AFTER THE MONEY IS GONE: BANK LIABILITY AND THE ALLOCATION OF CYBER-FRAUD LOSSES IN INDIA - RETHINKING FRAUDULENTLY INDUCED PAYMENTS, MULE ACCOUNTS AND VICTIM RESTITUTION

Dr. Cumaran Nadaradjan, B.A. LL.B., LL.M., Ph.D. (Law)
Advocate, Bar Council of Delhi, Delhi, India

ABSTRACT

India's digital-payment architecture has made transfers faster, cheaper and more accessible, but the same speed and finality have intensified a difficult consumer-protection problem: who should bear the loss when a customer technically authorises a payment only because a fraudster has manipulated, impersonated, threatened or deceived the customer? Existing Reserve Bank of India rules provide an important liability framework for unauthorised electronic transactions, yet fraudulently induced payments do not fit comfortably within a model built around the binary distinction between authorised and unauthorised transactions. At the same time, beneficiary accounts and money-mule networks frequently determine whether stolen funds can be dissipated before intervention is possible. The Supreme Court's continuing proceedings concerning victims of "digital arrest" scams have brought these problems into sharper focus, including the need to consider shared liability, victim compensation, mule-account controls and expedited restoration of recovered funds.

This article argues that loss allocation should not turn exclusively on whether the victim physically initiated or authenticated the transfer. It proposes a prevention-and-control approach that examines which participant was best positioned to identify and interrupt the fraud at each stage of the payment chain. The article develops a five-stage framework—Prevent, Intervene, Trace, Allocate and Restore—for allocating responsibility among customers, originating banks, beneficiary banks and other regulated participants. It contends that a calibrated shared-liability regime, supported by rebuttable presumptions and rapid restitution procedures, would better align incentives, consumer protection and payment-system integrity without converting banks into insurers against every act of deception.

Keywords: Cyber Fraud; Bank Liability; Digital Payments; Mule Accounts; Victim Restitution; Authorised Payment Fraud; Reserve Bank of India.

I. INTRODUCTION

India's payments revolution has altered the practical meaning of banking. A transfer that once required a visit to a branch, a written instrument and a measurable period for clearing can now be completed in seconds. That transformation has produced enormous social value. It has also changed the economics of fraud. Criminals no longer need to defeat a bank's internal security in every case. Increasingly, they defeat the customer.

The most difficult cases are not classic unauthorised withdrawals in which credentials are stolen and money moves without the customer's knowledge. They are transactions in which the customer is present, authenticated and apparently consenting, but the consent has been manufactured by fraud. A victim may believe that she is transferring money to a safe account identified by a police officer, complying with a regulator, protecting herself from arrest, assisting an investigation, paying a legitimate vendor, or moving funds at the direction of a trusted institution. Technically, she presses the button. Legally and economically, however, the transaction is the product of deception.

This distinction matters because the Reserve Bank of India's customer-liability framework is principally framed around "unauthorised electronic banking transactions." The framework provides important protections where the bank is at fault, where a third-party breach occurs, and where the customer reports promptly. But an authorised-looking transfer induced by social engineering can fall outside the intuitive centre of that regime. The resulting dispute is often reduced to a simple proposition: the customer authenticated the transaction, therefore the customer bears the loss.1

That proposition is increasingly difficult to defend as a complete rule. Modern payment fraud is an ecosystem failure. The victim may have been deceived, but the fraud may also depend on a newly opened beneficiary account, abnormal transaction velocity, inconsistent KYC information, repeated inbound transfers from unrelated persons, immediate dispersal of funds, or other signals visible to regulated entities. RBI has warned banks about money-mule accounts for years and has required diligence and transaction monitoring designed to reduce their use. The question is therefore not merely whether the customer made a mistake. It is whether the financial system allocated prevention duties to institutions capable of detecting risks that an

1 Reserve Bank of India, Customer Protection—Limiting Liability of Customers in Unauthorised Electronic Banking Transactions, RBI/2017-18/15, DBR No. Leg. BC.78/09.07.005/2017-18 (July 6, 2017).

individual consumer could not reasonably evaluate.2

The Supreme Court's continuing suo motu proceedings concerning victims of digital-arrest scams have transformed this from a theoretical debate into an immediate policy question. In August 2026 the Court directed the Inter-Departmental Committee to discuss a shared-liability and victim-compensation framework, to consult banks and intermediaries on technological prevention and recovery measures, and to ensure expeditious disbursement of recovered amounts to victims. Earlier proceedings had addressed temporary debit holds, money-mule detection and a uniform process for custody and restoration of funds.3 4

This article uses those developments as a point of departure but does not treat "digital arrest" as a self-contained offence or the sole factual setting. The deeper problem is broader: how should Indian law allocate losses from fraudulently induced electronic payments?

The article advances four propositions. First, the authorised/unauthorised distinction remains useful but is no longer sufficient as the exclusive organising principle for consumer liability. Second, responsibility should follow preventability: legal rules should ask which participant had the information, capability and reasonable opportunity to prevent or interrupt the loss. Third, beneficiary banks deserve greater attention because mule-account governance is central to the monetisation of cyber fraud. Fourth, restitution must be treated as a core component of payment-system governance rather than an incidental consequence of criminal investigation.

The article proposes a five-stage framework—Prevent, Intervene, Trace, Allocate and Restore ("PITAR")—for evaluating responsibility. The model does not impose automatic reimbursement in every fraud case. Instead, it creates structured duties, rebuttable presumptions and a fact-sensitive allocation mechanism. Its purpose is to move Indian law away from the false choice between total bank liability and total customer liability.

2 Reserve Bank of India, Operation of Bank Accounts—Money Mules, RBI/2010-11/383, DBOD.AML.BC.No.77/14.01.001/2010-11 (Jan. 27, 2011).
3 In Re: Victims of Digital Arrest Related to Forged Documents, Suo Motu Writ Petition (Crl.) No. 3 of 2025, order dated Feb. 9, 2026 (Supreme Court of India) (recording measures concerning MuleHunter.AI, temporary debit holds, NCRP-CFCFRMS and restoration of funds).
4 In Re: Victims of Digital Arrest Related to Forged Documents, Suo Motu Writ Petition (Crl.) No. 3 of 2025, order dated Aug. 4, 2026 (Supreme Court of India) (directing discussion of a shared-liability and victim-compensation framework and expeditious disbursement of recovered amounts).

