

Governed AI Agents in Consumer Banking: Data Stewardship and Data Quality Framework for Credit Card and Deposit Analytics

LIBERTY MUDZINGWA¹, TSUNGAI E TSAMBATARE², MELODY MASUNDA³, MUNASHE NAPHTALI MUPA⁴

^{1,2}American University

³Suffolk University

⁴Hult International Business School

Abstract- Consumer banks and credit unions are moving from conventional dashboards and rules-based automation toward AI agents that can query data warehouses, summarize portfolio signals, monitor exceptions, draft analytical narratives and support frontline and risk teams. The value of these systems depends less on the novelty of the model than on whether the underlying data assets are governed, classified, traceable, timely and fit for regulated decision support. This article develops a practical data-stewardship and data-quality framework for governed AI agents in consumer banking, with emphasis on credit-card and deposit analytics. The framework integrates three bodies of practice: data governance and data-quality scholarship, model-risk and AI-risk governance, and banking-specific expectations around adverse-action explainability, risk-data aggregation and operational resilience. A synthetic consumer-banking test bed is used to demonstrate a risk-weighted data-quality score, business-glossary coverage, data-quality rule monitoring, lineage completeness, access-classification discipline, and AI-agent output controls. Heat-map analysis shows that the highest residual-risk areas are lineage gaps, weak glossary coverage, disputes/chargebacks, credit-bureau attributes and unexplained agent responses. The simulated before-and-after analysis indicates that a governed control layer can raise risk-weighted data-quality scores, improve agent acceptance, reduce escalation and shorten exception aging when controls are implemented as operating mechanisms rather than policy documents. The article contributes a reusable operating model for credit unions, community banks and consumer-banking teams seeking to operationalize AI agents without weakening data accountability, consumer-protection obligations or executive decision reliability.

Keywords: AI Agents, Consumer Banking, Data Stewardship, Data Quality, Credit-Card Analytics,

Deposit Analytics, Model Risk Management, Business Glossary, Credit Unions, Governed Analytics.

I. INTRODUCTION

Consumer banking has become a data-intensive operating environment in which credit-card portfolios, deposit accounts, transaction streams, digital-channel interactions, disputes, complaints, collections, fraud signals and customer master data must be assembled into reliable analytical products. The technical promise of AI agents is that they can compress analytical work: they can answer portfolio questions, retrieve policy definitions, identify data anomalies, summarize credit-card and deposit trends, prepare exception narratives and assist with monitoring workflows. Yet the same capability creates a governance problem. An agent that can access multiple data sources and generate decisions or recommendations also inherits the defects, blind spots, lineage gaps and classification weaknesses of those sources.

The governance issue is especially acute in credit-card and deposit analytics. Credit-card analytics may support portfolio risk monitoring, line-management strategy, payment behavior analysis, delinquency management, dispute classification, interchange and fee analysis, complaints, fraud review, marketing eligibility and adverse-action narratives. Deposit analytics may support balance trends, account activity, cash-flow behavior, overdraft monitoring, segmentation, liquidity signals, product profitability, digital adoption and relationship management. In both product groups, the risk is not only whether an

AI agent returns a plausible answer. The deeper question is whether the answer is grounded in approved data definitions, current data, complete lineage, appropriate access entitlements, monitored data-quality rules and explainable analytical logic.

Data quality is not a narrow technology defect. It is an accountability mechanism. Wang and Strong (1996) established that data quality must be understood from the perspective of data consumers, while Wand and Wang (1996) framed data quality as a multidimensional problem involving accuracy, completeness, consistency and timeliness. Khatri and Brown (2010) further positioned data governance as a question of decision rights and accountabilities across domains such as data principles, data quality, metadata, access and compliance. These ideas become more important when a bank deploys AI agents because the agent becomes an accelerated consumer of enterprise data. If decision rights, data

definitions and quality rules are ambiguous, the agent scales ambiguity.

This article argues that governed AI agents in consumer banking require an operating model that joins data stewardship, quality monitoring, access classification, model-risk discipline and human review. The proposed framework is not designed for fully autonomous credit decisions. Rather, it is designed for credit unions, community banks and consumer-banking divisions that wish to use AI agents responsibly for analytics, portfolio monitoring, data stewardship, exception detection and executive decision support. The central claim is that agentic analytics can be made more reliable when the data layer is governed as a controlled production asset and when every agent output is linked to traceable data, measurable rule performance, business glossary definitions and risk-tiered human oversight.

Table 1. AI-agent use cases and governed data assets in consumer banking

AI-agent use case	Main data assets	Suggested risk tier	Analytical value	Primary governance risk
Credit-card portfolio monitoring	Credit-card accounts, transactions, payments, disputes, bureau attributes	Tier 2-3	Delinquency trend, payment behavior, exposure, dispute and chargeback analytics	Lineage, bureau data quality, adverse-action explainability, privacy
Deposit balance and cash-flow analytics	Deposit accounts, transactions, balances, digital interactions	Tier 2	Balance behavior, cash-flow volatility, overdraft and relationship analytics	Refresh timing, transaction classification, account status consistency
Business glossary and policy Q&A	Glossary, policy documents, data dictionary, approved metrics	Tier 1	Definition lookup, report interpretation, policy navigation	Version control, source authority, restricted policy material
Data stewardship exception triage	DQ rule results, ETL logs, lineage, issue-management records	Tier 2	Exception routing, root-cause summaries, data owner notifications	False alerts, stale metadata, unclear accountability
Executive dashboard narration	Certified dashboard outputs, semantic layer, KPI library	Tier 2	Automated narrative of portfolio and deposit trends	Unapproved metric definitions, omission of confidence limitations
Decision-adjacent credit support	Application data, bureau attributes, policy rules, explainability artefacts	Tier 3	Human-reviewed adverse-action support and model explanation workflow	Fair lending, ECOA/Regulation B, model validation, auditability

II. RESEARCH PROBLEM AND BANKING CONTEXT

The market context justifies this emphasis. The Consumer Financial Protection Bureau's 2025 consumer credit-card market review describes credit

cards as central to the financial lives of millions of consumers and examines market conditions using multiple data sources, while Federal Reserve G.19 consumer credit statistics continue to track revolving and non-revolving consumer credit as major categories of household borrowing. These public signals show why credit-card analytics are not peripheral operational reports; they inform consumer access, risk appetite, customer treatment, compliance posture and portfolio economics. Similarly, FDIC Quarterly Banking Profile materials show that bank performance, asset quality, deposits and earnings are monitored at industry level, confirming that deposit and credit information remain supervisory and managerial priorities.

