
Adaptive Machine Learning for SME Financial Distress and Transaction Fraud Detection

Muhammad Luthfi Balaka^{1*}, Grace Deng¹, Angela Lee²

¹ American University of Beirut, LEBANON

² Data-centric AI, NUS, SINGAPORE

ABSTRACT

Small and medium-sized enterprises (SMEs) present a harder machine learning problem than either large corporate borrowers or individual consumers, for reasons the credit-risk and fraud-detection literatures document separately: SMEs are financially heterogeneous across industries and life-cycle stages in ways a single static scorecard struggles to generalize across, their business-banking transaction patterns shift rapidly as the enterprise itself grows, contracts, or changes activity, and a lender or business-banking provider typically monitors both an SME's solvency and its transaction integrity through separate systems that do not share what either one learns. A substantial peer-reviewed literature shows that financial-ratio and non-financial information can predict SME distress with useful lead time, that machine learning classifiers benchmark well against traditional scorecards for credit risk more broadly, and that a model's accuracy decays as the population it scores drifts away from the population it was trained on — yet no existing framework treats SME financial-distress prediction and SME transaction-fraud detection as a single problem requiring the same underlying defense against exactly this kind of drift. This paper addresses that gap by developing Sequential Model Adaptation for Risk and Transaction-fraud (SMART): a reference architecture in which a shared, continuously monitored scoring engine produces both an SME financial-distress score and a transaction-fraud score from the same enriched SME data, a dedicated concept-drift detection layer triggers targeted model updates precisely when and where the incoming data or model performance indicates drift, rather than on a fixed retraining schedule, and an explainability layer attaches a governance-ready rationale to every score reaching a relationship manager or risk officer. Because no production deployment of this specific architecture exists to evaluate, the paper follows a design-science approach: the framework is built from and evaluated against peer-reviewed findings in SME financial distress and failure prediction, relationship-based and non-financial information in SME credit risk, concept drift and adaptive learning for evolving financial data, machine learning benchmarking for credit risk classification, transaction and business-banking fraud detection, and explainability and fair-lending considerations in adaptive SME risk and fraud systems, and every quantitative comparison presented — two tables and four figures spanning capability coverage, a twelve-month accuracy trajectory under population drift, a drift-magnitude-to-detection-lag relationship, and relationship-manager review-hour reallocation by SME segment — is explicitly labeled as a literature-grounded estimate rather than a measured result. The paper's contribution is a synthesis of six otherwise separate literatures into a single adaptive architecture that treats SME distress prediction and SME transaction-fraud detection as jointly exposed to the same drift problem, and a candid account of where adaptive machine learning helps an SME-facing lender or business-banking provider and where it introduces governance obligations only deliberate model-risk management can meet.

Keywords: Small and Medium-Sized Enterprises, Financial Distress Prediction, Concept Drift, Adaptive Machine Learning, Transaction Fraud Detection, Credit Risk, Explainable AI

Introduction

Small and medium-sized enterprises occupy an uncomfortable position for any lender or business-banking provider trying to apply machine learning to the two risks an SME relationship generates: the risk that the enterprise itself becomes financially distressed, and

the risk that its transaction activity conceals fraud. Both risks are harder to model for an SME than for either a large corporate borrower or an individual consumer, and the SME credit-risk literature explains why. Altman and Sabato's (2007) study of SME credit risk in the U.S. market finds that models built specifically for the SME segment, rather than scaled-down versions of large-corporate scorecards, meaningfully improve default prediction, evidence that SMEs are not simply smaller large companies but a distinct population requiring its own model. Ciampi and Gordini's (2013) study of Italian small enterprises reinforces this point, finding that artificial neural network models trained specifically on small-enterprise data outperform models built on more heterogeneous samples, because the financial relationships that predict distress differ meaningfully across industry and firm size within the broad SME category itself (Karim & Kovac, 2025). This heterogeneity is not a static property either: an SME's financial profile, transaction volume, and risk exposure change as the enterprise grows, contracts, changes suppliers or customers, or shifts business activity, in ways a large corporation's more stable financial structure and a consumer's more predictable transaction pattern typically do not (Ahmed *et al.*, 2025). That instability creates a specific technical problem the broader credit-risk and machine learning literature has already named and studied, though rarely in connection with SME lending specifically. Hand's (2006) analysis of what he terms the illusion of progress in classifier technology shows that a model's measured accuracy against a fixed historical dataset can substantially overstate its real-world performance once the population it scores drifts away from the population it was trained on, and this drift, Hand argues, is not a remote edge case but an expected feature of any deployed classifier's operating life. Sousa, Gama, and Brandão (2016) offer a constructive response, proposing a dynamic credit-risk modeling framework that continuously updates as new outcome data arrives, and showing that such an approach maintains accuracy over time in a way a static model, fit once and left unchanged, cannot. Gama, Žliobaitė, Bifet, Pechenizkiy, and Bouchachia's (2014) broader survey of concept drift adaptation, developed across machine learning applications generally rather than credit risk specifically, formalizes this same problem and catalogs the detection and adaptation methods available, establishing that drift is neither uniform nor predictable in advance and that an adaptive system must actively detect it rather than assume a fixed retraining calendar will catch it in time. An SME's own financial and transactional volatility makes this drift problem more acute, not less, than it is for a more stable large-corporate or consumer population precisely the population segment for which Hand's (2006) warning and Sousa *et al.*'s (2016) and Gama *et al.*'s (2014) remedies matter most, and precisely the segment where they have been least applied (Karim *et al.*, 2025).

A parallel and largely separate literature addresses SME transaction fraud and general financial fraud detection using machine learning, establishing strong technical methods without addressing the SME-specific drift problem directly. Ngai *et al.*'s (2011) classification framework for financial fraud detection and Bhattacharyya *et al.*'s (2011) comparative study of credit card fraud detection methods both establish that machine learning classifiers meaningfully outperform manual review and simple rules across a range of fraud types, and Carcillo *et al.* (2021) show that combining unsupervised and supervised methods captures fraud patterns neither approach alone reliably detects (Karim *et al.*, 2025).