II. THE LIABILITY GAP CREATED BY FRAUDULENTLY INDUCED PAYMENTS

The legal difficulty begins with vocabulary. A payment can be "authorised" in the narrow technological sense because the customer entered a PIN, approved an application prompt, supplied an OTP, or initiated a transfer. Yet the same payment may be involuntary in a broader legal sense because the customer's decision was procured by impersonation, fear or deception.

Traditional banking rules were developed around a different fraud model. The paradigm case involved a thief who acquired a card, password or cheque and initiated a transaction without the account holder's participation. Authentication therefore served as a useful proxy for consent. The stronger the authentication, the stronger the inference that the transaction was genuine.

Social engineering disrupts that proxy. The criminal does not necessarily attack the authentication mechanism; the criminal recruits it. The victim is induced to become the instrument through which the bank's security process is satisfied.

This creates what may be called the "authentication fallacy": treating successful authentication as conclusive proof that the economic risk should lie with the customer. Authentication establishes that a credential or device was used. It does not establish that the customer understood the true identity of the recipient, the true purpose of the transfer, or the fraudulent circumstances in which the instruction was generated.

The problem is especially acute in impersonation scams. A victim who believes that she is speaking with a police officer or regulator may experience the transaction as compelled compliance rather than voluntary payment. The criminal's objective is to preserve the appearance of customer authorisation precisely because that appearance weakens the victim's later claim against the bank.

Indian law should therefore distinguish at least three categories. The first is the genuinely unauthorised transaction, initiated without the customer's knowledge or approval. The second is the fraudulently induced transaction, initiated by the customer but procured through material deception, impersonation or coercive manipulation. The third is the ordinary authorised transaction in which the customer knowingly assumes the commercial risk of the payment.

These categories should not necessarily produce identical remedies. But the second category

cannot sensibly be collapsed into the third.

A customer who transfers money to a legitimate but unsuccessful investment has accepted a commercial risk. A customer who transfers money to a criminal impersonating a government official has not accepted the risk that the recipient is a fraudster in the same normative sense. The transaction is "authorised" only if authorisation is defined mechanically.

The challenge is to translate this insight into administrable law. If every allegation of deception automatically generated reimbursement, banks would face opportunistic claims and potentially unsustainable costs. Conversely, if authentication automatically defeated reimbursement, institutions would have insufficient incentives to invest in recipient-side controls and behavioural fraud detection.

A workable regime must therefore avoid absolutes. It should examine the quality of the deception, the warnings provided, the transaction's objective risk indicators, the customer's conduct, the bank's monitoring, the beneficiary account's characteristics and the speed of the response after notification.

III. RBI'S UNAUTHORISED-TRANSACTION FRAMEWORK: STRENGTHS AND LIMITS

The RBI's 2017 circular on customer protection was an important shift toward structured allocation of electronic-banking losses. It recognises zero liability where the bank's contributory fraud, negligence or deficiency causes the unauthorised transaction. It also provides zero liability in specified third-party breach situations when the customer reports promptly. Where customer negligence contributes to the loss, the customer bears loss until reporting, while subsequent loss is allocated to the bank. The circular further places the burden of proving customer liability on the bank.5

These rules embody several sound principles. They recognise that financial institutions are not passive conduits. They impose incentives to maintain secure systems. They value rapid reporting. They allocate proof burdens to the institution that possesses superior records. And they acknowledge that loss allocation can change over time: the customer may bear one period

5 Reserve Bank of India, Customer Protection—Limiting Liability of Customers in Unauthorised Electronic Banking Transactions, supra note 1.

of loss while the bank bears another after notice.

The weakness is not that the circular is poorly designed for its subject. The weakness is that its subject is "unauthorised" transactions.

Fraudulently induced payments expose the limits of that category. A bank may argue that the transaction was authorised because the customer initiated it. The customer may answer that the instruction was procured by criminal deception. The circular does not provide a sufficiently detailed framework for deciding when an authenticated transaction should receive enhanced protection because the customer's apparent consent was corrupted.

This gap is not merely semantic. It shapes incentives. If banks know that a successfully authenticated transfer ordinarily shifts the loss to the customer, the strongest economic incentive may be to perfect payer authentication while underinvesting in beneficiary-account intelligence. Yet social-engineering fraud often succeeds even when authentication works exactly as designed.

The better regulatory question is therefore: what could the system reasonably have known?

A bank may not know that a caller is impersonating a police officer. But it may know that a long-dormant customer is suddenly transferring most of her savings to several new beneficiaries. A beneficiary bank may not know the victim's circumstances. But it may know that a recently opened low-activity account is receiving large transfers from unrelated persons and immediately forwarding them onward.

The current framework should be extended by regulation or legislation to address this information asymmetry directly.

Such an extension would be consistent with the logic of the 2017 circular. The circular already rejects a simplistic rule that every loss falls on the account holder. It already differentiates among bank fault, customer fault and third-party breach. Fraudulently induced payments require the next step: a structured inquiry into shared preventability.

IV. MONEY MULES AND THE NEGLECTED ROLE OF THE BENEFICIARY BANK

The public narrative of cyber fraud usually begins with the victim and ends with the fraudster.

Banking law must pay equal attention to the account through which the fraudster receives value.

Money-mule accounts are the bridge between deception and monetisation. RBI described the problem as early as 2010, explaining that individuals may be recruited to receive transfers and forward funds, often for commission, and advising banks to adhere strictly to KYC, AML and transaction-monitoring requirements. Current KYC directions continue to require diligence and meticulous monitoring to identify accounts operated as money mules and to take appropriate action, including suspicious-transaction reporting.6 7

These obligations have profound implications for loss allocation.

Consider a victim who is deceived into transferring ₹20 lakh to an account opened days earlier with limited economic history. Within minutes the beneficiary account receives transfers from several unrelated persons, moves the funds into multiple downstream accounts, and retains only a small balance. If the originating bank had no realistic basis to identify the fraud before payment but the beneficiary bank's systems failed to respond to obvious mule indicators, a rule placing the entire loss on the victim misallocates responsibility.

