International Journal of Advanced Engineering, Management and Science (IJAEMS)
Peer-Reviewed Journal
ISSN: 2454-1311 | Vol-12, Issue-2; Mar-Apr, 2026
Journal Home Page: https://ijaems.com/
DOI: https://dx.doi.org/10.22161/ijaems.122.5

Models and Concepts of AI Agents in Financial Operations for Autonomous Payroll Processing

Shanmuka Siva Varma Chekuri

Data Engineer, American Software Group (ASG), United States, New Jersey

Received: 29 Jan 2026; Received in revised form: 25 Feb 2026; Accepted:02 Mar 2026; Available online: 06 Mar 2026

Abstract— *The study examines models and architectural concepts of AI agents embedded in financial operations to support autonomous payroll processing in multi-region, compliance-sensitive environments. The research focus lies on agentic patterns that orchestrate payroll data flows end-to-end over ACID-compliant lakehouse platforms, feature stores, and metadata-driven orchestration layers. The work synthesizes current literature on AI agents in finance, autonomous decision systems, data lakehouse architectures, and feature-store-centric machine learning pipelines, and combines it with design patterns emerging in production-grade payroll and HR systems. Particular attention is given to deterministic computation graphs, idempotent pipelines, concurrency-safe merge patterns, and real-time observability for reconciliation and audit. The article aims to formulate a conceptual model of an autonomous payroll stack, outline classes of domain-specific agents, and identify their limitations and governance needs. The material is intended for researchers and practitioners in AI, data engineering, and financial systems design.*

Keywords— ***autonomous payroll, AI agents, financial data engineering, lakehouse architecture, feature store, metadata-driven orchestration, idempotent pipelines, compliance automation, multi-agent systems, real-time reconciliation.***

I. INTRODUCTION

Rapid diffusion of AI agents into financial services reshapes how organizations design payment, settlement, and payroll infrastructures. Large institutions gradually replace rule-based batch pipelines with systems where autonomous software agents monitor data streams, enforce policies, and trigger transactions with limited human supervision. Recent studies underline that agentic AI in finance no longer affects only analytics; it structurally changes intermediation, risk transfer, and operational workflows in payments and asset management [1].

Global payroll represents a particularly demanding field for such transformation. Multi-country organizations operate diverse labor codes, tax regimes, benefits schemes, and collective agreements. Errors propagate into banking interfaces, general ledgers, and regulatory reporting, creating exposure that cannot be mitigated purely through manual control. Industry and academic work on generative and autonomous agents in finance points toward architectures where specialized agents continuously validate calculations, monitor anomalies, and maintain regulatory alignment across jurisdictions.

The purpose of the article is to systematize models and concepts of AI agents suited for autonomous payroll processing and to connect them with modern financial data-engineering patterns. The research pursues three tasks:

1) to classify functional archetypes of AI agents that act across the payroll lifecycle, from eligibility determination to payment reconciliation;
2) to relate these archetypes to a reference lakehouse- and feature-store-centric architecture that supports ACID transactions,

schema evolution, and multi-region data harmonization;

3) to assess the limitations, risk factors, and governance requirements that arise when such agents participate in production payroll flows under regulatory scrutiny.

The novelty of the work lies in unifying three lines of research that are usually treated separately: conceptual analyses of AI agents in finance and decentralized markets, surveys of data lakehouse architectures and sustainable financial data platforms, and technical work on feature stores and real-time ML pipelines.The article reinterprets these results through the lens of global payroll, where deterministic behavior, traceability, and explainability are not optional but foundational design requirements.

