Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Online Calibration of Account-Specific Risk Sensitivity from Local Behavioural Distances in Streaming Banking Transactions

Musa Olowo Alabi

amolowo@auchipoly.edu.ng

Auchi Polytechnic

Morufat Damilola Gbolagade

Al-Hikmah University

Research Article

Keywords: behavioural neighbourhood, fraud prevention, nearest neighbour, adaptive sensitivity

Posted Date: November 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11664687/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Online Calibration of Account-Specific Risk Sensitivity from Local Behavioural Distances in Streaming Banking Transactions

*Musa Olowo Alabi and Morufat Damilola Gbolagade

Department of Computer Science

Al-Hikmah University

Ilorin

Nigeria

This manuscript is currently undergoing peer review at the ABUAD International Journal of Natural and Applied Sciences (AIJNAS).

Author Note

Corresponding author: Musa Olowo Alabi. Email: amolowo@auchipoly.edu.ng

Abstract

Risk evaluation in digital banking can be affected by differences in the behavioural regularity of individual accounts. A single sensitivity parameter may respond differently to local behavioural distances across accounts. This prompted us to examine an online calibration procedure in which an account-specific risk sensitivity parameter is estimated from transaction-level nearest-neighbour distances within a local behavioural reference pool. During our exploratory analysis of IEEE-CIS fraud detection benchmark, 608 accounts (7.15%) were identified as eligible under criteria of at least 20 transactions and 30 days of history. Transactions were then represented through four behavioural deviation dimensions: recency, temporal, amount, and structural. Local behavioural distances of those deviation dimensions were used both for risk evaluation and for incremental calibration. In our experiment, a running mean was maintained from legitimate transaction observations without assuming a parametric distribution. Calibrated sensitivity varied substantially across behavioural strata, with mean values ranging from 0.7113 for high-activity accounts to 1.2534 for low-activity accounts, compared with the fixed reference of one. We also observed, at the transaction level, that an account with sensitivity of 0.6819 produced an adaptive risk of 0.7973 compared with 0.7284 under the fixed reference, while sensitivity above one reduced the corresponding risk. Interestingly, these results show that online calibration of risk sensitivity alters continuous risk while leaving the underlying local distances unchanged. The procedure provides a constant-memory mechanism for adapting risk sensitivity to account-specific behavioural regularity in streaming banking environments.

Keywords: behavioural neighbourhood, fraud prevention, nearest neighbour, adaptive sensitivity

Online Calibration of Account-Specific Risk Sensitivity from Local Behavioural Distances in Streaming Banking Transactions

1 Introduction

Digital payment systems increasingly support real-time transaction processing. This creates the need for fraud-risk decisions that can be made within the transaction-processing window. This perfectly matches the instant payment policies and standards issued by regulatory authorities such as the Central Bank of Nigeria, (2022) and Committee on Payments and Market Infrastructures, (2026). Behavioural characteristics provide evidence for such decisions because legitimate transactions from an account may exhibit recurring patterns in timing, amount, recency, and relationships with previously observed entities (Khan et al., 2024; Salami et al., 2025; Yampolskiy & Govindaraju, 2008).

By design, instance-based methods evaluate a new observation by comparing it with stored observations (Aha et al., 1991; Gonzalez et al., 2003). In his work, Cover and Hart (1967) established the classical nearest-neighbour (NN) rule, in which a case is classified according to its proximity to known cases. Looking closer at past work, we see that subsequent modifications have extended the behaviour of k-nearest-neighbour (k-NN) methods under different data conditions (Halder et al., 2024). The neighbourhood principle of k-NN provides a basis for local comparison. In our study, that principle is further restricted to the historical behavioural references of the account being evaluated rather than to a population-level reference space.

One subtle point worth noting is that the real the usefulness of a local distance depends not only on the magnitude of the distance but also on how it is converted into risk. If the same distance is observed for two accounts with different behavioural regularity, a fixed conversion may not produce the same risk response. Research on financial anomaly detection has similarly examined the calibration of decision thresholds in response to changes in transaction distributions (Huang et al., 2025). To illustrate, at one extreme, an account with a highly homogeneous transaction history exhibits strong behavioural regularity, such that fraudulent transactions that closely mimic legitimate activity may produce limited discriminating deviation. At the other extreme, an account with a highly variable transaction history generates larger natural deviations, making anomalous transactions more difficult to distinguish from normal behavioural variation. Between these extremes is the overlapping region: accounts occupy intermediate positions along the same continuum of behavioural regularity. A fixed sensitivity applied uniformly across this range may therefore be inappropriate for some accounts: excessive sensitivity for accounts near the variable end may increase false positives, whereas insufficient sensitivity for accounts near the homogeneous end may reduce the ability to distinguish small but meaningful deviations. These differences motivate the need to adapt sensitivity to the behavioural characteristics of the individual account.