The beneficiary bank is not necessarily at fault merely because its customer commits fraud. Banks cannot guarantee the honesty of every account holder. The relevant question is whether the institution complied with reasonable onboarding, monitoring and intervention duties.

This suggests a distinction between "recipient wrongdoing" and "recipient-bank control failure." The first is criminal conduct by the account holder. The second is a regulatory question concerning whether the bank's systems reasonably should have detected or interrupted the account's misuse.

Beneficiary-bank responsibility should therefore be based on evidence rather than strict liability. Relevant factors could include the age of the account, declared occupation and expected turnover, abrupt changes in transaction patterns, multiple unrelated inbound transfers, rapid pass-through behaviour, links to previously flagged accounts, inconsistent KYC data,

6Reserve Bank of India, Master Direction—Know Your Customer (KYC) Direction, 2016, para. 59 (as updated) (requiring diligence and meticulous monitoring to identify money-mule accounts and appropriate action including suspicious-transaction reporting).
7Reserve Bank of India, Operation of Bank Accounts—Money Mules, supra note 2.

repeated complaints, device or network anomalies, and failure to act on risk alerts.

The growth of data-sharing and mule-detection tools strengthens the case for such duties. In February 2026 the Supreme Court recorded that banks were using AI/ML tools for fraud-risk management, that a significant number had adopted MuleHunter.AI, and that RBI and the Indian Cybercrime Coordination Centre were working on suspect-registry data sharing. Once the regulatory ecosystem possesses better detection capacity, legal expectations should evolve accordingly.8

The central principle is straightforward: an institution should not bear liability merely because fraud passed through it, but neither should it enjoy categorical immunity when it ignored risks that its regulatory obligations required it to monitor.

V. DIGITAL ARREST AS A STRESS TEST, NOT A SEPARATE LIABILITY UNIVERSE

"Digital arrest" scams provide a vivid illustration of fraudulently induced payment, but the liability framework should not be limited to that label.

The scam typically combines impersonation of public authority, manufactured urgency, surveillance-like control, threats of arrest or prosecution, fabricated documents and instructions to transfer funds for supposed verification or protection. The victim may remain on a video call for hours, may be told not to communicate with family, and may be pressured to liquidate deposits or borrow money.

The distinctive feature is not simply deception. It is simulated state coercion.

Yet from a payment-law perspective, the transfer resembles other fraudulently induced payments. The customer may authenticate it personally. The bank's system may therefore record a valid instruction. The recipient account may nevertheless exhibit mule characteristics. The funds may then be rapidly layered.

The Supreme Court's intervention is important because it shifts attention from criminalisation alone to prevention, recovery and allocation. The August 4, 2026 order directed consultation

8In Re: Victims of Digital Arrest Related to Forged Documents, order dated Feb. 9, 2026, supra note 4.

with banks and intermediaries regarding technological steps to prevent digital arrest, recover defrauded amounts, assist investigation and comply with legal obligations. It also directed discussion of a shared-liability and victim-compensation framework.9

This is a significant conceptual development. It recognises that punishing fraudsters after the money disappears is not an adequate consumer-protection strategy.

The same reasoning applies to investment impersonation, romance-investment fraud, fake customer-care scams, business-email compromise, invoice redirection and other socially engineered transfers. A durable legal framework should therefore define the protected category functionally: payments induced by material fraudulent misrepresentation or coercive impersonation, subject to proof safeguards.

Digital arrest is the present stress test. It should not become the boundary of reform.

VI. THE CASE AGAINST AUTOMATIC CUSTOMER LIABILITY

Three arguments are commonly offered for placing the loss on the customer.

The first is consent: the customer instructed the bank to pay. The second is control: only the customer knew what the fraudster said. The third is moral hazard: reimbursement may reduce consumer caution.

Each argument has force, but none justifies an automatic rule.

Consent in payment systems is operational rather than philosophical. Banks need reliable instructions and cannot investigate the subjective motivation behind every transfer. But legal liability does not have to mirror operational settlement. A payment may be final between institutions while a later reimbursement right exists between customer and bank.

Control is also divided. The customer knows the conversation; the bank knows the transaction history. The beneficiary bank knows the recipient account. The payment system may know device, velocity and network indicators. No single participant possesses the entire picture. Liability rules should therefore encourage each participant to act on the information it uniquely

9 In Re: Victims of Digital Arrest Related to Forged Documents, order dated Aug. 4, 2026, supra note 3.

holds.

Moral hazard is real, but it can be managed through standards of reasonable customer conduct. A victim who knowingly ignores a prominent, transaction-specific warning may bear a greater share of loss than a victim who receives only a generic message. A customer who colludes with a fraudster should receive no protection. A customer who reports promptly should be treated more favourably than one who waits without justification.

The more serious moral hazard may arise on the institutional side. If authenticated fraud is always treated as customer loss, banks have weaker incentives to improve confirmation-of-payee systems, behavioural monitoring, beneficiary risk scoring, cooling periods for extreme anomalies and rapid cross-bank freezes.

A balanced regime should therefore recognise bilateral moral hazard. Consumers need incentives to exercise caution; banks need incentives to design systems that account for predictable human vulnerability.

VII. FROM NEGLIGENCE TO PREVENTABILITY

The language of negligence is familiar but sometimes too blunt for payment fraud. The better organising concept is preventability.

Preventability asks four questions. Who possessed relevant information? Who had the technical or legal capacity to intervene? What would a reasonable intervention have cost? How directly would the intervention have reduced the risk?

This approach avoids retrospective perfectionism. A bank should not be liable simply because fraud occurred. Fraud detection is probabilistic, and false positives can harm legitimate customers. The standard must remain reasonable.

At the same time, preventability is more demanding than mere technical compliance. A bank that mechanically completes KYC at account opening but ignores extraordinary post-opening activity may satisfy paperwork while failing the purpose of monitoring.

For originating banks, preventability may involve transaction-specific friction. A transfer representing 80 per cent of a retiree's liquid balance to a first-time beneficiary after a sudden

change in digital behaviour presents a different risk from a routine monthly transfer. Reasonable controls could include a tailored warning, a short cooling period, confirmation of beneficiary identity, or human review above defined risk thresholds.