II. MATERIALS AND METHODS

The study draws on a corpus of recent work dealing with AI agents in finance, autonomous decision systems, and financial data platforms. I. Aldasoro [1] analyses how artificial intelligence, including emerging AI agents, alters key financial system functions and stability concerns, providing a macro-level framing for agentic finance. K. Allam [2] examines cloud-based data hubs and SQL pipelines for real-time financial analytics, illustrating practical paths from monolithic warehouses to more flexible streaming and hub-and-spoke architectures in financial institutions. L. Ante [3] investigates autonomous AI agents in decentralized finance, identifying agent typologies, coordination patterns, and governance tensions that inform the design of autonomous financial workflows. J. de la Rúa Martínez [4] introduces the Hopsworks feature store and describes an FTI (feature–training–inference) architecture that relies on centralized feature governance and mixed offline/online serving, which is directly relevant for payroll-related models. A. A. Harby [5] and co-author present a survey of data lakehouse architectures in Information Systems, synthesizing how lakehouses merge warehouse reliability with data-lake scalability and how ACID transactions, unified metadata, and open table formats contribute to AI-ready platforms. I. Hettiarachchi [6] focuses on generative AI agents in finance and analyzes operational disruption scenarios, including multi-agent systems built on large language models. S. A. Ionescu [7] investigates sustainable data architectures for modern financial institutions, highlighting lakehouse-style designs, hybrid cloud deployment, and governance for resilient financial data infrastructures. W. Liang [8] and co-authors discuss scalable and secure feature stores in real-time ML pipelines, emphasizing feature governance, low-latency serving, and security constraints in production ML. K. U. S. Somarathna [9] proposes an agent-based simulation approach for human resource management as a complex system, providing conceptual foundations for multi-agent reasoning over workforce and payroll-related variables. Z. Yordanova [10] and Y. Hristizov explore the evolution of financial analysis toward AI agents, mapping how autonomous systems assume decision authority in financial workflows.

For the article, comparative analysis is applied across these sources to extract converging patterns of agent design, data-platform architecture, and governance in financial systems. Conceptual modeling is used to define a taxonomy of AI agent types for payroll operations and to map their interactions onto an ACID-compliant lakehouse with an enterprise feature store. Structural–functional analysis supports decomposition of the autonomous payroll lifecycle into discrete, agent-addressable tasks. Synthesis and abstraction form the basis for the reference architecture and for the conceptual figure that integrates findings from the surveyed works. Throughout, the reasoning is aligned with real-world constraints of regulated payroll environments, including idempotent processing, concurrency-safe merge patterns, multi-region data residency, and traceable computation graphs.

III. RESULTS

Research on AI agents in finance indicates that agentic systems progress from decision-support tools to semi-autonomous and fully autonomous operators embedded in financial infrastructure. Aldasoro et al. highlight that generative models and AI agents intersect with core payment and settlement functions, with particular relevance for information acquisition, signal extraction, and execution of rule-driven actions [1]. Yordanova and Hristizov describe a similar trajectory for financial analysis, where systems evolve from static analytics to goal-directed agents that interpret data, propose actions, and implement them within predefined bounds [10]. When these insights are transferred to payroll, they point toward a layered architecture where multiple specialized agents cooperate to maintain payroll correctness, compliance, and timeliness across regions.

First, the literature supports clear differentiation between strategic agents, operational agents, and infrastructure-level agents. Generative AI agents described by Hettiarachchi operate at strategic and tactical levels, synthesizing information, detecting patterns, and proposing interventions in areas such as trading and risk [6]. By analogy, an autonomous payroll environment benefits from strategic agents that monitor long-term cost dynamics, overtime trends, and regional exposure, while operational agents focus on rule execution, calculation validation, and workflow automation. Ante's typology in decentralized finance, which distinguishes AI agents by autonomy degree and governance arrangement, provides a template for defining payroll-specific agent roles across eligibility checks, accrual derivation, tax calculation, payment scheduling, and anomaly triage [3].

Second, data-engineering work on lakehouses and sustainable financial architectures indicates that autonomous agents reach full potential only when grounded in a unified, high-quality data plane. Harby and Zulkernine show that data lakehouses combine warehouse-style ACID guarantees with lake-style scalability and schema flexibility [5]. Ionescu demonstrates similar trends in financial institutions, where hybrid lakehouse designs simplify integration of transactional stores, risk systems, and historical archives while maintaining strong governance [7]. For payroll, such a lakehouse hosts canonical employee, contract, and transaction records across jurisdictions, with slowly changing dimensions (SCD-2) capturing historical versions of employment terms. Concurrency-safe merge patterns and ACID transactions protect against partial updates when agents simultaneously apply corrections, reclassifications, or back-dated adjustments.

