



FROM AUTHORISATION TO AGENCY: REGULATING MACHINE-EXERCISED PAYMENT AUTHORITY IN INDIA

Akshat Nigam¹, Arushi Rajagopala², Lavya Bhasin³, Yohann Titus⁴

¹ B.Sc. LL.B. (Hons.), Semester VIII The National Law Institute University, Bhopal, India

² B.Sc. LL.B. (Hons.), Semester VIII The National Law Institute University, Bhopal, India

³ B.Sc. LL.B. (Hons.), Semester VIII The National Law Institute University, Bhopal, India

⁴ B.Sc. LL.B. (Hons.), Semester VIII The National Law Institute University, Bhopal, India

Corresponding author:- Lavya Bhasin

Email: [lavyabhasin.bscllb@nliu.ac.in]

Phone: [9906006363]

ABSTRACT :

Agentic payment systems alter a familiar assumption in digital payments: that the person whose credentials authenticate a payment is also the person who has decided to make that particular payment. India's payment architecture already permits substantial forms of prior and delegated authority. UPI Circle permits full delegation to another user within defined limits, UPI Reserve Pay allows a customer to block funds for multiple future debits to a merchant. The legal difficulty therefore is not that Indian law requires a human to approve every digital payment.

The more precise difficulty is that agentic AI can be given authority to decide the transaction itself. The customer may authenticate a delegation, while the agent later selects the merchant, product, timing and transaction amount within a mandate. This separates four questions that existing payment regulation tends to connect - whether the customer was authenticated, whether the mandate was validly authorized, whether the agent was authorised to act, and whether the individual transaction fell within that authority. A fifth question follows - who bears the loss when an authorised agent makes a disputed or out-of-scope decision?

This article examines these questions through India's current authentication framework, e-mandates, UPI Reserve Pay, UPI Circle and the 2026 customer-liability amendments. It argues that existing mechanisms can accommodate bounded forms of agentic payment, but do not yet provide a sufficiently explicit legal architecture for machine-exercised payment authority. It proposes a three-tier framework distinguishing automated execution, bounded agentic authority and open-ended agentic authority, supported by verifiable mandates, agent identity, transaction-level decision records, evidentiary rules and liability allocation according to the source of failure. The objective is not to create a separate body of law for every AI payment, but to make explicit the legal significance of the machine's decision between the customer's delegation and the resulting payment.

Keywords: *Agentic AI, digital payments, UPI, authentication, delegated authority, e-mandates, UPI Reserve Pay, UPI Circle, payment liability, machine-exercised payment authority*

1. INTRODUCTION

The payment function of an AI agent is materially different from the use of AI merely to assist a customer. An ordinary recommendation system may identify products or compare prices while leaving the customer to make the final choice and payment. An agentic system can instead receive an objective, interpret it, gather information, make a selection, and initiate the resulting transaction. In the payment context, the relevant capability is therefore not simply automation. It is the delegation of judgment.

An agent may be instructed to find a suitable product, select a merchant, determine when to purchase, and remain within a monetary or categorical constraint. The customer may authenticate the initial delegation and then cease to participate in the transaction. The payment therefore takes a different form from the familiar sequence in which the customer authenticates a particular instruction immediately before the issuer processes it.

This distinction matters because payment regulation does not treat authentication as an abstract proof that money may eventually be spent. The Reserve Bank of India's Authentication Mechanisms for Digital Payment Transactions Directions, 2025 define authentication as the process of validating and confirming the credentials of the customer who is originating the payment instruction.¹ Where a machine later originates the instruction after the customer has authenticated only the delegation, the regulatory object being authenticated and the actor making the subsequent decision no longer coincide in the same way.

¹ Reserve Bank of India, 'Reserve Bank of India (Authentication mechanisms for digital payment transactions) Directions, 2025' RBI/2025-26/79, CO.DPSS.POLC.No S 668/02-14-015/2025-2026 (25 September 2025) para 5(1).

The agentic payment which includes the necessary element of authentication and delegation model does not necessarily make the payment unlawful. It does, however, make the legal analysis more granular. The customer may have authorised an agent without having selected the particular transaction. The agent may be identifiable without every transaction it generates being authorised. A transaction may also satisfy every technical payment check while remaining disputed because the customer alleges that the agent misunderstood or exceeded the mandate.

The central question is therefore not whether AI can technically send a payment instruction. It is whether the law can identify what, precisely, the customer authorised and can allocate the consequences of the machine's subsequent decision.