For beneficiary banks, preventability focuses on mule-account indicators and rapid dispersal. A system that receives repeated complaints about an account but allows continued outward transfers may present a stronger case for responsibility.

For customers, preventability concerns reasonable response to warnings and protection of credentials. The law should distinguish ordinary human susceptibility to sophisticated deception from gross disregard of clear and specific warnings.

Preventability thus provides a principled basis for shared liability without converting the inquiry into a moral judgment about whether the victim was "careless."

VIII. THE PROPOSED PITAR FRAMEWORK

This article proposes a five-stage model: Prevent, Intervene, Trace, Allocate and Restore.

The framework is sequential because cyber-fraud governance fails when institutions focus only on the final stage. By the time a victim enters litigation, the money may have passed through multiple accounts. Effective loss allocation begins before the transfer.

A. Prevent

Prevention concerns baseline duties before a suspicious transaction occurs. Originating banks should maintain behavioural fraud-detection systems proportionate to the scale and risk of their digital services. Beneficiary banks should maintain robust KYC, expected-activity profiling and mule-account monitoring. Customers should receive intelligible education and transaction-specific warnings rather than repetitive boilerplate.

Prevention should be judged prospectively. The question is whether the institution maintained reasonable systems, not whether hindsight reveals a signal that could theoretically have predicted the fraud.

B. Intervene

Intervention concerns the moment when risk becomes concrete. A high-risk transfer may

justify additional friction. The design challenge is proportionality. Requiring human verification for every UPI payment would destroy the utility of instant payments. Requiring additional verification for a rare, high-value, anomalous transfer may be justified.

The strongest interventions are contextual. "Do not share your OTP" is ineffective where the customer is not sharing an OTP. A warning that says, in substance, "Police, courts and RBI do not require you to move money to a safe account; if someone has instructed you to do so, stop this transfer" directly addresses the scam.

C. Trace

Once fraud is reported, time becomes the dominant variable. The legal system should impose interoperable duties to identify the first beneficiary account, downstream transfers and available balances. Banks should preserve relevant logs and respond rapidly to authorised freeze or hold mechanisms.

The February 2026 Supreme Court proceedings concerning a uniform NCRP-CFCFRMS process reflect the importance of coordinated custody and restoration.10

D. Allocate

If the entire sum is recovered, allocation may be straightforward. The hard cases involve unrecovered loss.

Allocation should examine compliance at each stage. A customer's reasonable conduct, the originating bank's intervention systems, and the beneficiary bank's mule controls should all matter.

E. Restore

Recovered money should not remain immobilised for months or years merely because criminal proceedings continue. Subject to verification of competing claims and lawful process, victim restoration should occur quickly. The Supreme Court's August 2026 direction that adjudicating authorities ensure expeditious disbursement reflects this principle.11

10 In Re: Victims of Digital Arrest Related to Forged Documents, order dated Feb. 9, 2026, supra note 4.
11 In Re: Victims of Digital Arrest Related to Forged Documents, order dated Aug. 4, 2026, supra note 3.

PITAR is therefore both preventive and remedial. It treats restitution as the endpoint of a continuous governance process.

IX. A SHARED-LIABILITY MATRIX

A shared-liability regime requires predictable presumptions.

The following model could be implemented through RBI regulation, supplemented where necessary by legislation.

Category One: Bank-System Failure

Where the loss results substantially from a bank's security deficiency, failure to implement mandatory controls, or disregard of a confirmed fraud alert, the bank should bear the loss.

Category Two: Third-Party Fraud with Reasonable Customer Conduct

Where the customer is deceived despite acting reasonably, reports promptly, and the transaction displays risk indicators that a regulated entity reasonably should have detected, the customer should presumptively receive reimbursement. Responsibility may then be apportioned among institutions according to control failures.

Category Three: Shared Failure

Where both the customer and one or more institutions materially contributed, liability should be apportioned. For example, a customer may have ignored a specific warning, while the beneficiary bank simultaneously failed to act on repeated mule alerts.

Category Four: Gross Customer Disregard

Where the customer deliberately circumvents clear, specific and repeated fraud warnings, misrepresents the transaction's purpose to the bank, or knowingly transfers funds despite credible notice that the recipient is fraudulent, customer liability may increase substantially.

Category Five: Customer Fraud or Collusion

No reimbursement should be available where the customer participates in the fraud or knowingly makes a false claim.

These categories should operate through rebuttable presumptions rather than rigid percentages. RBI could prescribe default allocations while allowing the Ombudsman or consumer forum to vary them where evidence justifies a different result.

The benefit of presumptions is procedural. Individual victims often lack access to bank risk logs, beneficiary-account history and fraud-detection records. If the consumer must prove precisely what the bank's algorithm should have detected, the right becomes illusory.

Banks should therefore bear the burden of producing relevant control evidence once the customer establishes a prima facie fraudulently induced payment and timely reporting.

X. ORIGINATING-BANK DUTIES

The originating bank is the institution with the longest relationship with the payer. It knows the customer's historical transaction patterns, typical beneficiaries, device history and account balances. That informational position supports several duties.

First, banks should deploy risk-based behavioural monitoring. The purpose is not to block unusual spending as such, but to identify combinations of indicators associated with coercive or socially engineered fraud.

Second, warnings should be specific. Generic cyber-safety messages create warning fatigue. A warning triggered by a first-time high-value beneficiary can ask whether the customer is being instructed by a police officer, regulator, investment adviser or customer-care representative.

Third, banks should develop escalation protocols for extreme anomalies. Human intervention remains valuable where automated signals identify a transaction as exceptionally risky.

Fourth, banks should record the warning and the customer's response. This evidence is essential to fair loss allocation.

Fifth, banks should make fraud reporting immediate and accessible. The 2017 RBI framework already emphasises prompt reporting channels for unauthorised transactions. Fraudulently induced payment rules should adopt the same philosophy.12

12 Reserve Bank of India, Customer Protection—Limiting Liability of Customers in Unauthorised Electronic Banking Transactions, supra note 1.

None of these duties implies that a bank must interrogate every customer. The objective is calibrated friction: the highest intervention where risk is highest.