Feature-store research adds a further layer to this picture. De la Rúa Martínez and co-authors describe an FTI architecture in which a feature store supplies consistent features to training and inference pipelines, eliminating offline/online skew [4]. Liang et al. emphasize that secure, real-time feature stores standardize feature computation, control access, and support low-latency inference in production ML pipelines [8]. In payroll automation, such a feature store aggregates contract attributes, time-sheet aggregates, historical payments, tax brackets, and compliance signals into reusable features. Validation agents then consume these features deterministically, applying rules and ML-based estimators within a computation graph whose nodes and edges are fully versioned and observable. This architecture supports idempotent pipelines: if an upstream data source sends duplicate or corrected events, agents recompute downstream states without double-paying or corrupting balances.

Figure 1 integrates these strands into a conceptual architecture for autonomous payroll on AI-ready financial infrastructure. The diagram follows the structure of lakehouse and feature-store architectures described by Harby & Zulkernine and de la Rúa Martínez et al., adapted to the payroll domain [4–5].

At the foundation, an ACID-compliant lakehouse consolidates raw HR, time-tracking, benefits, banking, and tax data into harmonized tables partitioned by region and period. Above this plane, the enterprise feature store exposes curated payroll features: gross-to-net factors, regulatory thresholds, eligibility indicators, and audit metrics. A metadata-driven orchestration layer defines deterministic computation graphs that describe how agents traverse states: ingestion, normalization, feature computation, rule evaluation, payment execution, and reconciliation. On top of this layer, several classes of AI agents operate:

- Ingestion and normalization agents align heterogeneous feeds, enforce schema evolution policies, and record lineage for each transformation step. Their behavior reflects practices described in sustainable financial data architectures and real-time financial analytics pipelines [2,7]
- Policy and compliance agents encode labor laws, tax rules, and internal policies as executable rule graphs. Building on ideas from agentic

finance, they interpret regulatory signals, update rule sets, and flag conflicts between jurisdictions [1,6].

- Calculation and anomaly-detection agents compute payroll outputs, compare them against historical baselines, and use ML models trained over feature-store data to identify anomalies such as outlier net pay, misclassified contracts, or missing contributions [4,8,10].
- Reconciliation and audit agents align payroll results with general ledger postings, bank movements, and third-party system data. Insights from work on autonomous agents and decentralized finance—where agents coordinate across multiple ledgers—inform the design of these reconciliation flows [3].