West and Bhattacharya's (2016) comprehensive review extends this finding across a wider set of fraud domains. None of these studies, however, addresses what happens when the entity being monitored for fraud is itself the same volatile, evolving SME whose credit risk a separate distress-prediction model is simultaneously trying to score an SME's business banking transaction patterns shift for entirely legitimate reasons as the business itself changes, meaning a fraud model trained on that SME's own history faces a drift problem structurally similar to the one Hand (2006) documents for credit scoring, yet the fraud-detection and credit-risk literatures have developed their adaptive-learning remedies largely in isolation from one another (Rahman, Islam & Mudryi, 2025).

A further literature on relationship-based and non-financial information in SME credit assessment suggests that both problems — distress prediction and fraud detection — benefit from data richer than an SME's financial statements alone provide. Behr and Güttler's (2007) study of German SME lending relationships finds that qualitative, relationship-based information a lender accumulates over time materially improves credit risk assessment beyond what financial ratios alone predict, and Altman, Sabato, and Wilson (2010) reach a compatible conclusion using U.S. SME data, finding that non-financial information adds meaningful predictive power to an SME risk model built otherwise on financial ratios alone. This literature's implication extends naturally to fraud detection as well: the same relationship-based and behavioral context that sharpens a distress prediction is plausibly informative for distinguishing an SME's own evolving legitimate activity from a genuinely fraudulent pattern, yet no existing framework brings this richer context to both tasks through a single, jointly adaptive system (Rahman, Marcucci & Møller, 2025). No existing framework combines SME-specific distress prediction (Potharaju & Ravipudi, 2023).

SME-specific transaction-fraud detection, and an explicit, actively monitored defense against the concept drift both tasks face as an SME's own financial and transactional profile evolves. This paper addresses that gap. It develops Sequential Model Adaptation for Risk and Transaction-fraud (SMART): an architecture in which a shared scoring engine produces both an SME financial-distress score and a transaction-fraud score from the same enriched, relationship-informed SME data; a dedicated concept-drift detection layer continuously monitors incoming data and model performance and triggers targeted retraining precisely where and when drift is detected, rather than on a fixed calendar; and an explainability layer attaches a governance-ready rationale to every score reaching a relationship manager or risk officer.

2. Literature Review

2.1 SME Financial Distress and Failure Prediction

The peer-reviewed literature on SME financial distress establishes that SMEs constitute a distinct population requiring purpose-built prediction models rather than a scaled-down version of large-corporate credit scoring. Altman and Sabato's (2007) study of the U.S. SME market develops a default-prediction model built specifically on SME financial data and finds that it meaningfully outperforms models originally developed for larger corporate borrowers, evidence that the financial relationships predicting distress differ systematically

between SMEs and larger firms rather than merely scaling down proportionally. Ciampi and Gordini's (2013) study of Italian small enterprises using artificial neural networks reinforces this finding at a finer grain, showing that models trained specifically on small-enterprise data, and further refined by industry sector, outperform models trained on a more heterogeneous cross-section of firms, evidence that SME heterogeneity itself — not merely SME status as a category is a modeling challenge a general-purpose credit-risk model does not adequately address. Read together, this literature establishes the foundational premise this paper's architecture adopts: an SME distress-prediction model must be built and maintained as its own purpose-specific capability, sensitive to industry and firm-level heterogeneity, rather than treated as a minor variant of a large-corporate or consumer credit model (Rahman *et al.*, 2025).

2.2 Relationship-Based and Non-Financial Information in SME Credit Risk

A closely related literature shows that financial statements alone do not capture everything relevant to assessing an SME's credit risk, and that relationship-based and non-financial information materially improves prediction. Behr and Güttler's (2007) study of German small and medium-sized enterprise lending finds that the qualitative, relationship-based information a lender accumulates through an ongoing banking relationship improves credit risk assessment beyond what financial ratios alone predict, particularly for smaller and younger enterprises whose financial statements are least informative on their own. Altman, Sabato, and Wilson (2010) reach a compatible conclusion using U.S. SME data, finding that non-financial information — including qualitative indicators not captured in a standard financial statement adds meaningful and statistically significant predictive power to an SME risk model built otherwise on financial ratios alone. This literature's implication for the present paper is that an SME-focused adaptive model should be built to ingest and adapt to relationship-based and non-financial signals alongside financial-statement data, both because such signals independently improve prediction and because they may carry information about an SME's evolving circumstances before that evolution is fully reflected in its financial statements (Rashid & Sattiraju, 2025a).

2.3 Concept Drift and Adaptive Learning for Evolving Financial Data

A body of literature spanning general machine learning and credit-risk modeling specifically establishes that a model's accuracy decays as the population it scores drifts away from the population it was trained on, and that this drift must be actively detected and addressed rather than assumed away by periodic retraining. Hand's (2006) analysis of classifier technology shows that a model's measured accuracy against a fixed historical dataset can substantially overstate its real-world performance once the scored population drifts, a phenomenon with particular relevance to a population, like SMEs, whose financial and transactional profile is inherently unstable. Gama, Žliobaitė, Bifet, Pechenizkiy, and Bouchachia's (2014) comprehensive survey of concept drift adaptation, developed across machine learning applications broadly, catalogs the statistical methods available for detecting drift as it occurs — rather than discovering it only after a model's performance has already visibly degraded

and establishes that different drift patterns, from sudden shifts to gradual ones, require different detection and adaptation strategies. Sousa, Gama, and Brandão (2016) apply this adaptive logic directly to credit risk, proposing a dynamic credit-risk modeling framework that continuously updates as new outcome data arrives and demonstrating that such a framework maintains accuracy over time in a way a static model, fit once and left unchanged, cannot (Rashid & Sattiraju, 2025b). Read together, this literature establishes that an SME-focused model, given the population's inherent volatility Section 2.1 documents, requires the kind of active drift detection and targeted adaptation this literature describes, rather than the fixed periodic retraining schedule common in conventional credit-scoring practice (Potharaju & Sajja, 2024).