XI. BENEFICIARY-BANK DUTIES AND MULE-ACCOUNT GOVERNANCE

Beneficiary-bank duties should become a distinct pillar of Indian cyber-fraud law.

KYC is often treated as an onboarding event. In reality, mule risk is dynamic. An account may appear ordinary when opened and become suspicious only after its transaction pattern changes.

The RBI KYC directions' emphasis on monitoring is therefore critical. A liability regime should ask whether the beneficiary bank maintained controls reasonably capable of identifying pass-through behaviour.13

Several indicators deserve attention: rapid movement of nearly all inbound funds; repeated credits from unrelated persons; transaction values inconsistent with declared profile; multiple newly added beneficiaries; unusual cash withdrawals; repeated fraud complaints; links to devices or addresses associated with flagged accounts; and abrupt activity in dormant accounts.

A beneficiary bank that receives a credible fraud notification should also have a duty to preserve remaining funds and cooperate with lawful tracing mechanisms without avoidable delay.

Importantly, the law should protect innocent account holders. Not every mule is equally culpable. Some are knowing participants; others may be deceived job seekers or persons whose accounts are compromised. Freezing and investigation should therefore preserve procedural safeguards.

Bank liability and mule criminal liability are separate questions. A bank can fail in monitoring even if the account holder's criminal intent is uncertain. Conversely, a bank may have complied reasonably even though the account holder is plainly guilty.

Keeping these inquiries separate prevents loss allocation from becoming dependent on the eventual criminal conviction of a mule.

13 Reserve Bank of India, Master Direction—Know Your Customer (KYC) Direction, 2016, supra note 7.

XII. CUSTOMER DUTIES WITHOUT VICTIM BLAMING

Consumer protection loses legitimacy if it treats customers as incapable of responsibility. It also loses legitimacy if it treats sophisticated deception as a moral failure by the victim.

The correct standard is reasonable consumer conduct in context.

Relevant factors should include the specificity of bank warnings, the sophistication of the fraud, whether the criminal impersonated a trusted institution, whether the customer was placed under sustained coercive pressure, whether the customer concealed information from the bank, and how quickly the customer reported after discovering the fraud.

Age or vulnerability should not automatically determine liability, but they may be relevant to the reasonableness of communication and safeguards. Payment systems designed for universal use cannot assume expert knowledge.

The law should also distinguish credential sharing from payment inducement. Telling a fraudster an OTP after a clear warning presents one type of risk. Personally initiating a transfer because a convincing impersonator claims that the money must be moved to prevent arrest presents another.

A fair system should encourage caution without requiring consumers to possess the fraud-intelligence capabilities of regulated institutions.

XIII. RESTITUTION, FREEZING AND THE PROBLEM OF PROCEDURAL DELAY

Cyber-fraud recovery is often discussed as though freezing the money solves the victim's problem. It does not.

A freeze preserves value. Restitution returns it.

Victims may face prolonged uncertainty while funds remain attached across multiple accounts, police stations and jurisdictions. The procedural system must distinguish preservation necessary for evidence from retention that serves no continuing purpose.

The Supreme Court's 2026 directions recognise this concern. The Court called for uniform

processes and later directed adjudicating authorities, including the RBI Ombudsman, consumer fora, courts and law-enforcement agencies, to ensure expeditious disbursement of recovered amounts.14

A statutory or regulatory restoration protocol should contain clear timelines. Once ownership is sufficiently established, no competing claim exists, and evidentiary preservation can be achieved through records rather than continued retention of money, restoration should be the default.

Partial recovery should also be distributed promptly rather than withheld until the entire transaction chain is resolved.

The framework should allow digital documentation of claims, standardised bank certificates, and inter-state coordination. Victims should not have to travel to distant jurisdictions merely because the mule account was opened elsewhere.

Restitution is not charity. It is the restoration of property wrongfully extracted through fraud.

XIV. THE RBI OMBUDSMAN AND CONSUMER FORA

India already possesses institutional mechanisms capable of resolving many payment disputes without ordinary civil litigation.

The Reserve Bank–Integrated Ombudsman Scheme, 2021 consolidated earlier ombudsman mechanisms and provides a centralised grievance-redress framework for regulated entities. Its importance will increase if fraudulently induced payments receive clearer substantive rules.15

The Ombudsman is well placed to decide fact-intensive allocation disputes if it receives access to standardised evidence: transaction logs, warnings shown to the customer, fraud-risk scores, beneficiary-account alerts, reporting timestamps and freeze responses.

Consumer fora also remain relevant where banking service deficiencies are alleged. But inconsistent standards can produce uncertainty.

14 In Re: Victims of Digital Arrest Related to Forged Documents, order dated Aug. 4, 2026, supra note 3.
15 Reserve Bank of India, Reserve Bank—Integrated Ombudsman Scheme, 2021, Ref. CEPD.PRD.No.S873/13.01.001/2021-22 (Nov. 12, 2021).

RBI should therefore issue a binding or highly structured liability code applicable across internal bank grievance systems and Ombudsman proceedings. Consumer adjudication could then evaluate deficiency against a clearer regulatory baseline.

The objective should be quick, evidence-based resolution rather than years of litigation over whether a technically authenticated payment was "authorised."

XV. COMPARATIVE LESSONS: USEFUL BUT NOT TRANSPLANTABLE

India is not the only jurisdiction confronting fraud in which the payer technically authorises the transfer.

Comparative experience demonstrates an important policy shift: payment systems increasingly recognise that strong authentication does not eliminate social-engineering risk. Some jurisdictions have moved toward reimbursement obligations, shared responsibility between sending and receiving institutions, or stronger confirmation and recipient controls.

India should study these approaches, but transplantation must be cautious.

The scale of Indian digital payments, diversity of users, prevalence of small-value transactions, structure of UPI, banking inclusion goals and enforcement capacity differ from other systems. A rule designed for a smaller payment market may create excessive friction if copied mechanically.

The transferable principle is not a particular reimbursement percentage. It is incentive alignment.

Where the payer's bank can detect behavioural anomalies, it should have reason to do so. Where the recipient bank can detect mule activity, it should have reason to do so. Where the consumer receives a meaningful warning, the consumer should have reason to heed it.