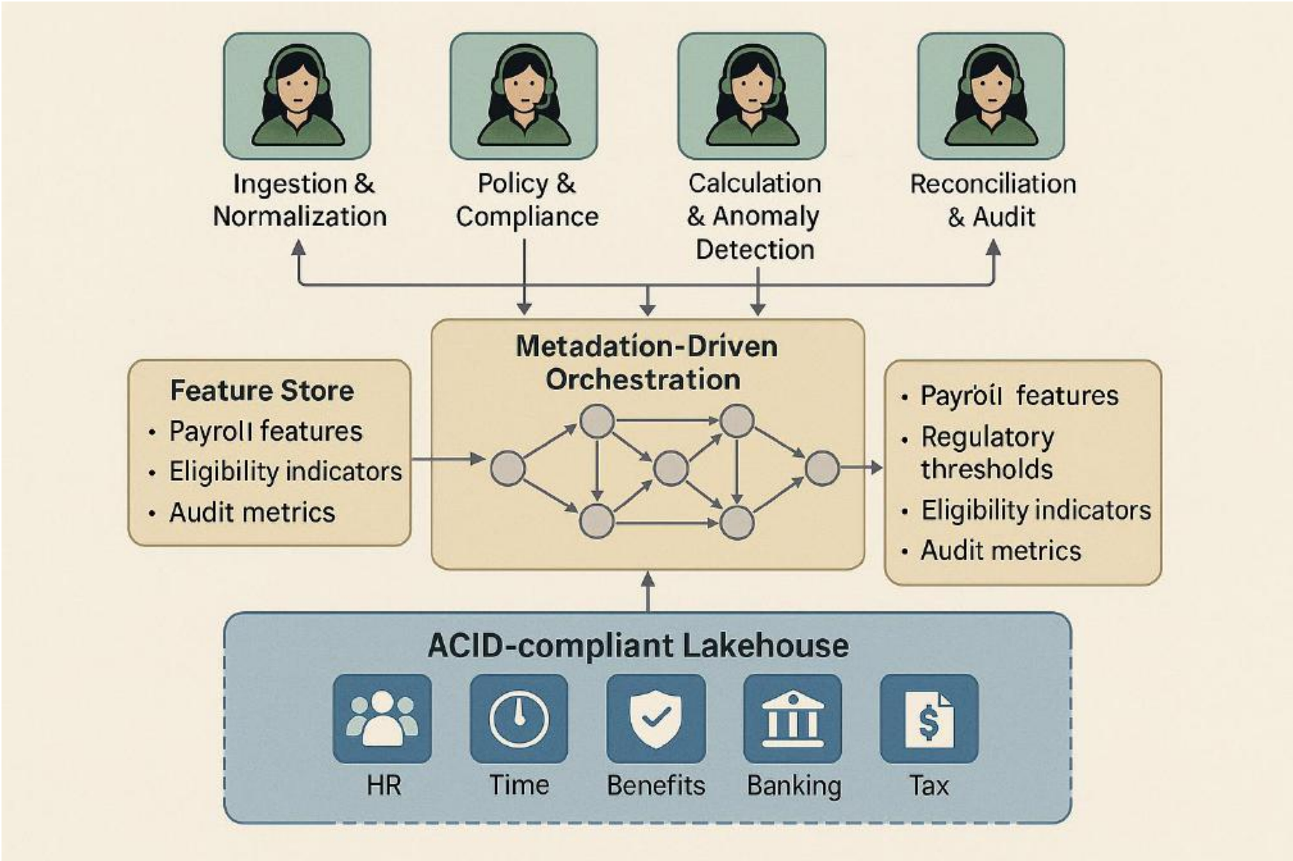

Fig.1. Conceptual multi-agent lakehouse architecture for autonomous payroll processing (compiled by author based on [4–5,7–8]).

Agent-based modeling work in HR, such as the simulation of human resource management strategies by Somarathna, shows how agent interactions and local decision rules generate emergent organizational dynamics [9]. Translating such ideas into payroll, one can model agents representing employees, managers, regulators, and systems, and simulate how modifications to policies, data-quality regimes, or scheduling logic influence error rates, rework, and compliance risk. This perspective motivates design of autonomous payroll agents not as isolated scripts but as entities embedded in a socio-technical system with feedback from human operators and regulators.

The surveyed literature further clarifies that such agentic payroll environments depend on real-time observability and lineage. Harby and Ionescu both emphasize unified metadata layers and catalogues that track table schemas, versions, and access patterns [5,7]. In a payroll setting, every decision made by an agent—such as assigning a tax code or resolving a conflict between overlapping allowances—must be traceable back to data versions, rule definitions, and model parameters in force at the decision time. Feature-store architectures add fine-grained lineage at the level of individual features and their transformation pipelines [4,8]. Together, these elements make post-hoc audit, rollback, and re-simulation feasible, which is central for highly regulated payroll operations.

Finally, several works highlight sustainability, security, and governance as cross-cutting concerns. Ionescu points out that financial institutions seek architectures that balance performance with energy consumption and operational resilience [7]. Liang et al. stress secure access control, encryption, and isolation in feature stores [8]. Aldasoro et al. and Hettiarachchi address systemic risk, regulatory adaptation, and governance for AI agents in finance [1,6]. For autonomous payroll, these insights translate into requirements for least-privilege access for agents, verifiable guardrails around payment initiation, and explicit alignment with responsible-AI frameworks. In combination, the materials enable a coherent conceptual result: an autonomous payroll system constructed as a set of domain-specific agents operating over an AI-ready, lakehouse-based data platform with feature-store-centric ML, deterministic orchestration, and full observability.

IV. DISCUSSION

The conceptual architecture outlined above presents both opportunities and constraints when reconciled with the surveyed literature. Works on AI agents in finance consistently signal a shift from passive analytics toward systems that act as economic participants. Aldasoro et al. describe AI agents as entities that influence price formation, risk distribution, and stability channels [1]. Yordanova and Hristizov extend this picture by arguing that AI agents increasingly undertake analytical, forecasting, and decision-making functions previously delegated to financial analysts [10]. In payroll, a similar migration from advisory dashboards to execution-capable agents is technically feasible but institutionally constrained by labor law, collective bargaining, and organizational risk appetite.

Table 1 summarizes how insights from the reviewed financial-agent literature map onto concrete responsibilities in the payroll lifecycle. The mapping synthesizes themes from generative AI agents in finance, decentralized-finance agents, and financial-analysis agents [1,3,6,10].