India's payment architecture already contains several forms of prior or delegated authority. Under the e-mandate framework, the first transaction and the creation, modification and revocation of a mandate are subject to additional-factor requirements, while subsequent recurring transactions fall within an established exemption from the requirement of additional-factor authentication.² UPI AutoPay has subsequently developed into a lifecycle-managed mandate system in which users can view, revoke, pause, modify and port mandates subject to the applicable framework.³

UPI Circle provides an even closer analogy to delegation. Full Delegation allows a primary user to authorise a secondary user to initiate and complete payments from the primary user's account within a predefined monthly limit. NPCI has also extended UPI Circle to IoT devices and software profiles, including AI profiles for limited users, although the 2025 addendum stated that debit transactions using IoT were to be initiated by explicit user action.⁴

UPI Reserve Pay provides another important comparison. It permits a customer to reserve an amount with a merchant and make multiple payments out of that reserved amount, reducing the need for repeated authentication while preserving customer controls over the reserve.⁵ These mechanisms demonstrate why it would be inaccurate to frame the legal problem as a simple prohibition on payments without contemporaneous human approval.

The problem is instead one of degree and object, because AutoPay largely concerns predetermined recurring execution and Reserve Pay concerns a customer-created financial block for future debits. The UPI Circle permits another user to act within delegated limits. Agentic payment systems can go further by delegating the selection of the actual transaction.

2. UNDERSTANDING AGENTIC PAYMENTS: WHAT HAS ACTUALLY CHANGED?

2.1 What distinguishes an agentic payment?

There are, as widely accepted, four states. First is automated execution - the customer has already determined the payment and software merely executes it. "Pay Netflix ₹649 every month" is principally an instruction for repeated execution. The relevant transaction parameters are substantially fixed by the customer.

Second is recurring or conditional payment: the customer creates a standing mandate under which future transactions occur according to predetermined conditions. This is the conceptual territory occupied by e-mandates and AutoPay.

Third is delegated execution: another actor is permitted to initiate a payment within defined limits. UPI Circle is an important example. The delegated person or device has operational authority, but the scope of permissible action is externally bounded.

Fourth is delegated decision-making. The customer says, in substance, "find a suitable badminton racquet for my brother and spend no more than ₹5,000."

The agent must decide which racquet, which merchant, and when to buy it. The customer has authorised a goal and constraints, not a fully specified transaction.

That fourth category is where the legal difficulty becomes distinctive. The payment is not merely the execution of an already-selected transaction. The machine's decision is an intermediate causal act that determines what the transaction will be. This fourth category is what agentic ai payments seeks to achieve, and what sets it apart is rather obviously, the amount of delegation that is done, of reasoning as well, to a certain extent.

The existing architecture should therefore be treated as a foundation rather than an obstacle. E-mandates and AutoPay demonstrate that a customer can validly authorise a series of future payments. Reserve Pay demonstrates a mechanism in which funds can be set aside for multiple future debits. Circle demonstrates that payment authority can be delegated to another actor. The 2025 IoT extension demonstrates that the delegated actor can also be a software or device profile rather than a conventional human user.⁶

3. THE AUTHENTICATION GAP: ARE AGENTIC PAYMENTS ALREADY COVERED BY EXISTING EXCEPTIONS?

3.1 General digital-payment authentication requirements

The RBI's Authentication Mechanisms for Digital Payment Transactions Directions, 2025 apply to payment system providers and payment system participants, including banks and non-bank entities, and to domestic digital payment transactions unless specifically exempted. They require all digital payment transactions to be authenticated by at least two distinct factors unless an exemption applies. At least one factor must be dynamically created or

² Reserve Bank of India, 'Reserve Bank of India (Authentication mechanisms for digital payment transactions) Directions, 2025' RBI/2025-26/79, CO.DPSS.POLC.No S 668/02-14-015/2025-2026 (25 September 2025) Annexure 1, Reserve Bank of India, 'Processing of e-mandate transactions for recurring transactions' DPSS.CO.PD.No 447/02.14.003/2019-20 (21 August 2019).

³ National Payments Corporation of India, 'Enhancement of UPI Autopay' NPCI/UPI/OC-223/2025-26 (7 October 2025).

⁴ National Payments Corporation of India, 'Addendum to NPCI/UPI/2024-25/OC 201 – Introduction of IoT devices & software on UPI Circle' NPCI/UPI/OC-201B/2025-26 (8 October 2025).

⁵ National Payments Corporation of India, 'UPI Reserve Pay' product guidance, National Payments Corporation of India, 'Enablement of UPI Mandate feature of Single Block Multiple Debits' NPCI/UPI/OC-200/2024-25 (31 July 2024), National Payments Corporation of India, 'Enhancements in UPI Single Block Multiple Debits (UPI Reserve Pay)' NPCI/UPI/OC-228/2025-26 (8 October 2025).

