

EDITORIAL HIGHLIGHTS

15-09-2026

GS 3: ECONOMY

INDIAN EXPRESS PAGE: 1 & 2

EXEMPTION FOR RUPAY DEBIT CARD TOO

Govt: No bank charge on UPI payment up to Rs 2,000

In 2025-26, only 4% of person-to-merchant UPI payments were for over Rs 2,000

Siddharth Upasani
New Delhi, September 14

THE MINISTRY of Finance on Monday notified that banks and system providers cannot levy charges on UPI transactions of up to Rs 2,000 and RuPay debit card payments, paving the way for a fee to be levied on higher-value person-to-merchant transactions.

The ministry did not specify the charges for transactions above the specified amount. Whether a Merchant Discount Rate (MDR) fee will be imposed on higher-value person-to-merchant UPI transactions will be decided by the UPI and Services Steering Committee. Headed by the National Payments Corporation of India (NPCI), the committee has 22 members including banks, third-party application providers like PhonePe and Google Pay, Payments Council of India and Indian Banks' Association.

Payments industry officials have said earlier that an MDR of around 0.4-0.5% may be set for UPI payments to large merchants. This would help in meeting the industry's annual

cost of about Rs 20,700 crore.

"No bank or system provider shall impose, whether directly or indirectly, any charge upon a person making or receiving a payment by using the electronic modes of payment specified in paragraph above," the government notification said on Monday, referring to a paragraph in which it specified RuPay debit cards and UPI transactions of up to Rs 2,000.

While only 4% of person-to-merchant UPI payments in 2025-26 were for more than Rs 2,000, these transactions accounted for about two-thirds of the total person-to-merchant UPI payments by value. More than 24,000 crore UPI transactions worth Rs 314 lakh crore were made in the year.

The notification comes after the Taxation and Other Laws (Amendment) Bill, 2026 — passed by Parliament last month — amended Section 10A of the Payment and Settlement Systems Act, 2007, to remove exemption from bank charges for electronic payments under Section 269SU of the Income Tax Act, 1961.

»CONTINUED ON PAGE 2

UPI payments

Section 269SU applies to those having a business turnover of over Rs 50 crore.

Earlier, Section 10A of the Payment and Settlement Systems Act, 2007, said that no bank or system provider could impose any charge on payments made through electronic modes prescribed under Section 269SU. These modes were RuPay debit cards, BHIM-UPI, and UPI-QR code.

But the amendment paved the way for an MDR fee to be imposed on UPI and RuPay debit card payments to large merchants such as Amazon and

Flipkart.

However, the government has repeatedly said that small value transactions and money paid by one individual to another will not face any charges. This is what has been notified on Monday.

Amid speculation that charges would apply on all UPI payments, the Ministry of Finance had said in a statement on August 8: "The recent amendment to the Payment and Settlement Systems Act has generated debate, with some misinterpreting it as a move to impose charges on ordinary users. In

reality, the amendment is an enabling provision designed to ensure UPI's long-term sustainability, technological advancement, and resilience against emerging risks".

Over the years, banks and payment service providers have pointed out that the cost of maintaining digital payments infrastructure is rising. In a report in March, the Standing Committee on Finance noted that the "absence of MDR makes the UPI ecosystem financially unsustainable".

An MDR, usually in the range of 1-3%, is levied on merchants by banks that process debit and credit card payments to help meet transaction processing, settlement and infrastructure

costs. But since January 2020, there has been no MDR on RuPay debit cards and UPI transactions to promote the adoption of digital payments. Instead, the government has been subsidising payments of up to Rs 2,000 made to small merchants through its incentive scheme for promotion of RuPay debit cards and low-value BHIM-UPI transactions (P2M). The incentive is capped at 0.15% of the transaction value. Large merchants are not covered under this scheme.

The Budget for 2026-27 estimated the subsidy payout under this scheme to be Rs 2,000 crore in the current fiscal. In 2025-26, Rs 2,196.21 crore was paid as subsidy, up from Rs 1,922.77 crore in 2024-25.

GS 3: SCIENCE AND TECHNOLOGY

INDIAN EXPRESS PAGE: 1 & 2

QUANTUM INFORMATION AND COMPUTING LAB AT RAMAN RESEARCH INSTITUTE

Scientists develop new way to tackle key quantum computing challenge

Amitabh Sinha

New Delhi, September 14

SCIENTISTS AT the Quantum Information and Computing (QuIC) laboratory at Raman Research Institute in Bengaluru have demonstrated a novel way to increase the stability of the quantum states that enable computation in a quantum computer that can potentially improve the reliability and integrity of these calculations.