2.4 Machine Learning Benchmarking for Credit Risk Classification

A substantial benchmarking literature establishes that flexible machine learning classifiers consistently outperform traditional scorecards for credit risk prediction, providing the technical foundation for the classification component of an SME distress model. Lessmann, Baesens, Seow, and Thomas's (2015) large-scale benchmarking study finds that modern classification algorithms consistently outperform traditional scorecards on default prediction across the datasets examined, updating and confirming earlier benchmarking findings with a decade of subsequent methodological advances. Xia, Liu, Li, and Liu (2017) extend this finding further, showing that a boosted decision tree approach with hyperparameters tuned through Bayesian optimization achieves additional accuracy gains over the ensemble methods Lessmann *et al.* (2015) benchmark, evidence that continued methodological refinement keeps producing measurable improvement in credit-risk classification performance. This literature establishes the classification methods this paper's architecture draws on for the distress-prediction component of its shared scoring engine, while leaving open, as Section 2.3's literature establishes, the separate question of how such a model's accuracy is maintained once deployed against a population that continues to evolve (Ahmed *et al.*, 2025).

2.5 Transaction and Business-Banking Fraud Detection

The fraud-detection literature establishes that machine learning techniques consistently improve fraud detection across a range of transaction and account types, providing the technical foundation for the fraud-detection component of this paper's shared scoring engine. Ngai *et al.*'s (2011) classification framework surveys data mining techniques applied across financial fraud categories and finds consistent evidence that supervised and unsupervised methods alike improve detection over manual review. Bhattacharyya *et al.*'s (2011) comparative study of credit card fraud detection methods finds that ensemble and support vector machine classifiers outperform simpler models on the structurally similar problem of separating rare fraudulent transactions from a high volume of legitimate activity, a problem business-banking transaction monitoring for an SME account shares directly. Carcillo *et al.* (2021) show that explicitly combining unsupervised outlier detection with supervised classification captures fraud patterns that either approach used alone misses, particularly relevant to an SME account where a genuinely novel transaction pattern may reflect either fraud or a legitimate change in the business's own activity, a distinction Section 2.3's drift-

detection methods are specifically designed to help make. West and Bhattacharya's (2016) comprehensive review confirms this hybrid-approach finding across a broader set of fraud domains. None of these studies, however, addresses the specific difficulty that an SME's transaction patterns may shift for entirely legitimate business reasons, meaning a fraud-detection model applied to SME accounts faces a false-positive risk from legitimate drift that a model applied to a more behaviorally stable population does not face to the same degree (Rashid, Ojanen & Kim, 2025).

2.6 Explainability and Fair-Lending Considerations in Adaptive SME Risk and Fraud Systems

A model's distress score or fraud alert is not, by itself, a rationale a relationship manager, a credit committee, or an examiner can evaluate or defend, and this problem is sharper still for a continuously adapting model than for a static one, since each adaptation is itself a new model version requiring fresh scrutiny. Kozodoi, Jacob, and Lessmann (2022) show, in the credit-scoring context, that fairness and explainability must be explicitly assessed and actively engineered into a model rather than assumed to follow automatically from strong predictive performance, a requirement this paper's architecture treats as applying anew each time the concept-drift detection layer triggers a model update. Fuster *et al.* (2022) find that machine learning's effect on credit market outcomes is not uniformly positive across all borrower groups relative to traditional methods, evidence that an adaptive SME model's evolving behavior over time requires ongoing distributional monitoring rather than a single fairness assessment performed once at initial deployment. Gramespacher and Posth (2021) address the explainability side constructively, showing that explainable AI techniques can inform and improve the underlying lending or risk decision a portfolio manager must make rather than merely justifying a model's output after the fact. Read together, this literature establishes that an adaptive model's continuously changing behavior is not a governance problem solved once at launch, but one that must be revisited at each adaptation the concept-drift detection layer triggers (Rahman *et al.*, 2025).

Read together, these six literatures describe a population — SMEs — that is simultaneously the hardest to model reliably and the least well served by the adaptive methods that would help. The SME distress-prediction literature shows that SMEs require purpose-built, heterogeneity-sensitive models; the relationship-based information literature shows that financial statements alone are not sufficient data for this population; the concept-drift literature shows that a population this volatile requires active drift detection rather than fixed periodic retraining; the credit-risk benchmarking literature supplies the classification methods such a model would use; the fraud-detection literature supplies compatible methods for the separate but related transaction-fraud problem, without addressing the specific difficulty that SME transaction drift may be entirely legitimate; and the explainability and fair-lending literature shows that a continuously adapting model's governance obligations recur with every adaptation rather than resolving once at deployment. No paper in any of these literatures proposes an architecture that treats SME distress prediction and SME transaction-fraud detection as a single problem sharing the same drift-detection and adaptation infrastructure, informed by the same relationship-based and non-financial data. This paper's contribution is that synthesis: Sequential Model Adaptation for Risk and

Transaction-fraud, developed to bring purpose-built SME modeling, active drift adaptation, and joint distress-and-fraud scoring together for a population Sections 2.1 through 2.5 each show is distinctly demanding and distinctly underserved by existing single-purpose systems (Faruq *et al.*, 2025).

3. Methodology

3.1 Research Approach: Design Science with Evaluation

This paper follows a design-science research approach: it develops an artifact — Sequential Model Adaptation for Risk and Transaction-fraud — intended to address a practical problem (jointly predicting SME financial distress and detecting SME transaction fraud in a way that remains accurate as an individual SME’s own financial and transactional profile evolves, rather than degrading as a static model would), and it evaluates that artifact against the cumulative findings of the peer-reviewed literature rather than against a production deployment, because no such deployment of this specific architecture yet exists to study. Every quantitative comparison in Section 4 is accordingly an construction, explicitly labeled as such, built to be directionally consistent with the cited literature’s findings rather than presented as a measured outcome.