⁶ National Payments Corporation of India, 'Addendum to NPCI/UPI/2024-25/OC 201 – Introduction of IoT devices & software on UPI Circle' NPCI/UPI/OC-201B/2025-26 (8 October 2025), NPCI BHIM Services Ltd, 'BHIM Goes Live with UPI Circle Full Delegation' (25 November 2025).

proven for non-card-present transactions.⁷

The Directions are significant for agentic payments for a second reason – they define authentication by reference to the customer who is originating the payment instruction.⁸ This wording fits naturally where the customer generates the payment instruction. It becomes less straightforward where the customer authenticates a delegation but an agent subsequently constructs and originates the payment instruction.

The distinction should not be overstated, as the Directions expressly preserve exemptions, and the existence of an exemption means that contemporaneous authentication is not a universal legal requirement. But it does mean that the legal basis for an agentic payment must be located either in an existing exemption or in a future regulatory mechanism that recognises the delegated instruction.

3.2 What exactly is being authenticated?

The agentic model reveals that “authentication” and “authorisation” are not interchangeable.

Authentication answers who or what supplied the credential being validated. Authorisation answers whether the actor was entitled to perform the act. A mandate answers what authority was granted. Attribution answers whether the transaction can legally be treated as an act performed pursuant to that authority. Liability answers who bears the loss if the transaction cannot be so characterised.

The customer may successfully authenticate a mandate. The agent may be registered and genuine. The payment may be technically processed through a valid rail. Yet the customer may still contend that the agent chose a transaction outside the mandate. The fact that the delegation was authenticated therefore cannot, by itself, settle the status of every subsequent decision.

The e-mandate framework is the strongest existing example of prior authorisation. RBI introduced e-mandates for recurring transactions with additional-factor authentication at registration, modification, revocation and the first transaction. Subsequent recurring transactions can be processed without additional-factor authentication within the applicable framework.⁹ RBI subsequently increased the limit for subsequent transactions without AFA to ₹1 lakh for subscriptions to mutual funds, insurance premium payments and credit-card bill payments.¹⁰

NPCI’s 2025 enhancement to UPI AutoPay also confirms that the framework is not simply a one-time authorisation followed by an uncontrolled debit. Payer-side systems must support lifecycle management, including viewing, revocation, pause, modification and portability, and further, payer-initiated operations require the UPI PIN under the applicable rules.¹¹

AutoPay therefore can accommodate a narrow form of agentic behaviour where the “agent” is only selecting whether and when to execute a transaction already defined by the mandate. It does not necessarily solve the broader problem where the agent selects the merchant or product.

The crucial distinction is predetermined execution versus delegated judgment. If the mandate says “pay this merchant this recurring amount under these conditions,” the subsequent payment is the execution of a previously specified authority. If the mandate says “find an appropriate product and buy it within ₹5,000,” the machine is determining the transaction that the payment system must treat as authorised. The latter cannot simply be assumed to be identical to AutoPay because both involve prior consent.

3.3 UPI Reserve Pay

UPI Reserve Pay, also known as Single Block Multiple Debits, provides an even more interesting case. NPCI describes Reserve Pay as allowing users to reserve a set amount from an eligible source of funds with a merchant and make multiple payments from the reserved amount. Its stated purpose includes reducing the friction associated with multiple OTPs while retaining customer control.¹² The 2025 enhancement extended Reserve Pay to UPI-permitted sources of funds and imposed operational limits including a maximum reserve of ₹10,000 per block and a maximum duration of 90 days.¹³

The mechanism therefore demonstrates that Indian payment infrastructure can support multiple future debits following an initial customer-created reserve. It is a clear answer to the proposition that every future debit necessarily requires a fresh human authentication event.

But Reserve Pay does not answer the separate question of substantive agent authority. The Reserve Pay materials make the reserve merchant-specific and provide the user with controls to create, modify and revoke it.¹⁴ The mechanism therefore resembles bounded payment authority attached to a defined merchant relationship. It does not, by itself, create a general legal status for an AI agent to decide which merchant or product should be paid.