A successful Indian model should therefore be indigenous in operation even if comparative experience informs its design.

XVI. OBJECTIONS TO SHARED LIABILITY

A shared-liability framework will attract several objections.

The first is cost. Reimbursement obligations may increase banks' fraud losses and ultimately the price of financial services.

The answer is that fraud already imposes costs; the present system merely concentrates many of them on victims. Liability reform reallocates costs toward actors capable of reducing them. If the rule causes banks to invest in prevention that costs less than the fraud it avoids, the social result is efficient.

The second objection is fraud by customers. Reimbursement systems can be abused.

That risk supports verification, not abandonment. Banks possess extensive transaction and device records. False claims should be investigated, and deliberate fraud should disqualify the claimant.

The third objection is payment friction. Additional checks can undermine instant-payment convenience.

This is why the proposed model is risk-based. Ordinary transactions should remain fast. Exceptional intervention should focus on exceptional risk.

The fourth objection is uncertainty. Apportionment may produce inconsistent outcomes.

Presumptions, evidentiary standards and published Ombudsman decisions can reduce inconsistency.

The fifth objection is that criminal law already provides remedies against fraudsters.

Criminal punishment does not answer the allocation question when the offender is untraceable or insolvent. Consumer-protection law exists precisely because contractual and regulatory relationships can allocate risks that criminal prosecution cannot practically repair.

The sixth objection is personal responsibility. Customers should be responsible for instructions they issue.

Personal responsibility remains relevant, but it should not become a doctrinal shortcut. Modern payment systems are jointly produced by consumers and regulated institutions. Responsibility should reflect that shared architecture.

XVII. IMPLEMENTING THE FRAMEWORK IN INDIA

Implementation should proceed in stages.

First, RBI should define "fraudulently induced electronic payment" for regulatory purposes. The definition should cover a payment initiated or authenticated by a customer where material deception, impersonation or coercive manipulation causes the customer to transfer funds to a person or account the customer would not otherwise have paid.

Second, RBI should establish minimum originating-bank intervention standards for high-risk transactions while permitting technological flexibility.

Third, beneficiary banks should face explicit mule-account monitoring and post-notification preservation duties tied to existing KYC/AML obligations.

Fourth, a shared-liability matrix should establish rebuttable presumptions. The rules should specify how customer conduct, warning quality, bank control failures and reporting speed affect allocation.

Fifth, evidentiary burdens should reflect information asymmetry. Banks should produce internal records relevant to the disputed control decision.

Sixth, the NCRP/1930 reporting ecosystem and banking systems should support immediate trace-and-hold communication.

Seventh, restoration procedures should impose timelines for returning verified recovered funds.

Eighth, RBI Ombudsman decisions involving significant fraud-allocation questions should be published in anonymised form to develop consistent principles.

Ninth, banks should report aggregate data on fraudulently induced payments, reimbursement outcomes, mule-account detection and recovery rates. Transparency would allow regulators to evaluate whether incentives are working.

Tenth, periodic review should measure both fraud reduction and unintended consequences such as excessive blocking of legitimate payments.

These reforms do not require India to abandon instant payments. They require instant payments to mature from a system designed primarily for speed and authentication into one designed also for adversarial manipulation.

XVIII. A PROPOSED DOCTRINAL TEST FOR ADJUDICATORS

Until a comprehensive regulatory framework is enacted, adjudicators will continue to face disputes under existing law. A structured test can improve consistency.

Step One: Characterise the Transaction

Was the payment genuinely unauthorised, fraudulently induced, or an ordinary authorised commercial payment?

Step Two: Identify the Deception

What material representation, impersonation or coercive technique caused the transfer? Is there credible contemporaneous evidence?

Step Three: Assess Customer Conduct

Did the customer act reasonably in context? What warnings were received? Did the customer knowingly circumvent safeguards?

Step Four: Assess Originating-Bank Controls

Was the transaction anomalous? Did the bank's systems flag it? Were warnings specific and proportionate? Was an available intervention unreasonably omitted?

Step Five: Assess Beneficiary-Bank Controls

Did the recipient account display mule indicators? Had it been previously complained against or flagged? Did the bank respond reasonably to transaction and post-notification risk?

Step Six: Examine Reporting and Recovery

How quickly was the fraud reported? What action was taken by each institution? Could additional loss have been prevented after notification?

Step Seven: Allocate Unrecovered Loss

Responsibility should correspond to causal and preventive contribution, subject to regulatory presumptions.

This test would allow courts and the Ombudsman to move beyond the sterile question, "Who pressed send?"

XIX. WHY LOSS ALLOCATION IS A CYBERSECURITY TOOL

Liability is often treated as an after-the-fact legal issue. In financial systems, it is also a cybersecurity design tool.

Institutions invest where law and economics create incentives. If a bank bears no meaningful cost when mule accounts repeatedly receive fraudulent transfers, mule detection competes internally with other priorities. If an originating bank faces no consequence for ignoring extreme anomalies because the customer authenticated the transfer, behavioural fraud controls may remain underdeveloped.

Conversely, unlimited liability can generate excessive blocking and cost.

The objective is therefore optimal pressure: enough liability to reward prevention, not so much that institutions must prevent every conceivable deception.

Shared liability is particularly suitable because cyber fraud is distributed. The criminal exploits human psychology, telecommunications, digital identity, bank accounts and payment rails. A legal regime focused on only one node will be circumvented.

Loss allocation can force each node to internalise a portion of the risk it can reasonably control.

This is why the Supreme Court's direction to consider shared liability is potentially more significant than a narrow compensation scheme. Properly designed, it can change ex ante behaviour.

XXI. THE PROBLEM OF CONSENT IN A HIGH-VELOCITY PAYMENT SYSTEM

The legal vocabulary of consent becomes unstable when payment instructions are generated

inside an environment deliberately engineered by a fraudster. Contract law ordinarily treats misrepresentation, fraud and coercion as circumstances capable of undermining the quality of consent. Payment systems, by contrast, require operational certainty. A bank cannot routinely pause settlement to determine whether the customer's reasons for paying are genuine. The mistake is to assume that because operational settlement must be fast, ultimate loss allocation must be equally mechanical.