Regulatory expectations also reinforce the need for governance. The CFPB has emphasized that creditors using complex algorithms in credit decisions must still provide specific reasons for adverse action.

Model-risk management guidance from U.S. banking regulators requires attention to conceptual soundness, implementation, validation, monitoring, governance and documentation. The Basel Committee's principles for risk-data aggregation and risk reporting similarly emphasize governance, accuracy, integrity, completeness, timeliness and adaptability in risk data. Although these sources were not written solely for AI agents, they provide a relevant governance baseline: regulated analytics must be explainable, reliable, monitored and appropriately controlled.

A further challenge is that AI agents shift analytics from static reporting to interactive, request-driven analysis. A dashboard exposes metrics that have usually passed through a release process. An AI agent, by contrast, can answer novel questions, combine data fields in unexpected ways, summarize unstructured policy text, route queries to tools, and generate explanations on demand. This raises three governance questions. First, what data assets is the agent allowed to use? Second, what quality evidence must exist before an answer is trusted? Third, how should the organization monitor the agent's outputs, exceptions and residual risk over time?

Table 2. Public and supervisory signals supporting the need for governed consumer-banking analytics

Context signal	Relevant evidence	Implication for governed AI agents
Consumer credit growth and credit-card relevance	Federal Reserve G.19 tracks household consumer credit and distinguishes revolving credit, which includes credit-card plans.	AI agents supporting credit-card analytics need high confidence in revolving credit, payment and account status definitions.
Consumer credit-card market scale	CFPB credit-card market reporting examines consumer credit-card conditions using multiple data sources and market segments.	Credit-card analytics affect consumer treatment, portfolio risk and compliance reporting.
Banking performance and asset quality	FDIC Quarterly Banking Profile summarizes earnings, loan and deposit activity, and asset quality for insured institutions.	Deposit and loan analytics are institution-critical and supervisory-relevant.
Adverse-action explainability	CFPB guidance states that creditors using complex algorithms must still provide specific reasons for adverse action.	Credit-card decision-adjacent agents require traceable features and reasons.
Model-risk governance	Banking model-risk guidance emphasizes conceptual soundness, validation, implementation, monitoring, governance and documentation.	AI agents that route analytical logic or support models should be inventoried, monitored and documented.

Risk-data aggregation	BCBS 239 principles emphasize governance, accuracy, completeness, timeliness and adaptability in risk data.	Agent outputs used in risk reporting must be grounded in reliable data aggregation practices.
-----------------------	---	---

III. LITERATURE REVIEW

The literature on data quality provides the foundation for the framework. Wang and Strong (1996) identified data quality as a consumer-centered concept, not merely a database attribute. Wand and Wang (1996) located quality problems in the representation of real-world states and highlighted dimensions such as completeness, accuracy and consistency. More recent surveys have shown that data-quality dimensions remain fragmented in practice and must be translated into actionable controls, checks and monitoring routines (Papastergios, Ehrlinger and Gounaris, 2025). For banking AI agents, the implication is direct: quality must be defined as fitness for a particular analytic use case, not as a generic certificate of database cleanliness.

Data governance scholarship provides the organizational layer. Khatri and Brown (2010) conceptualize data governance through decision domains and accountability structures. Weber, Otto and Österle (2009) show that one governance model does not fit all organizations and that governance design must reflect contingency factors. Abraham, Schneider and vom Brocke (2019) synthesize the data-governance literature and define governance as the exercise of authority and control over data management to increase value while minimizing cost and risk. In consumer banking, this means that credit-card and deposit data cannot be left solely to technology owners. Product, risk, compliance, operations, data engineering, analytics and model-risk teams must share a clear decision-rights model.

The AI governance literature adds a system-level risk perspective. Schneider et al. (2020) argue that AI governance concerns authority and control over AI systems, including data, models and AI-enabled products. NIST's AI Risk Management Framework structures AI risk management around Govern, Map, Measure and Manage functions. Kurshan, Shen and

Chen (2020) note that financial-services AI model governance faces challenges of effectiveness, cost, complexity and speed, and propose more automated monitoring and mitigation capabilities. Agentic AI extends these concerns because agents combine model behavior, tool access, data retrieval, task planning and generated responses.

Research on explainability and credit scoring is also relevant. Ribeiro, Singh and Guestrin (2016) proposed local explanation methods for understanding model predictions, while Rudin (2019) warned against relying on post hoc explanations for high-stakes decisions when interpretable models can be used. Bücken et al. (2022) developed dimensions for making machine-learning credit-scoring models transparent, auditable and explainable. These studies caution banks against treating explainability as a cosmetic reporting layer. For credit-card analytics, explainability must be designed into data definitions, features, model outputs, adverse-action support and agent narratives from the start.

Mupa et al. (2025) discuss the use of machine learning in actuarial risk management and identify data quality, interpretability and regulatory compliance as constraints on predictive modelling. Homwe et al. (2025) examine interpretable machine learning for audit planning in financial services and emphasize risk scoring, SHAP-type explanations, dashboarding and visual decision support. Nguruve et al. (2026) discuss operationalizing AI at scale through repeatable frameworks, adoption discipline and performance measurement. Nhova et al. (2026) develop a controls-first analytics framework in a regulated finance-adjacent environment, showing how continuous monitoring, KPI libraries and audit-ready evidence packs can connect governance to operational execution. These sources support the article's focus on practical, controlled, evidence-producing analytics rather than abstract AI enthusiasm.