This distinction also matters for the authentication question. The 2025 RBI Authentication Directions list existing exemptions from the two-factor requirement in Annexure 1. That list expressly includes subsequent recurring transactions under the e-mandate framework, selected prepaid instruments, NETC transactions, small-value offline digital payments, small-value contactless card transactions and specified travel bookings. Reserve Pay is not expressly identified by name in that published list.¹⁵ Accordingly, Reserve Pay is not itself an RBI-created authentication exemption. Rather, Reserve Pay is a payment architecture that reduces repeated authentication through a customer-created reserve, and its precise interaction with the authentication

⁷ Reserve Bank of India, ‘Reserve Bank of India (Authentication mechanisms for digital payment transactions) Directions, 2025’ RBI/2025-26/79, CO.DPSS.POLC.No S 668/02-14-015/2025-2026 (25 September 2025) paras 3, 4, 6(a)–(c).

⁸ *ibid* para 5(1).

⁹ Reserve Bank of India, ‘Processing of e-mandate transactions for recurring transactions’ DPSS.CO.PD.No 447/02.14.003/2019-20 (21 August 2019), Reserve Bank of India, ‘Processing of e-mandates for recurring transactions’ RBI/2022-23/73, CO.DPSS.POLC.No S-518/02.14.003/2022-23 (16 June 2022).

¹⁰ Reserve Bank of India, ‘Processing of e-mandates on cards for recurring transactions – Enhancing the limit’ CO.DPSS.POLC.No S-882/02.14.003/2023-24 (12 December 2023), Reserve Bank of India, Annual Report 2023-24, section on payment and settlement systems.

¹¹ National Payments Corporation of India, ‘Enhancement of UPI Autopay’ NPCI/UIP/OC-223/2025-26 (7 October 2025).

¹² National Payments Corporation of India, ‘UPI Reserve Pay’ product guidance, 1, 8–18.

¹³ National Payments Corporation of India, ‘Enhancements in UPI Single Block Multiple Debits (UPI Reserve Pay)’ NPCI/UIP/OC-228/2025-26 (8 October 2025) paras 4–5.

¹⁴ National Payments Corporation of India, ‘UPI Reserve Pay’ product guidance, 8–18.

¹⁵ Reserve Bank of India, ‘Reserve Bank of India (Authentication mechanisms for digital payment transactions) Directions, 2025’ RBI/2025-26/79, CO.DPSS.POLC.No S 668/02-14-015/2025-2026 (25 September 2025) Annexure 1.

Directions must be assessed through the applicable UPI and regulatory framework rather than inferred merely from its technical operation. Therefore, since Reserve Pay is not an identified exemption, it must employ multi-factor authentication in all its transactions.

3.4 P3P and Reserve Pay

Pine Labs' P3P illustrates the next step. Pine Labs describes P3P as an agent-native payment protocol that enables AI agents to initiate and complete payments without a human at the point of transaction. Its June 2026 launch materials state that P3P extends existing UPI mandate infrastructure, including One Time Mandates and Reserve Pay, so that a customer can authorise once and an agent can subsequently browse, select and pay.¹⁶ Pine Labs' current documentation states that Reserve Pay is available as a P3P payment method and that its system uses a customer-approved mandate before agent-triggered payment requests are completed.¹⁷

P3P is therefore evidence of a technological solution to the authentication and execution problem, but it is not itself an RBI regulation. It is an industry protocol built on payment infrastructure that its developer considers capable of supporting autonomous transactions. The legal question is whether the underlying payment rail permits the particular use case and whether the parties have allocated authority and liability adequately.

The distinction is especially important because P3P's architecture adds identity, delegated authorisation, spending controls and auditability around the payment rail. Those features may make the resulting transaction easier to prove and govern, but they do not, without regulatory recognition, alter the statutory or regulatory allocation of liability.

4. THE ATTRIBUTION AND LIABILITY GAP

Indian private law already recognises that authority can be granted to an agent and can be limited in scope. The Indian Contract Act, 1872 defines an agent and principal and recognises express and implied authority, the extent of an agent's authority, and the consequences where an agent exceeds authority.¹⁸ This provides a useful conceptual vocabulary, but payment-system liability is governed by specialised regulatory rules that cannot simply be replaced by ordinary agency doctrine.

The key point is that the customer's grant of authority and the agent's exercise of authority are distinct events. A machine may be an instrument through which the customer's agency is exercised, but the payment system still needs a rule for determining whether the particular output falls within the authority granted.