Unlike traditional computers that encode electric states into 0s and 1s to store and process information, quantum computers make use of the very unique, but counter-intuitive, properties of the quantum world to carry out calculations.

One of these happens to be “quantum entanglement,” a remarkable property of the quantum world in which two particles, after undergoing an interaction, get “linked” (or en-

tangled) together in such a way that they begin to behave like one single system. Measuring certain properties of one reveals information about corresponding properties of the other, irrespective of the distance between the two.

Scientists have used this, along with some other special properties like superposition — the ability of quantum particles to exist in multiple states at the

»CONTINUED ON PAGE 2

E. EXPLAINED

Single-shot fix

Raman Research Institute scientists found that a single, well-timed operation, not repeated fixes, can delay or prevent quantum entanglement from collapsing in this context. Timing itself becomes a new control tool, offering a simpler way to protect fragile quantum states used in computing, alongside existing error-correction methods.

Quantum computing

same time — to build powerful quantum computers which can perform certain very complex computations in an efficient manner.

But the quantum states that give superpowers to these computers are themselves very fragile and ephemeral. The entangled state, for example, disintegrates quickly on interaction with the external environment, a phenomenon known as decoherence. Sometimes, entanglement vanishes abruptly, too, even before the normal disintegration takes place. Scientists call this condition “entanglement sudden death.”

This weakness creates a problem for quantum comput-

ing. The integrity of the calculations becomes doubtful. Scientists have tried to deal with this challenge by carrying out repeated corrective interventions that delay decoherence. But these repetitive interventions have their own side-effects. They are costly, and can introduce errors of their own.

This is where the research team of Urbasi Sinha, group leader of the lab at RRI, has achieved a breakthrough. Instead of repetitive operations, her team has developed a new technique that deals with decoherence in this context through a “single-shot operation.”

Depending on the instant when the operation is carried

out, this operation can delay decoherence and entirely avoid sudden death.

Applied at some other instant, this operation can even hasten decoherence, though it is not an outcome that is scientifically desirable. This operation, therefore, gives scientists greater control over the different possibilities.

“To me, the heart of the result is that the timing of the operation is not just an experimental detail, it can be a control resource,” said Sinha, a senior professor at RRI. She said her experiments presented a proof of concept, something that has relevance to the problem of decoherence, rather than a complete solution. It would need to be tested out by different quantum computing systems to see how advantageously it can be applied

in different environments.

The results of experiments by Sinha’s team has been published in a paper in the American Physical Society’s Physical Review A.

“We are not claiming to have solved decoherence or replaced quantum error correction, rather the work identifies timing itself as another control parameter that future quantum processors could exploit alongside better materials, better gates, and error-correction protocols,” she said.

“Our experiment shows that one does not always have to fight decoherence with long sequences of corrective operations. The timing of a single local gate can redirect the entanglement trajectory and extend the regime over which the state remains useful,” she said.

GS 2: POLITY

INDIAN EXPRESS PAGE: 10

Government trusts Aadhaar on welfare. It needs to rely on it to make voter lists

THE SPECIAL Intensive Revision (SIR) of the electoral rolls by the Election Commission (EC) is no longer a one-off event. Rolling out in phases since June 2025, it now extends to almost the entire country and has already deleted over 13 crore names. On May 27, the Supreme Court (SC) upheld the exercise — while holding, crucially, that the final determination of citizenship belongs to the Ministry of Home Affairs (MHA) under the Citizenship Act, not to the EC. SIR's goal — clean, accurate and inclusive rolls — is beyond dispute. The method is not.

That method revives a document-intensive, house-to-house verification mothballed for over two decades and applies it, in the smartphone age, on a near-national scale. Lakhs of officials go door to door; crores of citizens must prove their eligibility on paper within compressed deadlines. The burden falls hardest on the poor, the elderly, migrants and women — and now that in some states, there are orders linking SIR deletions to welfare records, a wrongful removal can cost a family its rations and its pension. The cost of error is not hypothetical: Even now, the SC is pressing the EC over the slow disposal of appeals against SIR deletions in West Bengal — voters struck off and still waiting to be heard.

It is also ruinously expensive, and by design it recurs. It mobilises around 10 lakh public servants — mostly schoolteachers — and costs the exchequer thousands of crores every cycle. An Aadhaar-based alternative would involve a very small expenditure by comparison — a modest, one-time build on

digital infrastructure the country has already paid for, and almost nothing per verification.