Table 3. Literature streams informing the CB-GADS framework

Literature stream	Illustrative sources	Core contribution	Relevance to this article
Data quality	Wang and Strong (1996); Wand and Wang (1996); Papastergios et al. (2025)	Defines quality as multidimensional and consumer/use-case dependent.	Metrics must be tied to AI-agent use cases, not generic database hygiene.
Data governance	Weber et al. (2009); Khatri and Brown (2010); Abraham et al. (2019)	Defines decision rights, accountability and authority over data assets.	Every critical data element should have owner, steward, definition and rule evidence.
AI governance	Schneider et al. (2020); NIST (2023); Kurshan et al. (2020)	Connects data, models, systems, risk assessment and lifecycle monitoring.	Agent governance must integrate model, data, prompt/tool and output controls.
Explainability and credit scoring	Ribeiro et al. (2016); Rudin (2019); Bucker et al. (2022)	Shows the importance and limits of explanations in high-stakes models.	Credit-card agents require traceability, auditable features and human review.
Controls-first analytics	Homwe et al. (2025); Nhova et al. (2026); Nguruve et al. (2026)	Shows how analytics needs visual evidence, adoption discipline, KPI libraries and control evidence.	Consumer-banking agent adoption should be measured as operational control maturity.
Agentic AI in finance	Okpala et al. (2025); Axelsen et al. (2025); Kurshan et al. (2025); Aldridge et al. (2026)	Identifies agent workflows, financial-service use cases and new control gaps.	Autonomous workflows need scoped roles, audit logs, policy enforcement and monitoring.

IV. CONCEPTUAL FRAMEWORK:
 CONSUMER BANKING GOVERNED
 AGENT DATA STEWARDSHIP
 FRAMEWORK

The proposed framework is named the Consumer Banking Governed Agent Data Stewardship Framework (CB-GADS). It is built on eight components: data-asset ownership, business glossary discipline, data-quality rule design, lineage capture, access classification, model and agent monitoring, human review, and audit evidence. Each component is necessary because an AI agent can fail through multiple pathways. It can answer from stale data; combine fields with inconsistent definitions; expose restricted attributes; summarize incomplete transactions; generate an answer that is not reproducible; or produce a plausible narrative that cannot be traced to authoritative sources.

CB-GADS treats a consumer-banking data asset as production-grade only when it has an owner, steward, glossary definition, critical-data-element classification, lineage map, quality rules, access

class, refresh standard, exception route and evidence retention policy. This moves governance from a policy document to an operating mechanism. The framework also distinguishes between three AI-agent tiers. Tier 1 agents support low-risk information retrieval and glossary search. Tier 2 agents support portfolio analytics and exception summarization under human review. Tier 3 agents support regulated decision-adjacent workflows, such as credit-card adverse-action support, line-management analytics, fraud risk triage or complaint trend analysis, and require stricter documentation, monitoring, validation and approval.

The framework is deliberately product-aware. Credit-card analytics and deposit analytics share several data-governance needs, but their risk profiles differ. Credit-card analytics raise stronger concerns around credit policy, consumer protection, fair-lending proxies, adverse-action explainability, delinquency status, disputes, charge-offs and bureau attributes. Deposit analytics raise concerns around balances, transaction timing, overdraft behavior, account status, suspicious activity routing, product eligibility and

cash-flow interpretation. Therefore, a shared governance operating model must allow product-specific rule thresholds and evidence requirements.

Table 4. Components of the Consumer Banking Governed Agent Data Stewardship Framework

CB-GADS component	Control objective	Evidence artefacts	Why it matters for AI agents
1. Data-asset ownership	Named owner and steward for each critical data asset	Owner register; stewardship RACI; approval workflow	Prevents orphaned data assets used by agents without accountable oversight
2. Business glossary discipline	Approved definitions for product, risk, customer and operational terms	Glossary entries; version history; definition owner	Reduces semantic drift in agent-generated narratives
3. Data-quality rules	Automated checks for completeness, validity, consistency, timeliness and uniqueness	Rule library; exception logs; pass/fail thresholds	Creates measurable evidence that source data is fit for use
4. Lineage capture	Traceability from source system to table, semantic layer, dashboard and agent answer	Lineage map; source-to-target documentation	Supports reproducibility and audit review
5. Access classification	Data classification and entitlement matching for agent roles	Access matrix; masking rules; approval tickets	Prevents overexposure of sensitive customer or bureau attributes
6. Model and agent monitoring	Ongoing monitoring of answer quality, escalation, drift and incidents	Agent telemetry; model-monitoring dashboard; validation pack	Connects data quality to model/agent risk management
7. Human review	Risk-tiered human approval for high-impact outputs	Review logs; override reasons; quality sampling	Keeps regulated or decision-adjacent workflows under accountable judgment
8. Audit evidence	Retained evidence showing control design and operating effectiveness	Evidence pack; issue remediation record; audit trail	Makes governance defensible to compliance, internal audit and regulators

V. METHODOLOGY AND DATA ANALYTIC DESIGN

The empirical component of this article uses a synthetic consumer-banking data-quality test bed. No confidential employer, customer, credit-union or bank records are used. The test bed is designed to resemble the structure of credit-card and deposit analytics environments by modelling ten data domains across eight data-quality dimensions over a 12-month period. Six months represent a pre-control state and six months represent a post-control state after implementation of stewardship controls, glossary rules, access classification and monitoring routines. The analysis is illustrative but operationally grounded: the point is to demonstrate how a financial

institution can measure and govern data readiness for AI-agent analytics without disclosing proprietary information.

The synthetic test bed contains domain-month-dimension observations for customer master data, credit-card accounts, card transactions, card payments, deposit accounts, deposit transactions, digital-channel interactions, disputes and chargebacks, credit-bureau attributes, and product/pricing reference data. The dimensions are completeness, validity, consistency, timeliness, uniqueness, lineage, access classification and glossary coverage. These dimensions were selected because they connect classical data-quality

scholarship with the operational requirements of AI-agent analytics.