4.1 Classification under existing liability rules

The current customer-protection framework is organised around categories such as unauthorised transactions, bank negligence or deficiency, customer negligence and third-party breach. Under the RBI's existing framework, customers have zero liability in specified circumstances involving bank negligence and certain third-party breaches, while customer negligence can shift the loss to the customer until reporting. The burden of proving customer liability lies on the bank.¹⁹

The difficulty is that an agentic dispute can occur without fitting comfortably into those categories. Consider a customer who deliberately authorises an agent, the bank properly authenticates the mandate, the agent is genuine, and the payment rail functions correctly. The customer later says "I authorised the agent, but I did not authorise this transaction." There may be no bank authentication failure, no stolen credential, and no conventional third-party breach. The disagreement is about the scope of delegated judgment.

The forthcoming 2026 amendments make the classification architecture more explicit. The Commercial Banks – Responsible Business Conduct Third Amendment Directions, 2026 define authorised and unauthorised electronic banking transactions and introduce detailed categories for bank negligence, customer negligence and third-party breach. They also place the burden of proving customer liability on the bank.²⁰ However, these amendments apply to electronic banking transactions undertaken on or after 1 January 2027. As of October 2026, they are notified but not yet operative.²¹

The forthcoming framework is important because it shows how far RBI is prepared to go in allocating responsibility. It also exposes the remaining problem, which is that none of the core categories is expressly framed around a machine that was validly authorised to act but allegedly made the wrong decision within, or outside, a human mandate.

4.2 Scenario A: the agent acts within its mandate

Consider a hypothetical. A customer tells an AI agent to find a suitable badminton racquet for the customer's brother and spend no more than ₹5,000. The agent compares products, selects a racquet costing ₹4,700 from a merchant, and executes the payment. Authentication and downstream payment controls pass. The customer nevertheless disputes the purchase.

The first question is whether the customer's mandate can constitute authorisation of the transaction. On ordinary agency principles, authority can be expressed or implied and can be bounded by instructions.²² If the agent's decision falls within the objective limits of the mandate, the customer has a

¹⁶ Pine Labs, 'The AI Agent Can Now Pay. Pine Labs Launches P3P — India's First Agentic Payment Protocol Built on UPI' (11 June 2026).

¹⁷ Pine Labs, 'Pine Labs Payment Protocol (P3P)' and 'Pine Labs Payment Protocol – Quickstart' (accessed 7 October 2026).

¹⁸ Indian Contract Act 1872, ss 182, 186–188, 226–228.

¹⁹ 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 (6 July 2017) paras 6–12.

²⁰ Reserve Bank of India, 'Reserve Bank of India (Commercial Banks – Responsible Business Conduct) Third Amendment Directions, 2026' RBI/2026-27/167, DOR.MCS.REC.No 130/01-01-032/2026-27 (24 June 2026) paras 4(3A), 4(20B), 4(20C), 4(26B), 76K–76S.

²¹ *ibid* para 3(2).

²² Indian Contract Act 1872, ss 186–188.

strong basis to be treated as having authorised the agent's act. The customer's later change of mind should not automatically transform a validly delegated transaction into an unauthorised electronic banking transaction.

But the payment framework still needs a way to prove that the agent acted within the mandate. If the issuer sees only a successful payment message, while the agent operator alone possesses the natural-language instruction, product-selection history, rejected alternatives and decision log, the dispute cannot be resolved from the payment record itself.

4.3 Scenario B: the agent exceeds its mandate

Now alter the hypothetical. The customer limits the racquet purchase to ₹5,000, but the agent is compromised or otherwise causes a payment of ₹50,000. The customer never authorised a transaction of that value.

This transaction is much easier to characterise as outside the customer's authority. Yet the liability question remains difficult. If the customer's mandate was properly authenticated and the bank complied with the payment framework, the bank may argue that the customer voluntarily delegated authority. The customer may respond that the delegation was expressly limited and that the machine's excess was not authorised.

Agency law provides a useful private-law analogy. Where an agent exceeds authority, the principal is not necessarily bound merely because the agent acted in the principal's name. But a principal may be bound where the principal's words or conduct induced a third party to believe that the act fell within the agent's authority.²³ That doctrine does not automatically determine the RBI-regulated liability of an issuer, but it demonstrates why "the agent was authorised" is an incomplete legal conclusion. The scope of authority matters.

The payment system therefore needs to distinguish between an authorised agent acting within mandate, an authorised agent exceeding mandate, a compromised agent, a false or unregistered agent, and a bank or payment intermediary failing to enforce a known limit.