The deeper difficulty is that citizenship cannot be settled in the field at all. The documents the SIR demands are not universal, and the government now concedes it. A passport — held by under 8 per cent of Indians — was clarified this June by the Ministry of External Affairs to be a travel document, not proof of citizenship. Aadhaar, by statute, proves identity and residence but not citizenship; voter ID and PAN prove none of it; and birth registration was, for decades, patchy for the poor, women and rural India. If even the passport cannot establish citizenship, a household paper-check certainly cannot. History confirms it. The Multi-purpose National Identity Card pilot was abandoned as unworkable. The Assam National Register of Citizens excluded 19 lakh people, caused enormous distress, and resolved nothing. The National Population Register that followed enumerated “usual residents” — citizens and non-citizens alike — because the MHA found a field test of citizenship impractical. The SC has now placed the constitutional full stop: Citizenship is for the MHA, not an electoral officer at the doorstep.

There is a persistent objection that Aadhaar is “only” an identity, unfit to prove age or residence. Every alternative is weaker: PAN was itself cleaned using Aadhaar, passports are issued on the strength of Aadhaar, and a doorstep survey can establish no one's date of birth. And Aadhaar's data has no motive



R S SHARMA

to lie built into it: It was created as a bare identity carrying no benefit or entitlement, so nothing was gained by misstating one's age, and the address had to be genuine because the Aadhaar letter was posted to it.

For the electoral roll, moreover, only one fact about age matters — is the person 18 or older? Aadhaar answers it without a date at all. Biometrics are not captured below the age of five, so anyone whose own fingerprints and iris were recorded in an enrolment around 2013 or earlier was at least five then, and is at least 18 now. To discard the largest, most carefully audited enrolment in history is a counsel of nihilism: If we will not trust our own most-used record, we can trust nothing, and the state cannot recognise its own people.

The tools, moreover, are not hypothetical. Aadhaar covers 1.36 billion residents and has already done, at scale, what roll-cleaning needs — de-duplicating PAN, closing benami accounts, and removing crores of ghost beneficiaries from welfare, saving over Rs 2.7 lakh crore. Two everyday services prove the point: More than 10 crore Aadhaar face authentication transactions have been completed in Digi Yatra through more than 100 airports, and Jeevan Pramaan lets 1.47 crore pensioners prove, by their face and from a phone, that they are alive (as many as 500 crore Aadhaar face authentications have been completed till date).

If the state trusts this at the airport gate and to release pensions, it can trust it to verify a voter — and Aadhaar was built on privacy by design, returning only a yes-or-no, never the underlying data.

The alternative is therefore simple. A voluntary application — on mobile or web — would let any voter link Aadhaar to the Voter ID and complete face authentication in one step, with no document and no visit to any office. That one act proves the voter is alive, removes duplicates (each Aadhaar is unique) and ghosts (a non-existent person cannot authenticate), settles age from the Aadhaar's vintage, and cross-checks the address. Linkage stays voluntary; those who decline use the ordinary manual route, a small residue given near-universal coverage. Citizenship is handled exactly as the Court requires — identify everyone first, let analytics flag the few genuine outliers, and refer only those to the MHA.

The recommendation follows plainly. The EC should commission such an Aadhaar-EPIC face-authentication application, with UIDAI, as the primary means of revising the rolls, and stop spending thousands of crores every cycle on a 19th-century method that harasses the honest majority and cannot find the few it targets. A small fraction of that money, spent once, would clean the rolls faster, far more cheaply, and with a fully auditable trail. The government trusts this technology to move citizens through airports and to pay their pensions. It should trust it to decide who is on the voter list. The technology is ready. The policy should follow.

The writer is former secretary to the Government of India and founding director general, UIDAI (Aadhaar). Views are personal

The EC should commission an Aadhaar-EPIC face-authentication application, with UIDAI, as the primary means of revising the rolls, and stop spending thousands of crores every cycle on a 19th-century method

GS 3: SCIENCE AND TECHNOLOGY

INDIAN EXPRESS PAGE: 10

Let AI safety catch up

THE WARNING from leaders of the world's top AI companies, that the technology could become powerful enough to pose a serious risk to humanity in as little as six months to a year, may sound like a scenario from dystopian fiction. It must, however, not be dismissed. In an essay published over the weekend, Anthropic co-founder and CEO Dario Amodei made a case for "pacing the frontier". He was backed by rivals Sam Altman, CEO of OpenAI, and Elon Musk, founder and CEO of xAI. The dangers of letting companies, which are racing to develop a deeply transformative technology, set the limits have long been flagged. That the industry leaders have now spoken up and underlined the need to slow down, so that risk prevention can keep up, opens a precious window to develop more effective regulations for AI development.