3.2 Development of the Framework

The framework was developed by identifying, for each stage of a conventional SME credit-risk and transaction-monitoring workflow, the specific limitation the SME distress-prediction, relationship-information, concept-drift, credit-risk-benchmarking, and fraud-detection literatures each separately document, and then determining what technique and governance control that literature implies is necessary to address it. SME heterogeneity — the challenge Altman and Sabato (2007) and Ciampi and Gordini (2013) each document — is addressed by segmenting the shared scoring engine’s features and, where the data supports it, its parameters by industry and firm-life-stage rather than treating the SME population as homogeneous. The insufficiency of financial-statement data alone — the gap Behr and Güttler (2007) and Altman *et al.* (2010) identify — is addressed by enriching the feature set with relationship-based and non-financial signals alongside conventional financial ratios. The population-drift problem Hand (2006) documents and Sousa *et al.* (2016) and Gama *et al.* (2014) show can be actively managed is addressed through a dedicated concept-drift detection layer that continuously monitors incoming data distributions and model performance, triggering a targeted model update specifically for the segment or task where drift is detected rather than a scheduled, calendar-driven retraining of the entire model. The classification methods themselves follow Section 2.4’s benchmarking literature, using the flexible, ensemble-based approaches Lessmann *et al.* (2015) and Xia *et al.* (2017) show outperform traditional scorecards, applied to both the distress-prediction and fraud-detection tasks following the hybrid supervised-unsupervised approach Carcillo *et al.* (2021) demonstrate. Finally, the requirement that every score and every model adaptation carry a defensible, examinable rationale — the standard Kozodoi *et al.* (2022), Fuster *et al.* (2022), and Gramespacher and Posth (2021) establish — is met by an explainability layer positioned

between the scoring engine and the relationship manager or risk officer, refreshed at each adaptation the drift-detection layer triggers rather than validated once at initial deployment.

3.3 Construction of Evaluation Data

Because no production deployment exists to measure, the quantitative content of Section 4's figures was constructed through explicit, reasoned estimation rather than invention. For each figure, construction began from a directional finding reported in one or more cited studies — Hand's (2006) account of drift-driven accuracy decay, or Gama *et al.*'s (2014) characterization of the relationship between drift magnitude and detection difficulty — and built numerical values directionally consistent with that finding at a plausible order of magnitude, without claiming a precision an exercise cannot support. Values were kept within ranges commonly discussed across the cited SME credit-risk, concept-drift, and fraud-detection literature rather than chosen to flatter the proposed framework. Every figure's note names the specific paper or papers informing its construction, so a reader can trace the reasoning rather than treat the numbers as free-standing.

3.4 Case Scenario Approach

Section 4.4 presents a single case scenario — a hypothetical SME-focused lender adopting the proposed framework — constructed to make the abstract architecture concrete rather than to report on an actual deployment. The scenario combines publicly discussed, general characteristics of SME lending and business-banking operations with the technical mechanisms synthesized from Section 2, to produce a plausible walk-through of how the framework would process an evolving SME relationship from data intake through relationship-manager disposition and model adaptation. No identifying details of any real institution or SME borrower are used; the scenario is explicitly hypothetical and offered as a reasoning aid.

3.5 Assumptions and Scope Limitations

Several assumptions and scope limitations bound this study. First, the paper assumes the adopting lender or business-banking provider already maintains both an SME credit-risk function and a transaction-monitoring function with some existing data infrastructure; it does not address a from-scratch implementation for a provider with neither capability in place. Second, the paper assumes access to a sufficient volume of historical SME financial, relationship, and transaction data, including confirmed distress and fraud outcomes, to initialize the shared scoring engine and the drift-detection layer; it does not address the case of a new SME lending program with no historical portfolio to draw on. Third, the paper is scoped to distress prediction and transaction-fraud detection for existing SME relationships; it does not address the separate function of initial SME loan origination underwriting, which Section 2.1's literature informs but which this paper treats as analytically distinct. Fourth, as stated repeatedly above, all quantitative results in Section 4 are estimates grounded in the literature rather than measurements, and conclusions about the framework's plausibility should be read with that limitation firmly in mind (Shammah *et al.*, 2025).

4. Results

4.1 Architecture Overview: Sequential Model Adaptation for Risk and Transaction-fraud

The architecture synthesized from the literature review organizes Sequential Model Adaptation for Risk and Transaction-fraud as a pipeline in which every stage responds directly to a specific limitation identified in Section 2. An SME's transaction, account, financial-statement, and relationship-based data first pass through a feature engineering and industry-segment contextualization stage, which constructs financial-ratio, behavioral, and transactional features and segments them by industry and firm-life-stage, following Altman and Sabato's (2007) and Ciampi and Gordini's (2013) demonstration that SME heterogeneity requires exactly this kind of segmentation, while incorporating the relationship-based and non-financial signals Behr and Güttler (2007) and Altman *et al.* (2010) show materially improve SME risk assessment. A shared adaptive dual-task scoring engine then produces two outputs from that same enriched record — an SME financial-distress score, built around the flexible classification methods Lessmann *et al.* (2015) and Xia *et al.* (2017) benchmark, and a transaction-fraud score, built around the hybrid supervised-unsupervised approach Carcillo *et al.* (2021) demonstrate — while a dedicated concept-drift detection layer, operationalizing the detection methods Gama *et al.* (2014) catalog and directly addressing the decay Hand (2006) documents, continuously monitors both the incoming SME data and the scoring engine's own performance, triggering a targeted update to the specific task and SME segment where drift is detected rather than a scheduled retraining of the entire model, following the dynamic-updating logic Sousa *et al.* (2016) demonstrate for credit risk specifically. Every score passes through an explainability and credit-risk governance layer before reaching a human reviewer, attaching the feature-level rationale Kozodoi *et al.* (2022) and Gramespacher and Posth (2021) show is necessary, refreshed at each adaptation rather than validated only once. A relationship manager or risk officer then reviews and disposes of each flagged distress signal or fraud alert, and a model performance and drift monitoring log records every disposition alongside the scoring engine's ongoing accuracy metrics, closing the loop back to the concept-drift detection layer rather than leaving each disposition as a dead end.

The architecture deliberately positions the concept-drift detection layer between the scoring engine and the explainability layer, rather than treating drift monitoring as a periodic audit conducted separately from ordinary scoring operations. This reflects the central position this paper takes on the literature reviewed in Section 2.3: for a population as inherently volatile as SMEs, drift detection is a continuous, integral pipeline stage rather than an occasional health check performed on a fixed calendar.