A two-layer model is preferable. At the first layer, the payment instruction may remain operationally valid so that the payment system can function. At the second layer, regulatory law can determine whether the payer is entitled to reimbursement after fraud is established. This separation preserves settlement finality while recognising that authentication does not settle every question of consumer responsibility.

The distinction is familiar in other areas of law. A transaction may be effective between parties while generating a damages claim because of the circumstances in which it was procured. Likewise, a bank may correctly execute an authenticated instruction yet still face a regulatory reimbursement obligation if its own fraud controls fell below required standards.

This approach also prevents doctrinal confusion between "authority to execute" and "allocation of loss." The bank's authority to execute the customer's instruction need not mean that the customer must ultimately absorb every consequence of a criminally induced instruction.

For Indian law, this conceptual separation is especially valuable because it permits incremental reform. RBI need not rewrite the legal nature of every electronic transfer. It can instead create a distinct post-transaction liability regime for verified fraudulently induced payments.

Such a regime should require convincing evidence of inducement. Contemporaneous call records, messages, fabricated notices, complaint timestamps, transaction patterns and law-enforcement records can assist. The objective is not to accept every later assertion of regret, but to recognise a category of fraud that modern payment architecture has made economically significant.

The concept also clarifies the role of customer negligence. A customer may have been deceived and still have acted unreasonably. Those propositions can coexist. Fraudulent inducement should open the door to the allocation inquiry; it should not predetermine the result.

This is a more mature model of consent for high-velocity finance: operational authorisation for settlement, contextual responsibility for loss.

XXII. EVIDENCE, ALGORITHMS AND PROCEDURAL FAIRNESS

Any liability framework that relies on fraud-risk monitoring must confront a practical problem: the relevant evidence is largely controlled by banks.

Customers generally do not know whether a transaction generated an internal risk score, whether the beneficiary account had previous complaints, whether a device was linked to known mule activity, or whether an automated alert was suppressed. Without disclosure duties, a customer may be asked to prove institutional negligence using information the institution alone possesses.

The burden-of-proof principle in RBI's unauthorised-transaction circular provides a useful foundation. For fraudulently induced payments, a similar evidentiary architecture should apply once the customer establishes a prima facie case of fraud and timely reporting.16

Banks should retain and, in disputes, produce a defined fraud-decision record. That record need not disclose proprietary source code. It should identify the material risk signals considered, the warning presented, any alert generated, the action taken, and the reason for overriding or not escalating the transaction.

Algorithmic systems also require procedural discipline. A bank should not be permitted to defend a claim merely by stating that "the system did not flag the transaction." The legal standard is reasonableness, not algorithmic deference. Conversely, a court should not infer negligence simply because an algorithm failed to predict a novel fraud.

Regulators should audit whether models are calibrated to current fraud typologies and whether false-positive rates are managed proportionately. Models should also be evaluated for unequal impact on groups of customers whose legitimate transaction patterns may differ from the majority.

16 Reserve Bank of India, Customer Protection—Limiting Liability of Customers in Unauthorised Electronic Banking Transactions, supra note 1 (placing the burden of proving customer liability in unauthorised electronic banking transactions on the bank).

The same principle applies to MuleHunter.AI and similar tools. Their growing use may improve detection, but liability should not depend on a particular vendor or model. The duty is functional: maintain reasonable systems for identifying mule behaviour.

Procedural fairness also requires a route for innocent account holders to challenge erroneous mule classification. An aggressive anti-fraud system that freezes legitimate livelihoods without effective review can create a different form of harm.

The goal is therefore accountable automation. Fraud detection should inform legal responsibility, not become an opaque substitute for it.

XXIII. PAYMENT-SYSTEM DESIGN AND THE DUTY TO CREATE FRICTION

Digital payments have been designed around the reduction of friction. That objective is generally beneficial. But in cybersecurity, friction is not always inefficiency. Sometimes it is a safety feature.

A seat belt makes entering a vehicle slightly less convenient. Two-factor authentication adds a step to login. A cooling period delays certain financial actions. The relevant question is whether the safety benefit justifies the burden.

Payment regulation should therefore reject the assumption that every additional confirmation is undesirable. The challenge is to place friction intelligently.

Risk-based friction can be triggered by combinations rather than single factors. A first-time beneficiary alone may be ordinary. A high-value transfer alone may be ordinary. A new device alone may be ordinary. But a first-time beneficiary, new device, unusual hour, transfer of most available balance and recipient account with mule indicators may justify intervention.

The intervention itself can be graduated. A low-level risk may trigger a targeted warning. A higher level may require re-authentication after a short delay. An extreme level may prompt a call-back or temporary hold subject to rapid confirmation.

Crucially, warnings should be empirically tested. Banks should measure whether customers understand them and whether they reduce fraud. A legal requirement to display text is not equivalent to an effective warning.

The design principle should be "minimum effective friction." This aligns security with inclusion. Customers retain fast payments for ordinary activity while receiving stronger protection in exceptional circumstances.

Liability rules can accelerate this design evolution. If banks face some responsibility for preventable fraud, investment in intelligent friction becomes economically rational rather than merely reputational.

XXIV. TOWARDS A NATIONAL CYBER-FRAUD LOSS CODE

The long-term solution may require more than amendments to individual circulars. India should consider a consolidated cyber-fraud loss code issued by RBI under its banking and payment-system powers, supported by legislation if necessary.

The code should define transaction categories, establish baseline prevention duties, specify customer responsibilities, regulate beneficiary-bank obligations, prescribe reporting and hold procedures, create allocation presumptions, and establish restoration timelines.

A single code would reduce fragmentation. Today, relevant duties are dispersed across customer-protection circulars, KYC directions, payment-security rules, internal bank policies, ombudsman procedures and law-enforcement processes. Victims should not need specialist knowledge to understand whether they have a remedy.

The code should also create standard terminology. "Unauthorised transaction," "fraudulently induced payment," "money-mule account," "confirmed fraud notification," "recoverable funds" and "reasonable intervention" should have defined meanings.

It should require interoperable bank-to-bank fraud communication. A victim's originating bank should be able to transmit a standardised, authenticated fraud alert to the beneficiary bank immediately, without waiting for informal emails or branch-level escalation.

