughput figure and an INR crore amount involved in fraud are different financial objects even after unit conversion. A footnote should identify the relevant

reporting source rather than implying that all figures were retrieved directly from NPCI or RBI.

Annual series require a stable financial-year window. The review deliberately shows an April-September count as a six-month observation and an April-November count as an eight-month observation. Dividing each figure by the number of observed months can create monthly averages but cannot by itself correct for lagged reporting, seasonal payment patterns or changes in recording practice. Projecting either period to twelve months would introduce a forecast that the original report expressly did not make. The safest comparison is a matched period with the same reporting definition and a clear note about remaining limitations.

A simple ratio requires a matching numerator and denominator, a transparent conversion and an appropriately cautious label. The incident ratios quoted earlier use complete years from the same historical parliamentary vintage. They are written as reported incidents per million UPI transfers, not "fraudulent transfers per million" because one registered incident has not been shown to correspond to exactly one transaction. The monetary comparison retains "amount involved per million rupees transferred," not an individual expected loss. Source totals can be reclassified or revised; a later issue should either document those changes or preserve the earlier analytical vintage when presenting a historical calculation.

Mean amount involved per incident is a descriptive quotient of a reported money total and incident count. Without the underlying distribution, it cannot reveal the median amount, the share of losses borne by a small number of people or the experience of a typical affected household. An observed increase in that quotient may reflect a changing case mix as well as different loss severity. A privacy-preserving release with amount bands, repeat-victim indicators, times of occurrence and final outcomes would support more careful analysis, but those missing variables should not be manufactured from existing aggregate totals.

The source illustrates a further data hazard through a bank/financial-institution UPI category that contains far fewer cases than the CPFIR domestic UPI table for another stated coverage. Selecting whichever count appears more dramatic, or averaging the two, would obscure the unresolved reporting difference. A future authoritative bridge should clarify source reporting institutions, incident thresholds, category boundaries, transaction versus complaint units and whether one underlying event can enter multiple systems. Until those relationships are documented, the reporting universes remain separate. [1, 14]

Table -6: Minimum labels to retain when reusing a figure from the source review.

Chart or table label	Required qualification
UPI transaction activity	Period, count versus value, rounding
Reported fraud incidents	Source-defined reporting universe
Amount involved	Not necessarily unrecovered loss
Year-to-date observation	Start/end dates and duration
Derived reporting ratio	Matched units; not personal probability
Conceptual process diagram	Proposed framework, not measured frequency

Editorial checking is essential because a figure may be copied without its surrounding explanation. A reader should be able to answer whether it refers to UPI only or all digital payments; a full year or a partial period; transactions, reports or people; and observed values or a proposed teaching diagram. An online infographic should preserve the source caption and evidence date even if the chart itself is reduced in size. Otherwise a technically correct graphic can create an incorrect impression when stripped of the qualifiers that define it.

10. DISCUSSION AND RECOMMENDATIONS

The empirical observations establish substantial UPI use and sizeable amounts of fraud registered in a defined domestic series, but they do not provide a complete estimate of victim counts, net household losses or the effectiveness of any one safeguard. The payment system's aggregate throughput should not be used to trivialise severe individual loss; a high raw complaint total should likewise not be used to invent the probability that a given person will be victimised. These interpretations require attention to measurement and to the population actually observed.

The conceptual three-stage framework suggests complementary responsibilities. Users need understandable ways to verify an incoming claim and evaluate what a transaction actually does. Providers should make authorisation consequences and support routes intelligible. Merchants should reconcile credited payments and validate unusual changes. Investigating and regulatory systems should distinguish reports, funds held, recovered sums and final outcomes in accessible statistics. Each proposed improvement requires evaluation for both protection and unintended barriers.

For future public reporting, the review recommends publishing explicit reporting categories, coverage, revised definitions, financial-year or year-to-date period

boundaries and privacy-preserving amount bands. Better linked aggregate outcomes could support analysis of repeat reports, recovery and concentrated loss without exposing named victims. Geographic risk comparisons would additionally require compatible transaction-location denominators and verified fraud-location definitions. [1, 14]

For communication and institutional design, warning messages should distinguish a request to send money from a claim that money will be received; the recipient and purpose should remain readable at authorisation; and authentic support should be reachable after an alert. Research should test such messages with varied language and assistance needs while recording false positives and missed detections. This manuscript recommends those tests, not a numerical loss reduction or proven superiority of an untested interface.

11. CONCLUSION

This condensed review distinguishes UPI payment growth from administrative fraud records and individual consumer harm. Its documented figures establish a large scale of payments and reported domestic UPI fraud, but not a personal risk probability, a nationally representative account of each fraud type, or a guaranteed route to recovering money. The relevant prevention opportunities extend from independent verification of a contact, through interpretation of the payment being authorised, to accessible reporting and evidence-preserving follow-up.

Consumer awareness is one component of a broader response involving bank and app design, merchant reconciliation, technical monitoring and accountable complaint handling. Future work should seek well-defined reporting datasets and test proposed awareness and interface improvements under clear governance. The source report and this manuscript are review-based summaries; neither presents new fieldwork, an evaluated intervention, a live fraud dashboard or peer-reviewed original findings.

12. REFERENCES

[1] Ministry of Finance, Department of Financial Services, “Fraud in UPI Transactions,” Lok Sabha Unstarred Question 211, Annexures I-II, 25 Nov. 2024.

[2] Ministry of Finance / Press Information Bureau, “UPI completes 10 glorious years,” 30 Apr. 2026, annual UPI activity release.

[3] S. Sen, “UPI frauds peak in FY24, show signs of decline: Parliament data,” YourStory, 15 Dec. 2025. Secondary report.

[4] A. Manikandan, “India launches tap-and-pay option for real-time payments,” Reuters, 10 Sep. 2026. Secondary activity snapshot.

[5] D. Mungara, H. Sri Ramulu and Y. Acar, “Security and Privacy Advice for UPI Users in India,” 34th USENIX Security Symposium, pp. 6085-6103, 2025.

[6] Reserve Bank of India, “RBI releases booklet BE(A)WARE on modus operandi of financial frauds,” 7 Mar. 2022.

[7] Google Pay Help, “Protect yourself against fraud & unauthorized activity,” accessed 14 Sep. 2026.

[8] Google Pay Help, “Keep your payment info safe,” accessed 14 Sep. 2026.

[9] Indian Cybercrime Coordination Centre, “National Cybercrime Reporting Portal,” accessed 14 Sep. 2026, https://cybercrime.gov.in/.

[10] Reserve Bank of India, “Customer Protection - Limiting Liability of Customers in Unauthorised Electronic Banking Transactions,” 6 Jul. 2017.

[11] Reserve Bank of India, “RBI Issues Amendment Directions for Review of Framework of Limiting Customer Liability in Digital Transactions,” 24 Jun. 2026; stated effective date 1 Jan. 2027.

[12] Securities and Exchange Board of India, Annual Report 2025-26, ch. 8, Investor Protection, 2026.

[13] D. Dahiphale et al., “Enhancing Trust and Safety in Digital Payments: An LLM-Powered Approach,” arXiv:2410.19845, 2024. Preprint.

[14] Ministry of Finance, Department of Financial Services, “Cyber Frauds,” Lok Sabha Unstarred Question 1213, 8 Dec. 2025.

Source note: References and figures follow the bibliography and evidence snapshot of the supplied 62-page review. This condensed manuscript has not independently reverified every external source. Author and institutional details remain to be completed before submission.