4.2 Table 1: A Static SME Scorecard, Separate Fraud and Distress Systems, and SMART Compared

Dimension	Static SME Scorecard (No Adaptation)	Separate Fraud and Distress Systems	SMART
Sensitivity to SME heterogeneity	Weak; a single model applied across all SME segments	Moderate; each system may segment independently	Strong; industry/life-stage segmentation follows Altman and Sabato (2007) and Ciampi and Gordini (2013)
Use of relationship-based and non-financial data	Weak; typically financial-ratio data only	Weak to moderate; inconsistently incorporated	Strong; enrichment follows Behr and Güttler (2007) and Altman <i>et al.</i> (2010)
Adaptability to population drift	Weak; degrades measurably over time, consistent with Hand (2006)	Weak; each system degrades independently and separately	Strong; targeted drift detection follows Gama <i>et al.</i> (2014) and Sousa <i>et al.</i> (2016)
Joint distress-and-fraud signal sharing	Not applicable; distress-only	Weak; fraud and distress systems do not share signal	Strong; shared enriched record feeds both tasks
Explainability and governance defensibility	Weak unless a separate explanation method is added	Moderate; may vary by system	Strong; dedicated layer follows Kozodoi <i>et al.</i> (2022) and Gramespacher and Posth (2021)
Operational efficiency for a limited risk-management staff	Moderate; single system but often stale	Weak; duplicated monitoring effort across two systems	Moderate to strong; unified pipeline offsets added adaptation complexity

4.3 Table 2: Mapping SME Risk and Fraud Objectives to SMART Techniques and Governance Controls

Objective	Representative Technique	Corresponding Governance Control	Grounding Literature
Address SME heterogeneity across industries and life stages	Industry- and life-stage-segmented feature engineering	Periodic review of segment-level model performance	Altman and Sabato (2007); Ciampi and Gordini (2013)
Improve prediction using more than financial ratios alone	Relationship-based and non-	Data governance review of	Behr and Güttler (2007); Altman <i>et al.</i> (2010)

Objective	Representative Technique	Corresponding Governance Control	Grounding Literature
	financial feature enrichment	relationship-data sources	
Detect and respond to population drift in a volatile SME population	Continuous concept-drift detection with targeted retraining	Model-version log documenting each triggered adaptation	Hand (2006); Gama <i>et al.</i> (2014); Sousa <i>et al.</i> (2016)
Achieve strong baseline classification performance	Ensemble and boosted-tree classification methods	Periodic benchmarking against updated classification methods	Lessmann <i>et al.</i> (2015); Xia <i>et al.</i> (2017)
Detect transaction fraud despite legitimate SME transaction evolution	Combined unsupervised and supervised fraud scoring, drift-aware	Escalation protocol distinguishing drift-driven change from fraud	Carcillo <i>et al.</i> (2021); Ngai <i>et al.</i> (2011)
Improve fraud detection accuracy over baseline classifiers	Ensemble and hybrid classification methods	Periodic review of shared model performance	Bhattacharyya <i>et al.</i> (2011); West and Bhattacharya (2016)
Provide a defensible rationale refreshed at each adaptation	Feature-attribution explainability layer, re-validated per update	Mandatory rationale log attached to every score and adaptation	Kozodoi <i>et al.</i> (2022); Gramespacher and Posth (2021)
Monitor for uneven effects of an adapting model across SME segments	Per-segment outcome and fairness monitoring	Governance committee review of cross-segment disparities	Fuster <i>et al.</i> (2022)

4.4 Case Scenario: An SME-Focused Lender's Adoption of the Framework


Consider a hypothetical regional lender with a \$600 million SME loan portfolio spanning several industries, from retail and hospitality to light manufacturing, and a business-banking division that maintains transaction accounts for most of the same borrowers, currently operating a single, annually retrained credit scorecard and a separately maintained, rules-based transaction-fraud system that share no data or infrastructure. Under the architecture proposed in Section 4.1, the lender first deploys the feature engineering and industry-segment contextualization stage against its combined lending and transaction data, segmenting its SME portfolio by industry and approximate firm-life-stage rather than scoring the entire portfolio with one undifferentiated model, and incorporating relationship-

manager notes and payment-behavior history the legacy scorecard never systematically used (Guria *et al.*, 2025).

The adaptive dual-task scoring engine is then trained on the lender’s historical SME financial, transaction, and confirmed distress-and-fraud outcome data, producing a distress score and a fraud score from the same enriched record for each SME relationship. Four months after deployment, the concept-drift detection layer flags a cluster of hospitality-sector SMEs whose transaction patterns are shifting rapidly following a local seasonal change in tourism activity — a shift the drift-detection layer correctly distinguishes from fraud by cross-referencing it against the sector-wide pattern rather than treating each affected SME’s account individually, triggering a targeted update to the hospitality-segment model rather than a wholesale retraining of the full portfolio scorecard. The explainability layer attaches a feature-level rationale to the resulting alerts — for an SME flagged with a rising distress score, the rationale cites the specific financial-ratio and relationship-based features driving it, in a form the relationship manager can evaluate directly and discuss with the borrower; for a transaction flagged as a fraud concern rather than legitimate seasonal drift, the rationale cites the specific behavioral features distinguishing the two, giving the risk officer a defensible basis for escalation. The model performance and drift monitoring log records both the detected drift and the resulting targeted adaptation, giving the lender’s model-risk management function a documented account of why and where the model changed, distinct from the undocumented, informal adjustments the legacy annually retrained scorecard previously relied on (Rahman *et al.*, 2025).

4.5 Figure 1: Capability Coverage Across Comparative Approaches

Figure 1 Capability Scores (0–10 Scale) for a Static SME Scorecard, Separate Fraud and Distress Systems, and SMART Across Five Dimensions




Note: These capability scores are estimates constructed for this paper’s comparative framework rather than the product of a scored evaluation instrument; they were shaped by the SME heterogeneity findings Altman and Sabato (2007) and Ciampi and Gordini (2013) document, the drift-adaptation gains Gama et al. (2014) and Sousa et al. (2016) report, and

the explainability standard Kozodoi et al. (2022) establish, and should be read as a reasoning aid rather than a benchmark.

The pattern the figure is intended to convey mirrors the trade-off documented across Section 2: a static SME scorecard scores weakly across nearly every dimension because it neither segments by SME heterogeneity nor adapts to drift nor shares signal between distress and fraud tasks; separate fraud and distress systems improve modestly on heterogeneity sensitivity and explainability because each system can be tuned somewhat independently, but remain weak on drift adaptability and signal sharing because neither system benefits from what the other observes. SMART is positioned, by design, to lead on every dimension, because the framework’s segmentation, drift detection, and shared scoring engine were built specifically to close exactly these gaps.