A risk-weighted data-quality score is calculated for each domain and month. Completeness and validity each receive 18 percent weight because missing and invalid data can directly corrupt agent responses. Consistency and timeliness each receive 14 percent weight because cross-system mismatch and stale data undermine analytical reliability. Lineage receives 12 percent weight because an agent output must be reproducible and explainable. Uniqueness, access classification and glossary coverage receive the remaining weights because they govern duplicate records, entitlement discipline and semantic clarity. The score is not presented as a universal standard; it

is a practical scoring model that each institution can adapt to its risk appetite and product environment.

Three analyses are performed. First, a before-and-after comparison estimates domain-level improvement in risk-weighted data-quality scores after control implementation. Second, heat-map analysis identifies weak domains, dimensions and control families requiring priority remediation. Third, correlation analysis estimates how data-quality dimensions relate to AI-agent performance indicators such as answer acceptance, escalation rate and explainability pass rate. These analyses demonstrate how a data-stewardship program can create evidence for executives, auditors, model-risk reviewers and product owners.

Table 5. Synthetic consumer-banking data-quality test-bed design

Design element	Specification	Rationale
Observation unit	Domain-month-dimension record	Each row represents one data-quality dimension for one domain in one month.
Time horizon	12 months	Six months pre-control and six months post-control.
Domains	10	Customer, card, deposit, transaction, dispute, bureau, digital and product-reference domains.
Dimensions	8	Completeness, validity, consistency, timeliness, uniqueness, lineage, access classification, glossary coverage.
Scoring scale	0 to 100	Higher score means better control performance.
Risk-weighted score	Weighted sum of dimension scores	Weights reflect consumer-banking risk and agent-output sensitivity.
Agent metrics	Acceptance, escalation, explainability, override and refresh breaches	Used to connect data quality to agent performance and control outcomes.
Ethical constraint	Synthetic data only	No confidential customer or employer data is used.

Table 6. Risk-weighting model for consumer-banking data-quality score

Dimension	Weight	Rationale
Completeness	18%	Missing values can cause incorrect portfolio, customer and exception narratives.
Validity	18%	Invalid codes, dates and balances can break rules and distort agent outputs.
Consistency	14%	Cross-system mismatch weakens reconciliation and executive confidence.
Timeliness	14%	Stale data creates outdated risk and balance interpretations.
Uniqueness	8%	Duplicate records inflate counts, balances, disputes or customers.
Lineage	12%	Traceability is required for reproducibility, auditability and explainability.
Access classification	8%	Agents must not use or expose data outside approved entitlement scope.
Glossary coverage	8%	Approved definitions reduce semantic drift and unsupported narratives.

VI. RESULTS: DATA-QUALITY SCORING,
 HEAT MAPS AND AGENT-
 PERFORMANCE INDICATORS

The simulated analysis produced three main findings. First, improvements were largest where the pre-control state was weakest: disputes and chargebacks, credit-bureau attributes, digital-channel interactions and customer master data. These are also the domains in which AI agents are likely to be asked complex questions involving customer behavior, complaint context, bureau inputs, product eligibility and exception narratives. The finding supports a practical rule: data-quality remediation should not be spread evenly across all domains; it should be prioritized according to downstream decision risk and agent exposure.

Second, the heat maps show that lineage and glossary coverage are the most persistent gaps after basic completeness and validity rules are addressed. This is consistent with field experience in banking analytics. Many institutions can identify missing values and

invalid codes, but fewer can show end-to-end lineage from system of record to warehouse table, semantic layer, dashboard, feature store and AI-agent response. For governed agents, lineage is not optional. Without lineage, the institution cannot reproduce the answer, challenge the calculation, determine data freshness or explain why a narrative changed.

Third, the correlation analysis indicates that agent acceptance and explainability pass rates move strongly with validity, lineage and glossary coverage, while escalation rates move in the opposite direction. This result is intuitive and useful. Agents perform better when data is valid, traceable and semantically defined; they require more escalation when data lacks definitions, contains mismatches or cannot be mapped to approved sources. The analysis therefore supports a shift in performance management: agent quality should not be measured only by response accuracy or latency, but also by the health of the data assets supporting the response.

Table 7. Pre- and post-control risk-weighted data-quality score by data domain

Domain	Pre-control	Post-control	Delta	Relative improvement %
Card payments	73.7	83.6	9.9	13.4
Card transactions	79.2	89.5	10.3	13.0
Credit-bureau attributes	63.4	73.0	9.6	15.1
Credit-card accounts	75.9	86.8	10.9	14.4
Customer master	71.3	80.9	9.6	13.5
Deposit accounts	77.9	88.2	10.3	13.2
Deposit transactions	81.4	91.1	9.7	11.9
Digital-channel interactions	67.4	77.1	9.7	14.4
Disputes and chargebacks	64.7	75.1	10.4	16.1
Product and pricing reference	83.7	92.9	9.2	11.0