Table 1. Functional allocation of AI agents across the payroll lifecycle (based on [1,3,6,10])

Payroll lifecycle stage	Dominant agent capabilities	Illustrative literature link
Data ingestion and normalization	Schema alignment, anomaly detection in source feeds, policy-aware filtering	Generative AI agents for data preparation and monitoring [6,10]
Eligibility and accruals	Rule evaluation, scenario exploration, explanation generation	AI agents as decision-support systems in finance [1,10]
Gross-to-net computation	Deterministic rule execution, conflict resolution between overlapping rules	Autonomous agents coordinating financial operations [3,6]
Compliance and controls	Constraint checking, alert triage, justification logging	Agentic AI for regulatory functions and governance [1,6,10]
Payment execution and routing	Workflow orchestration, exception management, retry logic	Operational agents coordinating financial transactions [3]
Reconciliation and audit	Multi-ledger matching, variance analysis, trace reconstruction	Multi-agent coordination in financial ecosystems [1,3]

The table illustrates that many agent capabilities required for payroll already appear in more general financial domains, though they must be adapted to specific payroll artifacts such as pay periods, earnings codes, and jurisdiction-specific tax constructs. The literature on DeFi agents emphasizes coordination across multiple ledgers and governance frameworks, which parallels the need to reconcile payroll with ERPs, banking systems, and statutory portals [3]. Generative AI work, in turn, suggests that explanation-oriented agents can help bridge the gap between autonomous operation and human oversight by producing narrative justifications for complex decisions [6,10].

The second cluster of findings concerns the underlying data platform. Lakehouse surveys show that successful deployments rely on a convergence of ACID table formats, unified catalogs, and fine-grained governance, which together support both analytical and operational workloads [5,7]. Feature-store work adds another layer by demonstrating how shared, governed feature repositories enable consistent training and inference across batch and streaming contexts [4,8]. Table 2 compares three architectural patterns that emerge from the reviewed materials in relation to autonomous payroll: a traditional warehouse-centric stack, a lakehouse without a feature store, and a lakehouse with an enterprise feature store and metadata-driven orchestration.

Table 2. Architectural patterns for payroll automation and their suitability for AI agents (based on [2,4–5,7–8])

Architecture pattern	Data characteristics	Suitability for autonomous payroll agents
Legacy warehouse-centric ETL	Rigid schema-on-write, nightly batches, limited history of schema changes	Constrains agent autonomy; limited real-time feedback; brittle under schema evolution
Lakehouse without feature store	ACID tables over semi-structured data, scalable storage, partial streaming support	Supports foundational observability, but ML features are duplicated across teams; higher risk of offline/online skew
Lakehouse + enterprise feature store + orchestration	Unified tables with SCD-2 dimensions, governed feature entities, deterministic orchestration graphs	Aligns with agentic operation: reproducible decisions, idempotent pipelines, low-latency feature access, strong lineage

The comparison reflects arguments from Allam, Harby, Ionescu, de la Rúa Martínez, and Liang regarding the shift from rigid ETL pipelines toward flexible, metadata-centric architectures in finance [2,4–5,7–8]. In the first pattern, agents are constrained to operate on static snapshots; recomputation requires full batch reruns, and partial failures often force manual intervention. In the second pattern, agents gain real-time data access, but absence of a feature store fragments feature logic across codebases, undermining determinism. The third pattern brings together lakehouse guarantees, feature governance, and metadata-driven orchestration, creating an environment where agents can recompute fallible steps, replay transactions, and rely on a stable semantic layer.

Despite these advantages, the surveyed works on AI agents and financial architectures repeatedly caution about systemic risks and governance gaps. Aldasoro et al. draw attention to feedback loops between AI-driven decision systems and financial stability [1]. Hettiarachchi underscores that generative AI agents introduce new operational-risk dimensions around explainability, drift, and adversarial behavior [6]. Liang and colleagues stress attack surfaces around feature-store access, data poisoning, and leakage of sensitive features [8]. Translated to payroll, these concerns manifest as mis-classified workers, systematically biased rule encoding, or silent failures in anomaly detection during periods of regulatory change. Mitigation requires defense-in-depth: layered validation agents, independent monitoring services, and human approval gates for high-impact operations such as payment release and retroactive corrections.