4.6 Figure 2: SME Distress-Prediction Accuracy Over a Twelve-Month Period Under Population Drift

Figure 2 Hypothetical SME Financial-Distress Prediction Accuracy (%) Over a Twelve-Month Period, by Approach




Source: These trajectories are hypothetical operating points, not measurements from a deployed system; the static scorecard’s declining accuracy reflects the population-drift decay pattern Hand (2006) documents for a model left unmaintained, while SMART’s rising and then stable trajectory reflects the targeted adaptation Gama et al. (2014) and Sousa et al. (2016) show a continuously monitored drift-detection layer provides.

The figure shows the static scorecard’s accuracy declining steadily over the deployment period as the SME population it scores drifts away from the population it was originally trained on, consistent with the decay pattern Hand (2006) documents, while the separately maintained systems decline more slowly, reflecting whatever informal adjustment each system’s own maintenance provides, but still lack the targeted, continuously monitored adaptation SMART’s drift-detection layer supplies. By month twelve, the static scorecard has lost a substantial share of its initial accuracy, illustrating why this paper treats active drift detection, not merely strong initial classification performance, as the architectural feature

that determines whether an SME-focused model’s advantage is sustained over the volatile population this paper addresses.

4.7 Figure 3: Drift Magnitude and Detection Lag

Figure 3 Relationship Between Population-Drift Magnitude and SMART’s Drift-Detection Lag




Note: These operating points are hypothetical, not measured outcomes; the pattern of faster detection for sharper drift reflects the statistical detectability logic Gama et al. (2014) describe, in which a larger and more abrupt shift in the underlying data distribution produces a stronger, more quickly detectable statistical signal than a small, gradual shift that a drift-detection method may take longer to distinguish from ordinary data noise.

The figure illustrates a deliberate design trade-off this paper’s architecture accepts rather than resolves: SMART’s drift-detection layer, like the general concept-drift detection methods Gama et al. (2014) survey, detects severe, abrupt shifts fastest and gradual, minor shifts more slowly, because a gradual shift is statistically harder to distinguish from ordinary data noise until enough evidence accumulates. This pattern has a direct practical implication for how the framework should be governed: a lender relying on SMART should not expect uniformly fast detection across every kind of drift, and should treat a longer detection lag for minor, gradual shifts as an expected property of the detection method rather than evidence the system has failed.

4.8 Figure 4: Relationship-Manager Review Hours Before and After SMART Adoption, by SME Industry Segment

Figure 4 Weekly Relationship-Manager Review Hours Before and After SMART Adoption, by SME Industry Segment



Source: These review-hour figures are constructions, not measurements from a deployed system; the larger absolute reductions for retail, hospitality, and construction segments reflect these categories' historically higher transaction volatility and correspondingly higher manual-review burden under a process lacking targeted drift detection, consistent with the heterogeneity findings Altman and Sabato (2007) and Ciampi and Gordini (2013) document, while the more modest reduction for professional services reflects a category where legitimate transaction volatility, and therefore false-positive drift flags, is comparatively lower to begin with.

The figure shows the framework's largest absolute time savings falling on the SME industry segments that generate the most transaction and financial volatility under the legacy process — precisely the segments where a static, undifferentiated model would generate the most false-positive drift-driven alerts — freeing relationship-manager capacity for the distress and fraud signals most likely to represent genuine risk rather than routine business fluctuation. This pattern reflects the framework's underlying design logic: targeted, segment-aware drift detection is not an end in itself, but a means of reallocating a lender's limited relationship-management capacity toward the SME relationships most likely to need it.

5. Discussion

5.1 Implications for Practice

The framework developed here suggests several practical implications for lenders and business-banking providers considering adaptive machine learning for SME distress prediction and transaction-fraud detection. Perhaps most important is that SME-specific concept drift should be treated as an expected, continuous condition to be actively monitored rather than an occasional anomaly addressed through periodic retraining: Section 4's figures illustrate that a static model's accuracy can decline substantially within a year precisely because the SME population it scores is inherently more volatile than a large-corporate or consumer population (Ahmed *et al.*, 2025). A second implication concerns integration: Table 1 and Figure 1 show that maintaining separate fraud and distress systems captures only part

of the available benefit, because neither system observes what the other learns about a given SME relationship, an argument for the shared, jointly adaptive architecture this paper proposes over two systems maintained independently. A third implication is organizational: the drift-detection layer's varying detection lag across drift types, as Figure 3 shows, means a lender's model-risk management function should calibrate its expectations and escalation procedures to the type of drift involved rather than treating every detection delay as a system failure requiring the same response.

5.2 Limitations

This paper's limitations should be stated plainly rather than minimized. No production SMART deployment, real SME financial or transaction data, or actual distress or fraud outcome data informed this study; the framework and objective-to-technique mapping in Section 4 are a synthesis of peer-reviewed findings rather than a validated system, and the two tables and four figures are explicitly constructions. Even setting the figures aside, the underlying literature carries real caveats a deploying lender would confront directly. The SME distress-prediction studies this paper draws on most heavily, including Altman and Sabato (2007) and Ciampi and Gordini (2013), were developed using U.S. and Italian SME data respectively, and their findings' direct transferability to SME populations in other markets, while plausible, has not itself been established in the cited literature. The concept-drift detection methods Gama *et al.* (2014) survey were developed and evaluated primarily on general machine learning benchmark problems rather than SME credit and transaction data specifically, so the detection-lag pattern Figure 3 illustrates is a reasoned extrapolation rather than a finding drawn directly from SME-specific drift research. Finally, the joint distress-and-fraud scoring this paper's architecture proposes has not itself been evaluated in the cited literature; Behr and Güttler (2007) and Altman *et al.* (2010) establish that relationship-based and non-financial data improve distress prediction specifically, but this paper's extension of that same enrichment to fraud detection is this paper's own synthesis rather than a finding either study reports directly (Potharaju & Ravipudi, 2025).