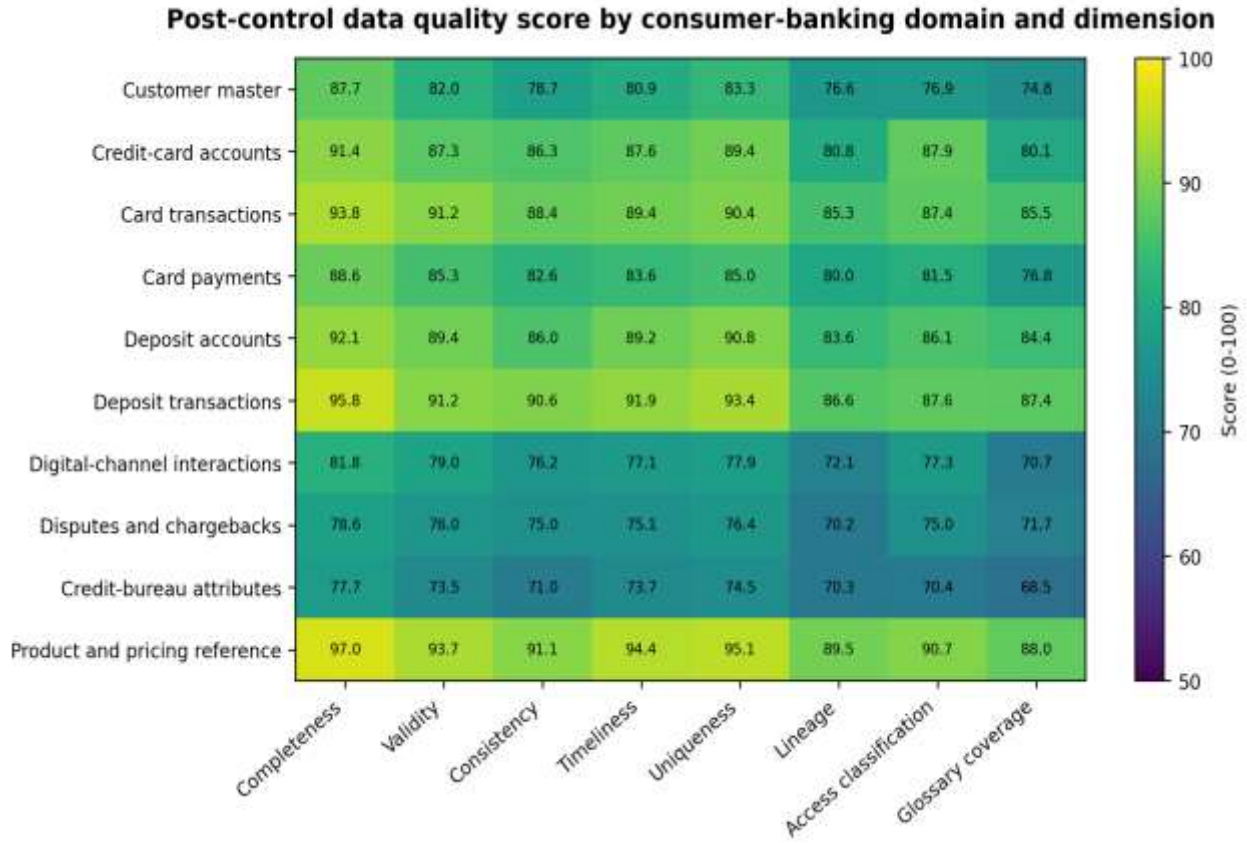


Figure 1. Heat map of post-control data-quality scores by domain and dimension. Higher values represent stronger control performance.

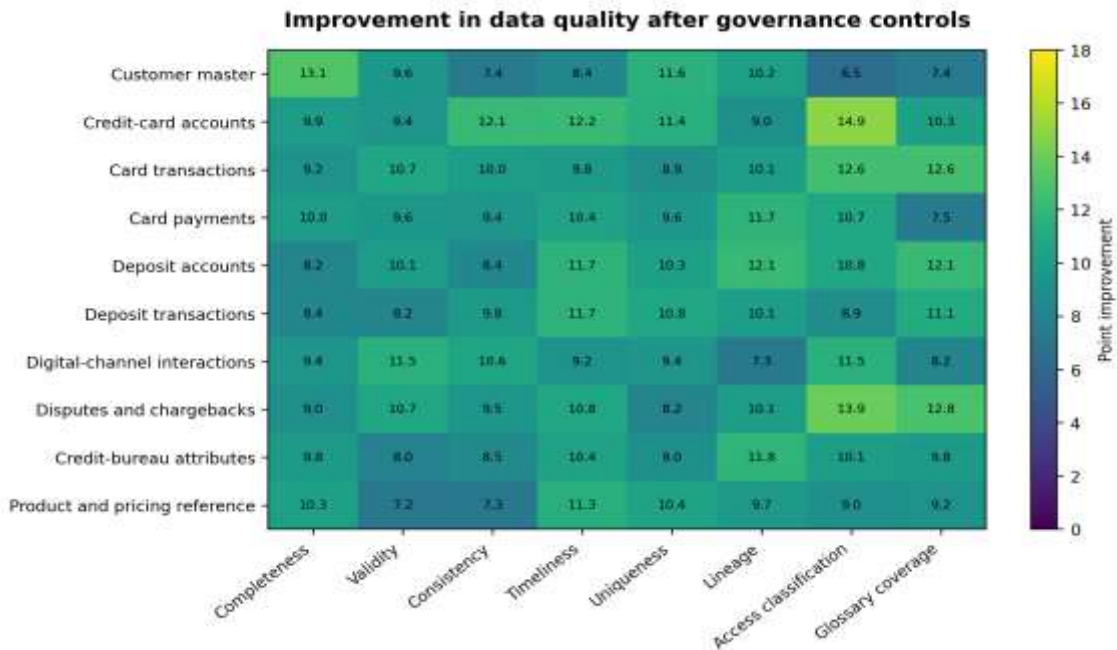


Figure 2. Heat map of point improvements after data-stewardship and data-quality controls. Larger values identify remediation impact.



Figure 3. Residual risk heat map by product and control family. Higher values indicate stronger priority for remediation and monitoring.



Figure 4. Correlation heat map linking data-quality dimensions to AI-agent performance indicators in the synthetic test bed.