A further dimension emerging from Ionescu's work concerns sustainability and resource efficiency [7]. Agentic systems in finance can demand intensive compute resources, particularly when generative models or simulation-heavy decision policies are involved. For a global payroll platform running on a lakehouse, design choices around feature computation frequency, model refresh cadence, and aggregation granularity influence both cost and energy consumption. Combining sustainable-architecture principles with feature-store governance encourages approaches such as incremental feature updates, adaptive sampling for anomaly detection, and load-aware orchestration—choices that preserve responsiveness for critical controls while limiting unnecessary resource use.

Finally, literature on agent-based modeling in HR suggests that socio-technical alignment matters as much as technical soundness [9,10]. Payroll processes involve HR, finance, compliance, and workforce representatives whose trust in autonomous agents depends on transparency, traceability, and controllable failure modes. The conceptual model derived in the results section acknowledges this by assigning explanation-oriented responsibilities to certain agents and by grounding all decisions in auditable data and rule artifacts. The surveyed materials thus jointly support a view of autonomous payroll not as a fully self-driving system, but as a layered, agent-enabled platform where automation and governance are deliberately co-designed.

V. CONCLUSION

The synthesis of recent research on AI agents in finance, decentralized markets, data lakehouse architectures, sustainable financial data platforms, and feature-store-centric ML pipelines suggests that autonomous payroll processing becomes practicable when three conditions are met. First, functional responsibilities across the payroll lifecycle are assigned to specialized agents whose behaviors are bounded by explicit rules, structured features, and deterministic computation graphs rather than opaque heuristics. Second, these agents operate over an AI-ready financial data platform that combines an ACID-compliant lakehouse, an enterprise feature store, and metadata-driven orchestration, enabling idempotent pipelines, concurrency-safe merge patterns, and comprehensive lineage. Third, governance, security, and sustainability measures are integrated into the design from the outset, addressing concerns highlighted in the literature regarding explainability, systemic risk, and resource consumption.

The tasks set in the introduction are thereby resolved: a taxonomy of agent types tailored to payroll operations is formulated, their interaction with modern financial data architectures is described, and their limitations and governance implications are analyzed using evidence from the surveyed works. For practitioners in data engineering and AI who design payroll and financial platforms, the article provides a conceptual blueprint that aligns agentic automation with regulatory requirements and enterprise-grade reliability, while leaving room for future research on empirical evaluation, formal verification, and large-scale deployment outcomes.

REFERENCES

[1] Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., & Stein, M. (2024). Intelligent financial system: How AI is transforming finance (BIS Working Paper No. 1194). Bank for International Settlements.

[2] Allam, K. (2024). Cloud-based data hubs and SQL pipelines for real-time financial analytics. International Journal of Emerging Trends in Computer Science and Information Technology.

[3] Ante, L. (2024). Autonomous AI agents in decentralized finance: Market dynamics, application areas, and theoretical implications. SSRN preprint. https://doi.org/10.2139/ssrn.5055677

[4] de la Rúa Martínez, J., et al. (2024, June). The Hopsworks feature store for machine learning. In Proceedings of the 2024 ACM SIGMOD International Conference on Management of Data (SIGMOD '24).

[5] Harby, A. A., & Zulkernine, F. H. (2025). Data lakehouse: A survey and experimental study. Information Systems, 127, 102460.

[6] Hettiarachchi, I. (2025). The rise of generative AI agents in finance: Operational disruption and strategic evolution. International Journal of Engineering Technology Research & Management, 9(4).

[7] Ionescu, S. A. (2025). Engineering sustainable data architectures for modern financial institutions. Electronics, 14(8), 1650.

[8] Liang, W., Taofeek, A., Rajuroy, A., Blessing, M., & Hamzah, F. (2025). Developing scalable and secure feature stores in real-time machine learning pipelines. Manuscript in preparation / preprint.

[9] Somarathna, K. U. S. (2020). An agent-based approach for modeling and simulation of human resource management as a complex system: Management strategy evaluation. Simulation Modelling Practice and Theory, 104, 102118.

[10] Yordanova, Z., & Hristizov, Y. (2025). The evolution of financial analysis: From manual methods to AI and AI agents. Economics – Innovative and Economics Research Journal, 13(3), 219–239.