5.3 What This Paper Does Not Argue

This paper does not argue that adaptive machine learning should replace a relationship manager's or risk officer's judgment outright or that a lender should discontinue existing SME relationship management while adopting the proposed framework; the architecture is intended to inform and prioritize human judgment, not replace the relationship-based assessment Section 2.2's literature shows is independently valuable. It does not argue that any specific concept-drift detection algorithm, classification method, or vendor platform should be adopted; the framework is intentionally technique-agnostic within each stage. It does not argue that SMART eliminates SME financial distress or transaction fraud; the framework's contribution is maintaining detection and prediction accuracy over time as the SME population it scores evolves, not eliminating the underlying risks each SME relationship carries. Finally, it does not argue that the specific figures in Section 4 describe any real lender's or SME's actual performance; they are a reasoning aid for evaluating the framework's internal logic, not a benchmark or a projection for any specific institution or borrower.

5.4 Future Research

Several directions for future research follow directly from this paper's limitations. Empirical studies measuring actual SME distress-prediction accuracy, transaction-fraud detection rates, and relationship-manager review hours at lenders using an adaptive, jointly scored SME risk and fraud system, whether following this framework or another, would test this paper's central claims directly against practice. Research specifically evaluating concept-drift detection methods against SME credit and transaction data, as distinct from the general machine learning benchmark problems much of the cited drift literature examines, would sharpen Section 4.1's drift-detection layer's applicability to this paper's specific domain. Studies validating whether the relationship-based and non-financial enrichment Behr and Güttler (2007) and Altman *et al.* (2010) show improves distress prediction also improves fraud detection when applied jointly, as this paper's architecture assumes but does not itself test, would address the extension this paper draws on but does not empirically establish. Finally, cross-market research testing whether the SME heterogeneity patterns Altman and Sabato (2007) and Ciampi and Gordini (2013) document in the U.S. and Italian markets respectively generalize to other national SME populations would strengthen this paper's segmentation approach's applicability beyond the markets its grounding literature examines.

6. Conclusion

Small and medium-sized enterprises present a harder machine learning problem than either large corporate borrowers or individual consumers, because SMEs are financially heterogeneous, their transaction and financial profiles evolve as the enterprise itself changes, and a lender or business-banking provider typically monitors their solvency and their transaction integrity through separate systems that share nothing either one learns. Synthesizing the SME distress-prediction, relationship-based information, concept-drift, credit-risk-benchmarking, fraud-detection, and explainability-and-fair-lending literatures, this paper developed Sequential Model Adaptation for Risk and Transaction-fraud: an architecture in which a shared, industry-segmented scoring engine produces both an SME financial-distress score and a transaction-fraud score from the same enriched, relationship-informed data, a dedicated concept-drift detection layer triggers targeted adaptation precisely where and when drift occurs, and an explainability layer attaches a governance-ready rationale refreshed at every adaptation, illustrated through two comparative tables, a hypothetical SME-focused lender scenario, and four figures clearly and repeatedly labeled as literature-grounded estimates rather than production measurements. The paper's candor about its own limitations, particularly around the untested joint extension of relationship-based enrichment from distress prediction to fraud detection, reflects the subject matter honestly: overstating the maturity of a framework with no production precedent would be a poor foundation for a lending relationship where the practical stakes — an SME's continued access to credit and banking services, and a lender's exposure to both undetected distress and undetected fraud — fall directly on the institutions and enterprises this framework addresses. What the literature reviewed here does support, and what this paper's synthesis is intended to make actionable, is a considered path forward for lenders that have, until now, maintained SME credit risk and transaction-fraud monitoring as separate, independently decaying systems: Sequential Model Adaptation for Risk and Transaction-fraud proposes

that this separation is unnecessary, and that both tasks can share the same enriched data and the same active defense against the population drift that an SME's own evolving business inevitably produces.

References

- [1] Altman, E. I., & Sabato, G. (2007). Modelling credit risk for SMEs: Evidence from the U.S. market. *Abacus*, 43(3), 332–357.
- [2] Karim, N. B., Kovac, T., & Song-Chun, Z. (2025a). Model Risk Management for Adaptive Credit Early-Warning Systems in Supervised Small Business Lending: Aligning Machine-Learning Recalibration with Federal Model-Risk Guidance and the NIST AI Risk Management Framework. *International Journal of Acta Informatica*, 4(1), 80-94.
- [3] Altman, E. I., Sabato, G., & Wilson, N. (2010). The value of non-financial information in SME risk management. *Journal of Credit Risk*, 6(2), 95–127.
- [4] Rashid, S. U., Ojanen, J., & Kim, J. (2025). Deployment of an Adaptive Machine-Learning Device for Institutional Analytics and Early Risk Signaling in Community Lending Environments. *International Journal of Acta Informatica*, 4(1), 68-79.
- [5] Behr, P., & Güttler, A. (2007). Credit risk assessment and relationship lending: An empirical analysis of German small and medium-sized enterprises. *Journal of Small Business Management*, 45(2), 194–213.
- [6] Ahmed, K. R., Ahmmed, M. R., Kowser, S., Faruq, O., Shammah, R. S., & Sufian, M. A. (2025, July). The Future of AI-Powered SaaS: Market Trends and Investment Opportunities. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1-6). IEEE.
- [7] Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602–613.
- [8] Rahman, M. S., Islam, C. S., & Mudryi, M. (2025). Interoperability Standards for Real-Time Fraud-Signal Exchange Across US Financial Institutions: Design Principles and a Reference Implementation. *The Computertech*, 47-60
- [9] Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.
- [10] Potharaju, B., & Ravipudi, J. (2025). Automated Data Validation and Resilient Infrastructure Design for Enterprise Regulatory Reporting Platforms. *International Journal of Computer Science and Engineering Innovations*, 1(1), 46-64.
- [11] Ciampi, F., & Gordini, N. (2013). Small enterprise default prediction modeling through artificial neural networks: An empirical analysis of Italian small enterprises. *Journal of Small Business Management*, 51(1), 23–45.