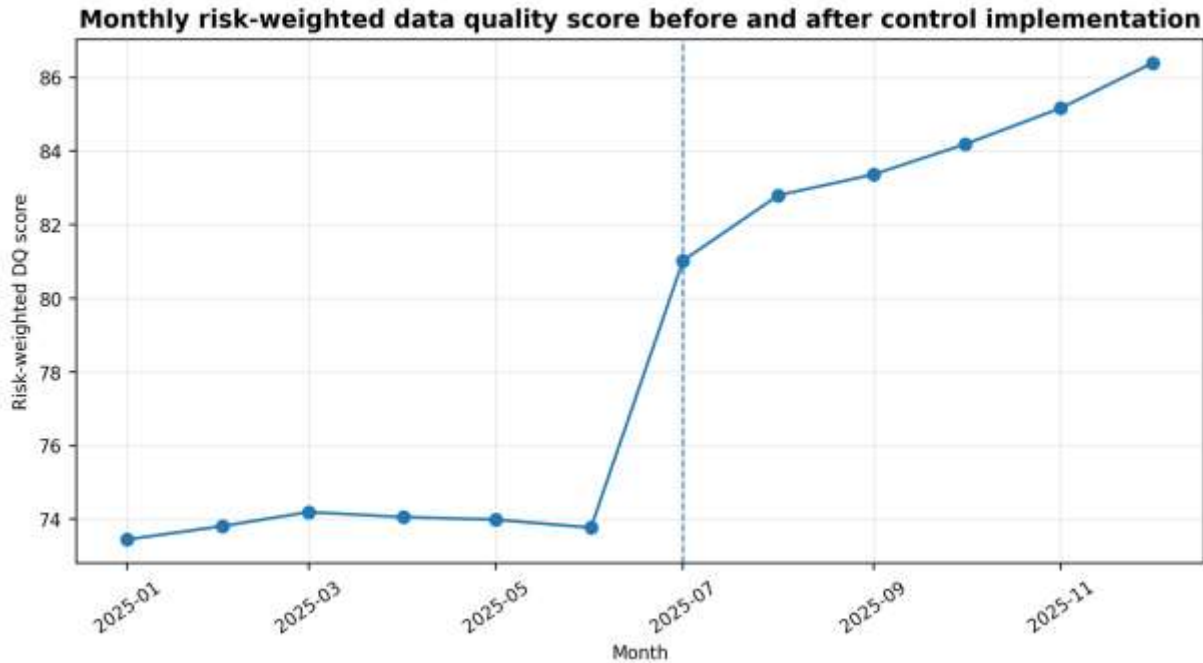


Figure 5. Monthly trend in average risk-weighted data-quality score before and after implementation of controls.

Table 8. Simulated AI-agent performance indicators before and after governance-control implementation

Metric	Pre-control	Post-control	Change
Agent answer acceptance (%)	70.4	84.7	14.3
Human override rate (%)	15.8	8.9	-6.9
Escalation rate (%)	21.6	11.3	-10.3
DQ exception aging (days)	6.8	2.7	-4.1
Explainability pass rate (%)	68.9	86.8	17.9
Dashboard refresh breaches/month	14.0	4.0	-10.0
Unauthorised data-access attempts blocked (%)	91.2	98.5	7.3

Table 9. Data-quality exception counts by domain and error type in the synthetic post-control monitoring sample

Domain	Missing mandatory field	Invalid code/value	Cross-system mismatch	Stale record	Duplicate key	Unclassified sensitive attribute	Unexplained agent response
Customer master	20.0	17.0	28.0	36.0	19.0	10.0	26.0
Credit-card accounts	19.0	24.0	29.0	15.0	12.0	14.0	23.0
Card transactions	10.0	22.0	25.0	20.0	10.0	14.0	23.0
Card payments	25.0	12.0	27.0	18.0	11.0	14.0	20.0
Deposit accounts	15.0	24.0	23.0	17.0	10.0	17.0	22.0

Deposit transactions	13.0	7.0	21.0	6.0	4.0	10.0	22.0
Digital-channel interactions	27.0	18.0	33.0	21.0	16.0	16.0	25.0
Disputes and chargebacks	33.0	19.0	35.0	22.0	17.0	24.0	29.0
Credit-bureau attributes	24.0	30.0	50.0	31.0	20.0	22.0	35.0
Product and pricing reference	16.0	12.0	10.0	17.0	6.0	7.0	19.0

VII. DISCUSSION

The results have several implications for consumer-banking governance. The first is that data stewardship must be attached to product outcomes. A business glossary is not valuable because it exists; it is valuable when a product manager, risk analyst, data engineer, compliance reviewer and AI agent use the same approved definition of delinquency bucket, available credit, posted balance, current balance, charge-off status, dispute reason, interest rate type, deposit hold, overdraft event or active account. The glossary should therefore be governed as a control instrument, not as documentation for documentation's sake.

The second implication is that data-quality monitoring must be continuous. Traditional periodic reconciliation is too slow for AI-agent workflows that may generate analysis on demand. Data-quality rules should be embedded into the ingestion, transformation, semantic and agent-answering layers. The evidence should be retained automatically: rule results, exception logs, refresh timestamps, lineage maps, access checks, prompt/tool records, human review notes and override decisions. This evidence pack becomes the practical bridge between innovation and auditability.

The third implication is that AI-agent access should be risk-tiered. Not every agent needs access to every data asset. A glossary-search agent can operate with low-risk metadata. A portfolio-insight agent may need aggregated credit-card and deposit data with masking. A decision-adjacent agent may require strict

access approvals, human-in-the-loop review, fairness testing, adverse-action traceability and model-risk oversight. Treating all agents as equivalent either overburdens low-risk use cases or under-controls high-risk ones.

The fourth implication is that model-risk management and data governance should converge. Many model-risk programs evaluate models after development, while many data-governance programs monitor tables and definitions separately. AI agents do not respect those boundaries. They combine data retrieval, model inference, natural-language generation, tool calling and user interaction. A governed-agent program should therefore maintain an integrated inventory linking use case, agent owner, data assets, critical data elements, model components, prompts/tools, quality rules, validation evidence, access controls and monitoring dashboards.

The fifth implication concerns credit unions and community banks. Smaller institutions may not have the resources of large banks, but they can still implement a disciplined version of CB-GADS. In fact, a lighter but rigorous framework may be more valuable than a complex enterprise program that never reaches business users. The minimum viable control set should include a named data owner and steward for each critical asset, a glossary for high-risk metrics, basic completeness and validity rules, refresh monitoring, access classification, lineage for priority reports and a human-review workflow for AI-agent outputs used in executive or regulated contexts.

