



# THE RISE OF AGENTIC AI IN PROCURE-TO-PAY: FROM BASIC AUTOMATION TO AUTONOMOUS FINANCE AGENTS

## Abstract

The procure-to-pay cycle has been automated in successive waves, each one more sophisticated than the last. It has eliminated manual data entry, routed approvals at scale, and accelerated invoice processing across global operations. What has remained out of reach, until now, is comprehensive adaptability. It still remains unable to handle novel conditions, reason across incomplete information, and act without waiting for a human to resolve every exception. Agentic AI presents an elegant solution to this shortcoming. Autonomous finance agents can perceive conditions, reason across data sources, adapt to new scenarios, and coordinate with other agents to complete end-to-end workflows. Agentic AI in procure-to-pay makes continuous, intelligent P2P process automation achievable.



Rule-based systems require exhaustive scenario mapping. They work precisely, and only within the conditions they were programmed for. When stakeholders submit an invoice in an unfamiliar format or a purchase order with a pricing deviation, the rule hits a constraint, and the exception lands in a human review queue. Robotic process automation improved

throughput on high-volume, predictable tasks but inherited the same brittleness. It mirrors human screen interactions without a semantic understanding of what it is doing. Something as simple as a UI change in a supplier portal could derail the entire process. The deeper problem is data latency. In most global P2P environments, ERPs

synchronise overnight. Regional finance systems push data into central platforms on set schedules. The consequence is that all the decisions in procurement, such as approvals, cash flow forecasting, and supplier risk assessments, are being made on outdated information. The automation may not solve the fundamental problem if it has obsolete data.

## Agentic AI: what it is, and what it is not

The term "agentic AI" is being applied to a wide range of products, and not all of them deserve it. The popularly accepted taxonomy is:

- Workflow automation follows predefined rules deterministically.
- A copilot assists a human who retains decision authority.
- An agent pursues a goal across multiple steps and chooses some of its own actions.

Most of what is marketed as agentic finance AI is copilot-grade or smart automation. This distinction is critical for executives, because the evaluation questions differ depending on which category a tool belongs to.

### **An autonomous accounts payable agent has four distinguishing capabilities.**

- It operates independently without constant oversight.
- It evaluates complex conditions and selects responses.
- It decomposes broad objectives into ordered task sequences and modifies those sequences when conditions change.
- It accumulates learning from operational experience to improve future performance.

These properties, combined with the conversational directive comprehension that large language models have made broadly accessible, position agentic

systems as something qualitatively different from their predecessor technologies.

However, certain categories of P2P work should always retain a human authorisation checkpoint:

- Releasing payments
- Posting material journal entries
- Changing vendor banking details.

These are irreversible actions that are precisely the targets of fraud. An agent can prepare, validate, and route these actions, but a human should authorise them. The distinction between preparation and authorisation is the correct architecture for autonomous finance at enterprise scale.

## Agentic AI across the P2P cycle

Industry research on agentic AI in P2P maps six core application areas where autonomous agents deliver measurable operational value. Each represents a distinct improvement over what rule-based automation or RPA could achieve.

### Multi-document invoice matching

Traditional three-way matching required manual comparison of purchase orders, goods receipts, and invoices. Autonomous accounts payable agents process all three simultaneously, extract relevant data fields, identify discrepancies, interpret contextual factors like split shipments or approved quantity variances, and route exceptions with explanatory context. Organisations with fully optimised AP automation report fewer duplicate or incorrect payments, and best-in-class AP teams run lower invoice exception rates.

### Anomaly detection and fraud prevention

Autonomous agents monitor transaction flows continuously rather than through periodic batch reviews. They identify signals that humans and rule sets may miss, like payment amounts clustering below approval thresholds, unusual transaction

timing patterns, and duplicate supplier identifiers with slight variations in address data. With fraud today involving some form of AI, static rule-based detection is increasingly inadequate against the methods that fraudsters are deploying.

### Contract compliance monitoring

Agents continuously compare actual transaction data against contract terms, identify price deviations and fulfilment shortfalls, track volume commitments against rebate thresholds, and initiate renewal or re-sourcing workflows.

### Autonomous negotiation and vendor management

AI-led vendor management agents interface with supplier platforms, transmit requirements, and evaluate submissions against multi-criteria frameworks that extend beyond unit price to delivery terms, specification certifications, and historical performance. They also conduct sequential negotiation cycles within defined governance boundaries. If they find transactions exceeding the authorisation thresholds or involving strategic supplier

relationships, they escalate to human review automatically.