4.4 Loss allocation

The customer ordinarily interacts with the issuer or payment app, not with every component in the chain. Under the current RBI framework, the bank bears specified losses arising from its own negligence or deficiency and certain third-party breaches, and the bank carries the burden of proving customer liability.²⁴ The forthcoming 2026 framework retains the basic principle that the bank must establish customer liability and further requires detailed classification and complaint handling.²⁵

That framework is not a simple rule that "the issuer always pays." Rather, it allocates customer-facing responsibility through defined categories. The agentic problem is that the cause of the loss may be a failure in the delegation layer: the agent was genuine but misunderstood the mandate, the agent operator failed to maintain a spending limit, the agent credential was compromised, or the system generated an instruction that did not correspond to the customer's authorised intent.

A regulatory framework for agentic payments should therefore preserve the customer-facing protection of the payment system while permitting the issuer to seek recourse against the party responsible for the relevant failure.

4.6 Evidentiary and burden-of-proof problems

The evidentiary question is central because an agentic payment can be technically valid but substantively disputed.

At minimum, the relevant evidence may include the original mandate, the identity of the agent, the authority granted, monetary, categorical and temporal limits, the natural-language instruction or structured intent, the agent's interpretation of that instruction, the options considered, the selected merchant and product, the decision timestamp, the resulting transaction, authentication records, subsequent modifications, and evidence of compromise.

The current customer-liability framework already recognises the importance of records. The 2017 RBI circular requires banks' communication systems to record the time and date of delivery of alerts and receipt of customer responses, and places the burden of proving customer liability on the bank.²⁶ The 2026 amendment directions similarly contemplate supporting records such as transaction logs when complaints are rejected.²⁷

Agentic systems require a more granular evidentiary chain. A transaction log that records only "₹4,700 paid to Merchant X" cannot establish whether the machine acted within a mandate that said "buy a suitable racquet under ₹5,000." The legally significant record must connect the mandate to the machine's decision.

This suggests a principled evidentiary rule for future regulation. Where an agentic transaction is disputed, the regulated entity or agent operator seeking to rely on the customer's prior delegation should be required to produce the authoritative mandate and the records necessary to demonstrate that the transaction fell within it. This would not necessarily mean imposing liability on the issuer for every agent failure. It would mean placing the evidentiary burden on the party that controls the relevant information and seeks to characterise the transaction as authorised.

That proposal is consistent with, but more specific than, the existing principle that the bank bears the burden of proving customer liability. It is also one of the areas in which "verifiable intent" can acquire legal significance - not because a technical log is itself a source of authority, but because it may provide the evidence needed to prove whether authority existed.

²³ Indian Contract Act 1872, ss 227–228, 237–238.

²⁴ 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 (6 July 2017) paras 6–12.

²⁵ Reserve Bank of India, 'Reserve Bank of India (Commercial Banks – Responsible Business Conduct) Third Amendment Directions, 2026' RBI/2026-27/167, DOR.MCS.REC.No 130/01-01-032/2026-27 (24 June 2026) paras 76K–76R.

²⁶ 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 (6 July 2017) paras 5, 12.

²⁷ Reserve Bank of India, 'Reserve Bank of India (Commercial Banks – Responsible Business Conduct) Third Amendment Directions, 2026' RBI/2026-27/167, DOR.MCS.REC.No 130/01-01-032/2026-27 (24 June 2026) para 76S.

5. EXISTING AND EMERGING REGULATORY RESPONSES

Mastercard's February 2026 announcement of its first authenticated agentic commerce transaction in India provides another model. Mastercard reported transactions involving cards issued by Axis Bank and RBL Bank, payment aggregators including Cashfree Payments, Juspay, PayU and Razorpay, and merchants including Swiggy, Instamart, Vodafone Idea, Tira and Zepto.²⁸ The initiative was presented as part of Mastercard's Agent Pay framework and emphasised tokenisation and authentication.

The deeper legal question can be reduced to three propositions: did the consumer authorise this, did the agent follow the consumer's instructions, and can the parties prove it? Industry "verifiable intent" systems are potentially strongest on the third question. They can generate a record connecting identity, mandate, transaction parameters and execution.

Such records should receive legal significance, but not because a technical protocol can determine legal authority by itself. Their role should be evidentiary. A cryptographically protected record can make it easier to establish what was authorised, what the agent decided and whether a limit was exceeded. The regulatory framework should therefore specify minimum integrity and retention requirements rather than simply declaring a particular industry protocol legally conclusive.

The position changes when the agent is operated by the regulated bank itself. A distinction between customer-operated, merchant-operated and bank-operated agents is analytically important. The latter raises model-risk and operational-risk questions in addition to the payment-authority problem.