Table 10. Implementation roadmap for governed AI agents in consumer banking

Implementation phase	Core activities	Minimum evidence	Primary accountable functions
Phase 1: Inventory and risk tiering	Create agent inventory; classify use cases; map data assets; assign owners	Approved inventory; risk-tier register; data owner list	Data governance lead with product, risk and IT
Phase 2: Glossary and CDE definition	Define critical data elements for card and deposit analytics; approve metric definitions	Glossary; CDE catalogue; metric approval minutes	Product owner and data steward
Phase 3: Rule library and monitoring	Implement completeness, validity, consistency, timeliness, uniqueness and access checks	Rule library; dashboards; exception queue	Data engineering and stewardship
Phase 4: Lineage and access controls	Map source-to-agent lineage; restrict agent access to approved assets; implement masking	Lineage diagrams; access matrix; approval tickets	Data architecture, IAM and security
Phase 5: Agent validation and human review	Define answer-quality tests, escalation rules and human approval thresholds	Validation test pack; review workflow; override log	Model risk, compliance and business owner
Phase 6: Evidence pack and continuous improvement	Retain telemetry, rule results, issue remediation and audit evidence	Audit-ready evidence pack; KPI dashboard; remediation backlog	Internal audit, compliance, data governance

Table 11. Illustrative RACI matrix for governed credit-card and deposit AI agents (R=Responsible, A=Accountable, C=Consulted, I=Informed)

Function	Asset ownership	Glossary/CDEs	DQ rules	Access/lineage	Agent monitoring	Human review	Evidence pack
Product owner	A	C	C	C	A	C	C
Data owner	A	A	A	A	C	C	C
Data steward	R	R	R	R	C	C	R
Data engineer	C	C	R	R	C	C	R
Model risk management	C	C	C	C	R	R	C
Compliance/fair lending	C	C	C	C	C	A	C
Information security/IAM	C	C	C	A	C	C	C
Internal audit	I	I	I	I	I	I	A

VIII. PRACTICAL POLICY RECOMMENDATIONS

- Do not deploy consumer-banking AI agents against uncertified data assets. Certification should require ownership, glossary coverage, access classification, refresh standards, lineage and rule monitoring.
- Separate low-risk glossary agents from higher-risk portfolio, exception and decision-adjacent

agents. Each tier should have different approval, logging and review requirements.

- Treat business glossary coverage as a control, not a documentation exercise. Terms used by an agent should map to approved definitions and versioned sources.
- Connect model-risk management and data governance inventories. Agent use cases should show the model components, data assets,

prompts/tools, critical data elements, validation results and monitoring evidence.

- Measure agent performance with data-health indicators. Acceptance rate, escalation rate and response latency are insufficient without data-quality scores, lineage coverage and explainability pass rates.
- Retain an audit-ready evidence pack for high-risk agents. The pack should include rule results, exception logs, prompt/tool audit logs, human reviews, override reasons, validation results and remediation tickets.
- For credit-card use cases, maintain special controls over bureau variables, adverse-action reason codes, credit-line management variables, delinquency status and dispute/chargeback attributes.
- For deposit use cases, maintain special controls over balance timing, account status, transaction classification, overdraft events, holds, product pricing and digital-channel metadata.

IX. LIMITATIONS AND FUTURE RESEARCH

This article is limited by its use of a synthetic data-quality test bed. The simulation is useful for demonstrating measurement logic, heat-map design and governance priorities, but it is not a substitute for empirical validation using de-identified institutional data. The scoring weights are judgment-based and should be recalibrated by each institution according to product risk, regulatory obligations, data architecture and AI-agent use cases. The article also focuses primarily on credit-card and deposit analytics rather than mortgage, auto, small-business, wealth, fraud or BSA/AML workflows, although the framework can be adapted to those environments.

Future studies should test CB-GADS in live institutional settings, compare human-only analytics with agent-assisted analytics, evaluate fairness and adverse-action outcomes in credit-card decision-adjacent workflows, quantify the cost of remediation versus reduction in exception burden, and examine how agent telemetry can be incorporated into enterprise model-risk dashboards. A further research opportunity is to compare centralized, federated and

data-mesh-style stewardship models for community banks and credit unions with limited data-governance staffing.

X. CONCLUSION

AI agents can improve the speed, coverage and accessibility of consumer-banking analytics, but they cannot compensate for weak data governance. In credit-card and deposit analytics, the risk of poor data quality is not merely a technical inconvenience. It can affect portfolio interpretation, customer treatment, consumer-protection compliance, executive decisions, operational resilience and institutional trust. The appropriate response is not to avoid AI agents altogether, but to govern them through a data-stewardship and data-quality operating model that makes each output traceable, controlled and reviewable.

This article developed the CB-GADS framework for governed AI agents in consumer banking. The framework integrates data ownership, glossary discipline, data-quality rules, lineage, access classification, model and agent monitoring, human review and audit evidence. The synthetic data analysis demonstrated how risk-weighted data-quality scoring, heat maps and correlation analysis can identify weak domains and quantify improvements after controls are implemented. The results support a practical conclusion: AI-agent performance should be measured together with the health of the data assets, definitions and controls that support it.

Future research should test the framework using de-identified institutional data, compare alternative weighting schemes, evaluate fairness and adverse-action explainability in credit-card use cases, and study how agent telemetry can be incorporated into model-risk dashboards. For now, the central recommendation is clear. Consumer-banking AI agents should be deployed only where data stewardship, quality evidence, access control, lineage and human accountability are sufficiently mature to support reliable analytics.