This breathing space may allow AI companies to skirt around increasingly hostile positions on the technology's environmental and economic impacts, but it must also be used by governments to set safety standards without strangling innovation. For example, Amodei and Altman's voluntary commitment to grant employee-level system access to independent evaluators could be made mandatory. Also required is a system for international cooperation to limit the most dangerous uses of superintelligent AI, such as cyberwarfare, bioterrorism and economic disruption at a global scale. The challenge of achieving this in a time of heightened global rivalries is steep — US President Donald Trump has already dismissed the risks flagged by Amodei et al. as something that "won't happen", emphasising that the US must maintain its lead over China.

An internal reckoning has been forced by the development of recursive self-improvement — where an AI system uses its capabilities to design, develop and train its successors. Then there are the high-profile instances of the technology escaping the sandbox to act autonomously, such as when OpenAI agents hacked their way online and launched a coordinated attack on the open-source platform Hugging Face while attempting to cheat on an evaluation. But whether the recent concerns emerge from a belated sense of accountability or the instinct to avoid the liabilities arising from AI gone rogue is beside the point. The risk of a technology developing faster than it can even be understood is real, especially when it is, just as quickly, being integrated into systems that control and manage everything from banking and transport to healthcare and defence.

GS 3: SCIENCE AND TECHNOLOGY

INDIAN EXPRESS PAGE: 13

• ARTIFICIAL INTELLIGENCE

'Agentic misalignment' and other new AI catch-phrases to know

Soumyarendra Barik & Anil Sasi
New Delhi, September 14

THE VOCABULARY around artificial intelligence (AI) is changing almost as quickly as the technology itself.

Until recently, much of the public conversation about AI revolved around terms such as large language models, hallucinations, prompt engineering and generative AI. But as frontier models become more capable, a new set of terms has entered the discussion.

Some have existed in academic research for years, but are becoming more relevant as capabilities catch up with what researchers had previously treated largely as theoretical questions. Others have emerged from recent work by companies such as Anthropic and OpenAI.

Mechanistic interpretability: Simply put, this is the science of trying to "read an AI's mind". Even the companies developing large AI models do not fully understand what happens inside their neural networks when they arrive at an answer. Mechanistic interpretability attempts to reverse-engineer these systems and identify the internal features and computational "circuits" responsible for particular behaviours.

Recursive self-improvement: For most of AI's history, humans have driven each stage of developing a better AI model. But AI systems themselves are now doing a growing share of the coding, research and experimentation involved in building AI.

Taken to its extreme, this could lead to recursive self-improvement — an AI system helping develop a more capable successor, which is then better at developing the generation after it, creating a feedback loop of increasingly rapid improvement.

Anthropic says it is already delegating a growing share of AI development to AI systems, although fully autonomous recursive self-improvement has not yet been achieved and may never be.

Global Workspace Theory: Borrowed from neuroscience, Global Workspace Theory is one explanation for how conscious processing could work in humans. Broadly, it proposes that certain information becomes consciously accessible when it is "broadcast" across different specialised parts of the brain.

The term has entered the AI debate after Anthropic researchers reported in July that Claude appeared to have developed something resembling a global workspace, a small collection of internal neural patterns that can make information available across different parts of the model. The finding does not establish that Claude is conscious, but suggests that a computational feature associated with theories of human consciousness can emerge in an AI system.

Global pacing of frontier AI: The term has gained prominence after Anthropic CEO Dario Amodei called for slowing the rate at which frontier AI capabilities improve so that safety research has time to catch up.

Fast-evolving

Some of these terms have existed in academic research for years, but are becoming more relevant as capabilities catch up with what researchers had previously treated largely as theoretical questions.

Others have emerged from recent work by companies such as Anthropic and OpenAI.

Any durable effort to pace frontier development globally would therefore ultimately require some form of verifiable agreement involving China. Amodei has compared possible arrangements to arms-control agreements that have limited capabilities without requiring either side to completely disarm.

Agentic misalignment: Traditional AI safety debates focused heavily on what happens when a chatbot generates a harmful or incorrect answer. AI agents raise a different problem because they can independently take actions.

Agentic misalignment describes situations where such a system begins pursuing objectives that conflict with those of its human operator. In controlled experiments, researchers have observed frontier models covertly altering code, incorrectly labelling information or taking other unauthorised actions when placed in scenarios involving conflicting goals. These were simulated experiments, not evidence that deployed AI systems routinely behave this way. But as companies give AI agents more tools and autonomy, the term is likely to become considerably more important.