RBI's June 2026 draft Guidance on Regulatory Principles for Model Risk Management is relevant, although it remains a draft and should not be treated as binding law. The draft applies to regulated entities and to models including AI/ML systems. It proposes that regulated entities remain accountable for model outcomes, use board-approved model-risk frameworks, undertake independent validation, and consider the degree of reliance and autonomy placed on AI outputs when assessing model risk.²⁹

For a bank-operated payment agent, the logic is particularly strong. If the bank itself supplies the machine making the purchasing decision, it cannot sensibly treat the machine as an independent external actor for all purposes. The bank controls the model, the payment infrastructure and the customer relationship. A future agentic-payment framework should therefore impose a higher governance burden on bank-operated agents, including documented authority rules, model validation, monitoring, incident response, override mechanisms and auditability.

6. TOWARDS A REGULATORY FRAMEWORK FOR AGENTIC PAYMENTS

6.1 Categorising agentic payments

The proposed framework should begin with three categories.

Category I - Automated execution. The customer has substantially predetermined the transaction and the system executes it. Existing e-mandate and recurring-payment rules should ordinarily govern.

Category II - Bounded agentic authority. The customer authorises an agent to select or determine transaction parameters within objectively verifiable limits. Existing delegated-payment infrastructure should be extended to this category, subject to additional requirements for agent identity, mandate integrity and decision records.

Category III - Open-ended agentic authority. The agent receives broad discretion to determine material transaction parameters without sufficiently defined constraints. This category should not automatically receive the benefit of existing mandate exemptions. It should require either a specific regulatory authorisation or a transaction-specific authentication mechanism.

This classification avoids two extremes. It does not treat every automated payment as a novel AI transaction, but it also does not assume that every prior mandate can support unrestricted machine discretion.

6.2 Requirements for valid agentic authorisation

A legally recognised agentic mandate should be capable of being reconstructed after the transaction. At minimum, it should record the customer identity, agent identity, scope of authority, monetary ceiling, temporal validity, merchant or category restrictions where applicable, the authentication used to establish the delegation, and the mechanism through which the agent's decision is linked to the resulting transaction.

The framework should also require integrity of the record. A mandate that can be silently altered after a dispute is not meaningful evidence of authority. Changes should therefore generate immutable or otherwise independently verifiable records showing who changed the mandate, when, and what changed. These are proposed requirements, not existing law. Their purpose is to transform the vague proposition "the customer authorised the AI" into a legally verifiable statement about the content and exercise of authority.

6.3 Possible authentication models

Three regulatory options are available.

First, RBI could extend existing e-mandate exemptions to bounded agentic payments. The advantage is continuity. India already has a mature model for prior authorisation and subsequent execution. The weakness is that e-mandates were designed principally for recurring transactions, not for agents that choose products and merchants.

²⁸ Mastercard, 'Mastercard demonstrates readiness for a new era of AI-powered payments in India with its first authenticated agentic commerce transaction' (2026).

²⁹ Reserve Bank of India, 'Guidance on Regulatory Principles for Model Risk Management, 2026' (draft, 24 June 2026) paras 8–10, 49–54.

Second, RBI could create a specific exemption for verified AI agents. Under this model, the exemption would attach not simply to the transaction but to an agent that satisfies registration, identity, mandate and audit requirements. This would provide regulatory certainty but would require a new supervisory architecture.

Third, RBI could recognise machine credentials as an authentication factor. The 2025 Directions already define authentication factors broadly enough to include possession, knowledge and inherence credentials, but the existing framework remains centred on credentials of the customer originating the payment instruction.³⁰ A machine credential could therefore be recognised only through an explicit regulatory framework that defines its relationship with the customer's delegation.

The third model is conceptually attractive for agentic systems but should not be adopted merely to remove friction. A machine credential authenticates the machine. It does not, without an associated mandate, prove that the machine was authorised to make the particular transaction.

6.4 Recognising machine-exercised payment authority

The central reform should therefore be recognition of a distinct category: payment made pursuant to verified machine-exercised delegated authority.

Such a category should require four elements. First, a verified principal must authenticate the delegation. Second, the agent must possess a persistent and verifiable identity. Third, the authority must be bounded in a manner capable of automated enforcement. Fourth, the transaction must be accompanied by a record sufficient to establish that the agent's decision complied with the mandate.

The regulatory category should not attempt to determine whether an AI model is "intelligent" enough to qualify. The relevant legal fact is the allocation of decision-making authority.

6.5 Evidentiary requirements

Agentic payment systems should maintain a minimum evidentiary chain consisting of the mandate, agent identity, authentication record, applicable limits, the operative instruction or intent, material changes to the instruction, the agent's final decision, the resulting transaction data and an integrity mechanism linking those records.