REFERENCES

- [1] Abraham, R., Schneider, J. and vom Brocke, J. (2019) 'Data governance: A conceptual framework, structured review, and research agenda', *International Journal of Information Management*, 49, pp. 424-438. doi: 10.1016/j.ijinfomgt.2019.07.008.
- [2] Aldridge, I. et al. (2026) 'Agentic artificial intelligence in finance: A comprehensive survey'. Available at: <https://arxiv.org/abs/2604.21672>.
- [3] Axelsen, H., Licht, V. and Damsgaard, J. (2025) 'Agentic AI for financial crime compliance'. Available at: <https://arxiv.org/abs/2509.13137>.
- [4] Basel Committee on Banking Supervision (2013) Principles for effective risk data aggregation and risk reporting. Basel: Bank for International Settlements. Available at: <https://www.bis.org/publ/bcbs239.pdf>.
- [5] Biecek, P. et al. (2021) 'Enabling machine learning algorithms for credit scoring: Explainable artificial intelligence methods for understanding complex predictive models'. Available at: <https://arxiv.org/abs/2104.06735>.
- [6] Bücke, M., Szepannek, G., Gosiewska, A. and Biecek, P. (2022) 'Transparency, auditability and explainability of machine learning models in credit scoring', *Journal of the Operational Research Society*, 73(1), pp. 70-90.
- [7] Consumer Financial Protection Bureau (2022) Consumer Financial Protection Circular 2022-03: Adverse action notification requirements in connection with credit decisions based on complex algorithms. Available at: <https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/>.
- [8] Consumer Financial Protection Bureau (2025) The Consumer Credit Card Market: Report to Congress. Washington, DC: CFPB. Available at: https://files.consumerfinance.gov/f/documents/cfpb_consumer-credit-card-market-report_2025.pdf.
- [9] Federal Deposit Insurance Corporation (2025) Quarterly Banking Profile. Washington, DC: FDIC. Available at: <https://www.fdic.gov/quarterly-banking-profile>.
- [10] Federal Reserve Board (2026) Consumer Credit - G.19. Washington, DC: Board of Governors of the Federal Reserve System. Available at: <https://www.federalreserve.gov/releases/g19/>.
- [11] Federal Reserve Board, Office of the Comptroller of the Currency and Federal Deposit Insurance Corporation (2026) Revised guidance on model risk management. Washington, DC: Federal Reserve/OCC/FDIC.
- [12] FinRegLab (2025) The next wave arrives: Agentic AI in financial services. Washington, DC: FinRegLab. Available at: <https://finreglab.org/research/agentic-ai-in-financial-services/>.
- [13] Homwe, T., Mupa, M.N., Matope, A., Mlambo, N., Chingezi, L. and Chihota, T.A. (2025) 'Interpretable machine learning for audit planning: Improving misstatement and compliance risk detection in financial services', *World Journal of Advanced Research and Reviews*, 28(2), pp. 925-933. doi: 10.30574/wjarr.2025.28.2.3779.
- [14] Khatri, V. and Brown, C.V. (2010) 'Designing data governance', *Communications of the ACM*, 53(1), pp. 148-152. doi: 10.1145/1629175.1629210.
- [15] Kurshan, E., Shen, H. and Chen, J. (2020) 'Towards self-regulating AI: Challenges and opportunities of AI model governance in financial services'. Available at: <https://arxiv.org/abs/2010.04827>.
- [16] Kurshan, E., Balch, T. and Byrd, D. (2025) 'The agentic regulator: Risks for AI in finance and a proposed agent-based framework for governance'. Available at: <https://arxiv.org/abs/2512.11933>.
- [17] Maurino, A. (2025) 'How data quality affects machine learning models for credit risk assessment'. Available at: <https://arxiv.org/abs/2511.10964>.

- [18] Mupa, M.N., Tafirenyika, S., Nyajeka, M.R., Moyo, T.M. and Zhuwankinyu, E.K. (2025) 'Machine learning in actuarial science: Enhancing predictive models for insurance risk management', *Iconic Research and Engineering Journals*, 8(8), pp. 493-504.
- [19] National Institute of Standards and Technology (2023) *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: NIST.
- [20] Nguruve, T., Chitate, I.C., Munjoma, M.B., Makanjera, R., Mutimaamba, V.A., Masunda, M., Mhike, T.I. and Mupa, M.N. (2026) 'Operationalizing AI at scale: Repeatable frameworks for integration, adoption and performance measurement across enterprise and startup environments', *World Journal of Advanced Research and Reviews*, 30(2), pp. 688-695. doi: 10.30574/wjarr.2026.30.2.1236.
- [21] Nhova, M., Mupa, G., Tsveta, L., Dima, J., Chingezi, L., Madzinga, G.T. and Mupa, M.N. (2026) 'Enterprise risk management in health coverage operations: Integrating premium billing controls-first in analytics framework for eligibility-to-invoice integrity', *World Journal of Advanced Research and Reviews*, 30(2), pp. 1179-1185. doi: 10.30574/wjarr.2026.30.2.1269.
- [22] Okpala, I., Golgoon, A. and Kannan, A.R. (2025) 'Agentic AI systems applied to tasks in financial services: Modeling and model risk management crews'. Available at: <https://arxiv.org/abs/2502.05439>.
- [23] Papastergios, V., Ehrlinger, L. and Gounaris, A. (2025) 'Unfolding data quality dimensions in practice: A survey'. Available at: <https://arxiv.org/abs/2507.17507>.
- [24] Ribeiro, M.T., Singh, S. and Guestrin, C. (2016) 'Why should I trust you? Explaining the predictions of any classifier', *Proceedings of NAACL-HLT 2016: Demonstrations*, pp. 97-101. doi: 10.18653/v1/N16-3020.
- [25] Rudin, C. (2019) 'Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead', *Nature Machine Intelligence*, 1, pp. 206-215. doi: 10.1038/s42256-019-0048-x.
- [26] Schneider, J., Abraham, R., Meske, C. and vom Brocke, J. (2020) 'AI governance for businesses'. Available at: <https://arxiv.org/abs/2011.10672>.
- [27] Wand, Y. and Wang, R.Y. (1996) 'Anchoring data quality dimensions in ontological foundations', *Communications of the ACM*, 39(11), pp. 86-95. doi: 10.1145/240455.240479.
- [28] Wang, R.Y. and Strong, D.M. (1996) 'Beyond accuracy: What data quality means to data consumers', *Journal of Management Information Systems*, 12(4), pp. 5-33. doi: 10.1080/07421222.1996.11518099.
- [29] Weber, K., Otto, B. and Österle, H. (2009) 'One size does not fit all: A contingency approach to data governance', *ACM Journal of Data and Information Quality*, 1(1), pp. 1-27. doi: 10.1145/1515693.1515696.