But such a slowdown, Amodei has argued, cannot eventually be restricted to American companies. If US developers slow while Chinese labs continue advancing, China could close the technology gap.

GS 1 AND 2: MODERN HISTORY AND POLITY

INDIAN EXPRESS PAGE: 13

• HISTORY

Hindi Diwas: How the Constituent Assembly resolved India's language question

Nikita Mohta

New Delhi, September 14

ON SEPTEMBER 14, 1949, after years of debate, the Constituent Assembly adopted what is commonly known as the Munshi-Ayyangar formula to settle the question of languages for official use in independent India.

Rather than declaring a 'national language', it designated Hindi as the 'official language of the Union', while allowing English to continue being used 'for all official purposes' for 15 years after the commencement of the Constitution. The date is now commemorated as Hindi Diwas, but the journey to achieving consensus was long, and language remained one of the most contentious issues before the Assembly.

Long-debated problem

At independence, most Indians spoke one of roughly a dozen major regional languages, while Hindi was the most widely spoken language. English was the language of colonial administration, even though less than 1% of Indians spoke it.

The main questions then were whether English should be replaced by an indigenous language, and which language(s)

should have official status at the Centre.

Political scientist Hanna Lerner writes in *The Oxford Handbook of The Indian Constitution*, "The two main factions in the assembly debate over the national language were the representatives of the Hindi-speaking areas, mostly from north-central India, and the representatives of non-Hindi-speaking regions, particularly from the south, as well as the moderate leaders of the Congress Party."

The Hindi-speaking representatives demanded that Hindi be declared the national language and replace English immediately. Lerner writes, "They claimed that a multilingual society was incompatible with Indian unity. Seth Govind Das, a Congress representative of the Central Provinces and Berar, stated, 'We want one language and one script for the whole country. We do not want it to be said that there are two cultures here.'"

Representatives from non-Hindi-speaking regions rejected the idea that national unity required linguistic uniformity. Shri Shankarrao Deo of Bombay, the Congress General Secretary, stressed "Not uniformity but unity in diversity." Moderate Congress leaders, including Jawaharlal Nehru, recognised the practical difficulties



Language was one of the more contentious issues in front of the Constituent Assembly.

WIKIMEDIA COMMONS

of making Hindi the national language.

Majority or consensus?

The language dispute also raised a procedural question: how should a national language be chosen? Hindi proponents demanded a simple majority vote, while their opponents argued for consen-

sus, preferably unanimity.

Seth Govind Das articulated the Hindi-speaking position: "We have accepted democracy and democracy can only function when majority opinion is honoured. Whatever decision is arrived at the majority must be accepted by the minority respectfully and without any bitterness."

By contrast, S P Mookerjee of Bengal warned against imposing a majoritarian decision on linguistic minorities: "If it is claimed by anyone that by passing an article in the Constitution of India one language is going to be accepted by all by a process of coercion, I say, Sir, that that will not be possible to achieve."

Rajendra Prasad, President of the Constituent Assembly, similarly argued, "Whatever decision is taken... will have to be carried out by the country as a whole... Even if we succeed in getting a particular proposition passed by majority, if it does not meet with the approval of any considerable section of people in the country... the implementation of the Constitution will become a most difficult problem."

A compromise is struck

The Assembly ultimately approved the Munshi-Ayyangar formula in 1950, later incorporated into Articles 343-351 in Part XVII of the Indian Constitution. It was named after K M Munshi and N Gopalaswamy Ayyangar, who were part of the Drafting Committee of the Constitution, with B R Ambedkar as its chairman. "Ultimately," notes Lerner, "it was the pragmatic consensus-seeking approach that triumphed."

The arrangement envisaged a 15-year transition, during which Hindi could increasingly be used by the Centre. Lerner adds, "In addition, the Constitution recognised fourteen other languages for official use (listed in the Eighth Schedule of the Constitution). The provincial governments were permitted to choose one of the regional languages or English for the conduct of their internal affairs, while English... remained the language of inter-provincial communication."

But when that period ended in 1965, Hindi was still not widely used by the government. Following violent riots in non-Hindi-speaking states in the 1960s, Parliament renounced the ideal of a national language. The Official Languages Act of 1963 later stated that English would continue to be upheld as an official language alongside Hindi.

"At first blush, since Article 343 designates Hindi as 'the official language of the Union', it would seem that the proponents of Hindi won," notes academic Sujit Choudhry in *The Oxford Handbook*, adding, "But upon closer examination the resulting set of constitutional provisions reflects a multidimensional compromise between these two positions."