Retention periods should be calibrated to the applicable payment and dispute regime. The important principle is that the evidence must remain available long enough for a customer dispute to be meaningfully adjudicated.

The evidentiary burden should also follow control of information. Where the bank or regulated payment participant relies on an agentic mandate to establish that a transaction was authorised, it should be capable of producing the mandate and relevant transaction evidence. Where a third-party agent operator controls the decision records, regulation should require that operator to preserve and furnish them to the regulated entity within the applicable dispute process.

6.6 Liability allocation

Liability should follow the source of failure rather than the mere fact that AI was involved.

If authentication or mandate creation failed because of an issuer or regulated payment participant, the existing customer-protection framework should apply.

If an agent was falsely registered or its identity controls failed, responsibility should lie with the entity responsible for registration and verification.

If the mandate itself was improperly created, altered or recorded, the responsible payment or agent operator should bear the resulting loss subject to the applicable customer-protection regime.

If an agent exceeded a clearly enforceable monetary or categorical limit, the framework should treat that as an out-of-authority transaction rather than automatically attributing the transaction to the customer merely because the customer created the agent.

If the agent was compromised after valid delegation, liability should depend on where the security failure occurred and whether the relevant entity maintained the safeguards required by the regulatory framework.

If the issuer or payment system failed to enforce a machine-readable limit that it was required and technically capable of enforcing, the issuer or responsible participant should bear the resulting regulatory liability.

Customer negligence should remain relevant where the customer deliberately or negligently compromised credentials, ignored mandated safeguards or knowingly expanded the agent's authority. But a customer's initial decision to use an agent should not, by itself, constitute negligence when the agent was otherwise operated within the prescribed regulatory safeguards.

7. CONCLUSION

India does not need to begin the legal analysis of agentic payments from the proposition that digital payments require fresh human authentication for every transaction. That proposition is already disproved by the country's own payment architecture. E-mandates permit subsequent recurring payments without additional-factor authentication within the applicable framework. UPI Reserve Pay permits multiple debits against a customer-created reserve. UPI Circle permits delegated payment authority, and NPCI has extended that concept to devices and software profiles.³¹

³⁰ Reserve Bank of India, 'Reserve Bank of India (Authentication mechanisms for digital payment transactions) Directions, 2025' RBI/2025-26/79, CO.DPSS.POLC.No S 668/02-14-015/2025-2026 (25 September 2025) paras 5(6), 6(a)–(c).

³¹ Reserve Bank of India, 'Reserve Bank of India (Authentication mechanisms for digital payment transactions) Directions, 2025' RBI/2025-26/79, CO.DPSS.POLC.No S 668/02-14-015/2025-2026 (25 September 2025) Annexure 1, National Payments Corporation of India, 'UPI Reserve Pay' product guidance, NPCI BHIM Services Ltd, 'BHIM Goes Live with UPI Circle Full Delegation' (25 November 2025).

The legal novelty of agentic payments lies elsewhere. An AI agent can be given authority not merely to execute a payment already selected by the customer, but to decide what the payment should be. That changes the relationship between authentication and authorisation. The customer can authenticate a delegation, the agent can make a decision, the payment system can authenticate or process the resulting instruction, and the customer can still dispute whether that particular decision was within the authority granted.

The central analytical distinction is therefore between authentication of the customer, authorisation of the mandate, authorisation of the agent, authorisation of the individual transaction, attribution of the agent's decision, and allocation of liability. Existing law contains pieces of each question, but does not yet express them as a single framework for machine-exercised payment authority.

The emerging industry response is already moving in that direction. P3P demonstrates a live attempt to combine mandate-based payment rails with agent identity, delegated authorisation and auditable execution. Mastercard has demonstrated authenticated agentic commerce.

The appropriate regulatory response is consequently not a blanket prohibition on autonomous payment. Nor is it a blanket extension of existing authentication exemptions. A better approach is a three-tier framework that treats automated execution, bounded agentic authority and open-ended agentic authority differently. Bounded agentic payments should be capable of operating without fresh transaction-level human approval where the delegation is authenticated, the agent is identifiable, the authority is objectively bounded, and the system can produce a reliable record showing compliance with the mandate. Open-ended authority should require stronger regulatory safeguards.

The principal contribution of this approach is to shift the legal focus from whether a human clicked "pay" to whether the system can establish a legally meaningful chain from human delegation to machine decision to payment. India already knows how to authorise a payment in advance. The unresolved question is how to authenticate, attribute and allocate responsibility for a decision that the customer authorised a machine to make, and this question can only be answered with time, as technology progresses, and the legal framework follows.