The code should establish a safe-harbour mechanism for temporary, good-faith debit holds imposed under defined fraud criteria, coupled with review rights for account holders. Banks may otherwise hesitate to act quickly for fear of wrongful-freeze liability.

Finally, the code should require public reporting. Aggregate statistics on fraud type, recovery,

reimbursement, mule detection and dispute outcomes would allow Parliament, RBI and researchers to evaluate whether the framework works.

A national code would not eliminate cyber fraud. It would make the allocation of its consequences more coherent, transparent and preventive.

XXV. CALIBRATING LIABILITY: PROPORTIONALITY, CAUSATION AND SYSTEMIC INCENTIVES

A workable reimbursement regime must distinguish between responsibility for creating the fraud and responsibility for failing to interrupt it. Banks do not cause the initial deception in most socially engineered payment frauds. The criminal does. Yet private law and regulation routinely allocate losses to parties who did not originate wrongdoing when those parties were better positioned to prevent or absorb a foreseeable risk. The relevant inquiry is therefore not moral blame alone. It is the relationship among causation, control, regulatory duty and prevention capacity.

Causation should remain an important limiting principle. A bank should not bear loss merely because a better system can be imagined after the event. The customer should identify a plausible control failure connected to the loss, after which the institution should produce the information necessary to show whether reasonable controls operated. If a transfer was entirely consistent with the customer's history and the beneficiary account displayed no identifiable risk indicators, institutional liability may be weak even though the customer was deceived. By contrast, where the transfer was exceptionally anomalous and the recipient account had already generated fraud complaints, the causal connection between omitted intervention and loss becomes stronger.

Proportionality is equally important. Controls appropriate for a ₹500 transfer are not necessarily appropriate for a ₹25 lakh transfer that empties a savings account. The regulatory standard should permit institutions to scale intervention to risk. This is particularly important in India, where payment systems serve an enormous range of consumers and transaction values. Excessive universal friction could undermine financial inclusion; insufficient targeted friction can expose the same consumers to catastrophic loss.

Loss severity should also matter to system design even if it does not independently establish

negligence. A small mistaken payment may be recoverable through ordinary complaint processes. A transaction that consumes retirement savings or the proceeds of a home sale can produce irreversible household consequences. High-severity transactions justify stronger preventive architecture because the cost of a false positive is comparatively small relative to the potential harm.

The allocation regime should recognise the difference between first-party and system-wide learning. An individual consumer may encounter a sophisticated scam once in a lifetime. A national bank may observe thousands of related attempts. That difference gives institutions a superior capacity to identify patterns, update warnings and redesign controls. Legal rules should harness that informational advantage rather than assume that every customer can independently keep pace with rapidly changing fraud typologies.

At the same time, banks should receive regulatory credit for demonstrably effective prevention. A shared-liability framework should not operate only as punishment. Institutions that implement strong confirmation systems, rapid mule detection, effective customer warnings and cross-bank response mechanisms should face lower exposure where those controls function reasonably. This creates competition around safety rather than merely around transaction speed.

Insurance may also play a supporting role. Banks can distribute fraud risk across large customer bases more efficiently than individual victims can absorb catastrophic losses, and specialised insurance products may help institutions manage residual exposure. But insurance should not substitute for prevention. A regime that simply transfers bank liability to insurers without changing control incentives would leave the underlying vulnerability intact.

Finally, allocation rules should be periodically recalibrated. Fraud methods evolve in response to regulation. Criminals will test thresholds, adapt transaction values, recruit different kinds of mules and imitate new institutions. RBI should therefore review reimbursement and detection data at regular intervals and adjust presumptions when evidence shows that a rule is either ineffective or producing excessive disruption.

The central regulatory objective is not to identify a permanently innocent class and a permanently responsible class. It is to distribute duties so that each participant has an incentive to reduce the portion of risk it can realistically control. Consumers control some behaviours.

Originating banks control transaction interfaces and behavioural monitoring. Beneficiary banks control account onboarding and transaction surveillance. Payment-system operators control network rules and information flows. Law enforcement controls investigative and freezing mechanisms. A coherent regime aligns all of them.

This approach also answers the concern that shared liability is inherently vague. Liability becomes predictable when the law specifies duties, records compliance, establishes rebuttable presumptions and links financial responsibility to identifiable control failures. The alternative—placing almost all authenticated-fraud loss on the victim—may be simpler, but simplicity is not the same as fairness or efficiency. A mature digital-payment system should prefer principled allocation over convenient allocation.

XXVI. CONCLUSION

India's digital-payment success should not be measured only by transaction speed, volume and accessibility. A mature payment system must also decide fairly what happens when speed and trust are weaponised against its users.

The traditional authorised/unauthorised distinction remains necessary, but it is no longer sufficient. Social engineering allows criminals to manufacture apparently valid instructions while avoiding direct technical compromise. When the law treats authentication as conclusive allocation of risk, it mistakes a security mechanism for a theory of responsibility.

The better approach is shared preventability.

Customers should exercise reasonable caution. Originating banks should identify exceptional behavioural risk and provide meaningful intervention. Beneficiary banks should take mule-account governance seriously as an ongoing monitoring duty. Reporting systems should trace funds in real time. Recovered money should return to verified victims without unnecessary procedural delay. Unrecovered losses should be allocated according to evidence of preventable failure rather than a single fact about who clicked the payment button.

The proposed PITAR framework—Prevent, Intervene, Trace, Allocate and Restore—offers a structure for that transition. It does not guarantee reimbursement for every victim and does not make banks insurers against all deception. It instead recognises that modern cyber-fraud prevention is a shared institutional function.

The Supreme Court's ongoing digital-arrest proceedings have created an unusual regulatory opportunity. By directing consideration of shared liability, technological prevention, mule-account controls and victim compensation, the Court has placed the central question squarely before India's financial system: after the money is gone, who should bear the loss, and why?

The durable answer should not depend on sympathy, technological formalism or institutional convenience. It should depend on capability, reasonable prevention, causal contribution and prompt response.

If India adopts that principle, liability law can do more than compensate yesterday's victim. It can redesign incentives to prevent tomorrow's fraud.