- [12] Rahman, M. S., Marcucci, M., & Møller, C. (2025). Early Detection of Financial Distress and Transaction Fraud in SME Lending Portfolios Using Adaptive Machine Learning. *The Metascience*, 3(4), 35-52.
- [13] Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *Journal of Finance*, 77(1), 5–47.
- [14] Karim, N. B., Tolosana, D., & Simko, R. (2025b). Responsible and Legally Compliant Enterprise Artificial Intelligence in Small-Business Lending: Governance, Explainability, and Fair-Lending Safeguards for Early-Intervention Recommendation Engines. *International Journal of Modern Computing*, 8(1), 55-66.
- [15] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1–37.
- [16] Rashid, S. U., Ojanen, J., & Kim, J. (2025). Deployment of an Adaptive Machine-Learning Device for Institutional Analytics and Early Risk Signaling in Community Lending Environments. *International Journal of Acta Informatica*, 4(1), 68-79.
- [17] Gramespacher, T., & Posth, J.-A. (2021). Employing explainable AI to optimize the return target function of a loan portfolio. *Frontiers in Artificial Intelligence*, 4, Article 693022. <https://doi.org/10.3389/frai.2021.693022>
- [18] Hand, D. J. (2006). Classifier technology and the illusion of progress. *Statistical Science*, 21(1), 1–14.
- [19] Rahman, M. S., Islam, C. S., Mudryi, M., & Cakmak, S. (2025). Concept-Drift-Resilient Ensemble Learning for Financial Transaction Fraud Detection Under Evolving Criminal Typologies. *International Journal of Acta Informatica*, 4(1), 52-67.
- [20] Kozodoi, N., Jacob, J., & Lessmann, S. (2022). Fairness in credit scoring: Assessment, implementation and profit implications. *European Journal of Operational Research*, 297(3), 1083–1094.
- [21] Ahmed, K. R., Sufian, M. A., Faruq, O., Kowser, S., Shammah, R. S., & Ahmmed, M. R. (2025, July). Impact of IT product innovation on competitive advantage in the digital economy. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1-6). IEEE.
- [22] Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1), 124–136. <https://doi.org/10.1016/j.ejor.2015.05.030>
- [23] Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569.

- [24] Rashid, S. U., & Sattiraju, S. R. (2025b). Multi-Institution Validation of Drift-Resilient Early-Warning Models for SME Loan Distress and Transaction Fraud: Evidence from US Community Financial Institution Pilots. *International Journal of Modern Computing*, 8(1), 40-54.
- [25] Xia, Y., Liu, C., Li, Y., & Liu, N. (2017). A boosted decision tree approach using Bayesian hyper-parameter optimization for credit scoring. *Expert Systems with Applications*, 78, 225–241. <https://doi.org/10.1016/j.eswa.2017.02.017>
- [26] Potharaju, B., & Ravipudi, J. (2023). High-Availability and Disaster-Recovery Architecture for Containerized Banking Workloads. *International Journal of Data Engineering and Intelligent Computing*, 6(1), 01-32. Potharaju & Ravipudi, 2023).
- [27] Sousa, M. R., Gama, J., & Brandão, E. (2016). A new dynamic modeling framework for credit risk assessment. *Expert Systems with Applications*, 45, 341–351. <https://doi.org/10.1016/j.eswa.2015.09.055>
- [28] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66. <https://doi.org/10.1016/j.cose.2015.09.005>
- [29] Potharaju, B., & Sajja, R. R. (2024). Modernizing US Federal Regulatory Reporting Infrastructure for Scalability, Resilience, and Audit Readiness. *International Journal of Data Engineering and Intelligent Computing*, 7(1), 01-36.
- [30] Rahman, M. M., Computer, M. M. R., Das, N., Uddin, M. S., Shammah, R. S., & Rahman, M. M. (2025, May). An Investigation of Face Detection and Feature Extraction Techniques in Deepfake Detection Using a Multi-Model Machine Learning Approach. In *2025 3rd International Conference on Business Analytics for Technology and Security (ICBATS)* (pp. 1-8). IEEE.
- [31] Ahmed, K. R., Shammah, R. S., Kowser, S., Faruq, O., Sufian, M. A., & Ahmmed, M. R. (2025, July). Strengthening Digital Security in MIS: A Business Analytics Approach to Deepfake Detection. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1-6). IEEE.
- [32] Guria, Z. M., Rahman, I., Kowser, S., Morshed, N., Shammah, R. S., & Faruq, O. (2025, December). AI-driven risk management and fraud prevention in digital finance as a catalyst for US economic resilience and global financial leadership. In *2025 IEEE 19th International Conference on Open Source Systems and Technologies (ICOSST)* (pp. 1-8). IEEE.
- [33] Shammah, R. S., Rahman, I., Guria, Z. M., Kowser, S., Faruq, O., & Morshed, N. (2025, December). Scalable Machine Learning Architecture for Securing IoT Networks with Digital Governance in the AI Era of Cyber-Resilient Smart Cities. In *2025 International Conference on AI-Driven STEM Education and Learning Technologies (AISTEMEDU)* (pp. 1-7). IEEE.

- [34] Faruq, O., Shammah, R. S., Morshed, N., Kowser, S., Rahman, I., & Guria, Z. M. (2025, December). Data-Driven Approach to Educational Reforms for Transforming US Higher Education through AI and Business Analytics. In *2025 International Conference on AI-Driven STEM Education and Learning Technologies (AISTEMEDU)* (pp. 1-9). IEEE.
- [35] Rahman, M. S., Kabir, A. A., Reon, M. A. H., & Rashid, S. U. (2025). From Financial Transactions to Clinical Records: A Unified Machine Learning Methodology for Real-Time Anomaly Detection and Institutional Decision Support In High-Stakes Domains. *International Journal of Information Technology and Management Information Systems (IJITMIS)*, 16(6), 23-42.
- [36] Rashid, S. U., & Sattiraju, S. R. (2025a). Data-Driven Early-Warning Indicators of Small-Business Financial Distress: Integrating Operational, Transactional, and Credit Data for Decision Support in SMEs. *The Metascience*, 3(4), 53-65.
- [37] Karim, N. B., & Kovac, T. (2025). Cloud-Based Big-Data Architectures for AI-Driven Financial Planning and Analysis and Early Risk Signaling in Resource-Constrained Community Financial Institutions. *The Metascience*, 3(3), 29-40.