### Payment and working capital optimisation

Autonomous payment agents reconcile competing treasury objectives, such as early payment discount capture, available capital retention, and supplier relationship health, dynamically. Simultaneously, the agents adjust disbursement schedules in response to real-time cash position rather than batch-processed approximations. The downstream effect on working capital optimisation is that organisations with fully optimised AP operations report on-time payments.

### Supplier risk monitoring

Rather than periodic vendor assessments, agents monitor supplier financial indicators, operational performance metrics, adverse media signals, and regulatory compliance status continuously, maintaining fluid risk profiles and triggering alerts and risk mitigation recommendations when thresholds are breached.



## The adoption gap and the barriers behind it

The numbers on agentic AI adoption are striking, and the gap between intent and execution is equally striking. Nearly all (99%) organisations plan to put autonomous agents into production, according to KPMG. Only 11% have done so. That gap represents a data, governance, and change management problem.

### Data quality

Data quality is the constraint, uniformly across every source surveyed. Autonomous systems trained on incomplete, inconsistent, or siloed data produce unreliable outputs that erode the user trust on which adoption depends. An agent is only as accurate as the transaction history, supplier records, and contract databases it draws on.

In most enterprises, these are fragmented across systems that were not designed

to interoperate. A mere 2% of companies had sufficient AI guardrails as of 2025, and 95% reported one or more AI incidents as a consequence. No less than 77% of those incidents led to financial losses.

### Model drift

Agents trained on historical procurement data progressively drift out of calibration as market conditions, supplier behaviours, and organisational structures change beyond their training boundaries. Detecting drift requires continuous performance monitoring, and correction requires structured retraining frameworks.

### Explainability

Every time an agent dismisses a supplier proposal, flags an invoice for investigation, or modifies a disbursement schedule, finance professionals and auditors expect justifiable reasoning. Regulatory

frameworks in most jurisdictions are moving toward requiring algorithmic transparency for consequential financial decisions. An explainability deficit jeopardises compliance and control.

### The human dimension

Procurement and AP professionals accustomed to direct transaction oversight resist transferring that authority to systems whose reasoning they cannot follow. Middle management fears losing visibility into decision-making that has historically defined their roles. Implementations that treat this as a change management communication exercise rather than a structural redesign of roles and incentives tend to achieve adoption on paper and resistance in practice.





## A phased approach to autonomous P2P

The path from pilot to full autonomy follows a logical maturity progression, with each stage building the evidence base that justifies the next.

### Pilot stage

The initial deployment is deliberately narrow in scope, limited to low transaction volumes and decisions that can be reversed without consequence. At this stage, agents provide recommendations only, and human operators validate every output before any action is taken. The objective is to confirm technical operability and build user familiarity.

### Assisted operations

As confidence grows, coverage can expand. Agents take on routine analysis, exception flagging, and data aggregation, while humans retain final decision authority at every step. Performance metrics at this stage should focus on recommendation

accuracy and efficiency gains rather than straight-through processing rates, since the value being generated is in the quality of AI-assisted decisions, not in the volume of decisions the AI makes alone.

### Semi-autonomy

Agents are granted independent execution authority for high-confidence, routine transactions that fall within clearly defined parameters. Complex, ambiguous, or high-value cases escalate automatically to human review. The criteria for escalation are documented and auditable. Human oversight shifts from reviewing individual transactions to managing the exceptions the system surfaces.

### Full autonomy

The final stage extends autonomous coverage comprehensively, with periodic governance review replacing transaction-level oversight. At this point,

the governance infrastructure, comprising continuous monitoring, audit logging, model performance tracking, and override mechanisms, becomes the primary mode of human control, and the procurement professional's role is to govern the system's performance rather than its individual outputs.

AI agents' ability to automate tasks is doubling every three to seven months at current adoption rates. The ability of autonomous agents to handle tasks is expanding rapidly, and organisations that establish working deployments now accumulate training data, governance experience, and institutional confidence that later adopters cannot develop. First movers achieve 2.3x ROI on agentic AI investment within thirteen months on average. However, organisations investing at the same point in time but later in adoption maturity achieve far less.

## End note

Agentic AI in procure-to-pay is an urgent competitive differentiator for the organisations that have moved past the pilot stage. The gap between those organisations and those still running rules-based automation is widening on every

operational metric, whether it be invoice exception rates, payment accuracy, cycle times, working capital performance, or fraud loss. Closing that gap permanently requires a data foundation capable

of supporting autonomous systems, a governance framework that does not confuse absence of oversight with efficiency, and a clear position on which decisions stay with humans.

For more information, contact [infosysbpm@infosys.com](mailto:infosysbpm@infosys.com)



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.