Existing fraud-detection research has examined improvements in behavioural representations, anomaly scores, classification models, and decision thresholds. Recent work has also considered dynamic calibration of anomaly-detection thresholds when transaction distributions change (Huang et al., 2025). However, adapting a decision threshold is distinct from adapting the sensitivity through which a local behavioural distance is converted into risk. In particular, the account-level maintenance of a sensitivity parameter derived from locally retrieved behavioural distances

requires further examination. This distinction motivates our present study. We isolate the account-specific sensitivity component of a larger localised ego-graph fraud prevention model.

Accordingly, this study aims to develop and empirically evaluate an online procedure for maintaining account-specific risk sensitivity from local behavioural distances. Specifically, it examines whether account sensitivity varies across behavioural strata, whether it can be maintained incrementally from transaction-level observations, and how adaptive sensitivity changes transaction-level risk relative to a fixed $\gamma = 1$ reference.

Our study focuses specifically on one element within a broader, localised ego-graph fraud prevention model: account-specific sensitivity parameter (γ_a). This parameter is used to convert local behavioural distance into transaction risk. The wider model evaluates transaction risk before authorisation and maps the resulting risk score in [0,1] to three policy actions: Allow, Step-Up Authentication, or Block. The calibration uses a local Behavioural Reference Pool (BRP) constructed from the account's previously observed legitimate behavioural references. The model first match incoming transaction against a relationship-local candidate set. It then compares the candidates with the transaction through four behavioural deviation dimensions: recency, temporal, amount, and structural deviation. These deviations produce local distance observations for maintaining account-specific sensitivity.

Rather than assuming a particular probability distribution for local distance, we tracked the empirical mean incrementally using Welford's algorithm (Welford, 1962). This allows us to update the account-specific sensitivity parameter without storing the complete sequence of the observations. This provides us with a constant-memory update mechanism for maintaining the parameter at the account level.

Three research questions guide our study:

1. Do eligible accounts exhibit heterogeneous account-specific risk sensitivity across behavioural strata?
2. Can account-specific sensitivity be maintained incrementally from transaction-level local behavioural distance observations?
3. How does adaptive sensitivity alter transaction-level risk relative to a fixed $\gamma = 1$ reference when the same underlying local distances are used?

Our study offers two core contributions. The work first introduces an online calibration procedure designed to maintain an account-specific risk sensitivity parameter directly from local behavioural distances. Following this, we empirically track how this calibrated parameter fluctuates across distinct behavioural strata -specifically High Activity, Long History, Random Baseline, Short History, and Low Activity- and assess how these variations impact continuous risk when underlying distances are held constant. Consequently, our focus is mechanistic and empirical, rather than algorithmic; our study bypasses proposing new estimators or distance functions to highlight how blending existing components yields dynamic, account-specific risk modulation in streaming environments.

Because this study focuses strictly on evaluating a continuous risk score, classification performance - which depends entirely on downstream threshold policies - remains outside our scope. Similarly, the fixed ($\gamma = 1$) condition serves purely as a mathematical reference to isolate adaptive sensitivity effects, not as a competing baseline. To ground these mechanics, the evaluation simulates a streaming environment chronologically using the IEEE-CIS Fraud Detection benchmark.

The remainder of this paper is structured as follows. Section 2 presents the research design, online calibration workflow, account qualification criteria, behavioural deviation and distance definitions, account-specific update procedure, and adaptive and reference conditions. Section 3 reports and

discusses the results. Section 4 concludes. Section 5 lists abbreviations, and Section 6 presents the declarations.

2.2 Online Calibration Workflow

Triggered by the arrival of a transaction (q) associated with account (a), the calibration process first maps the transaction to its local behavioural context: Distance D, Location L, and Recipient Z. The system then pulls relationship-local candidate observations from the account's behavioural reference profile, directly comparing them against this incoming transaction's behavioural context.

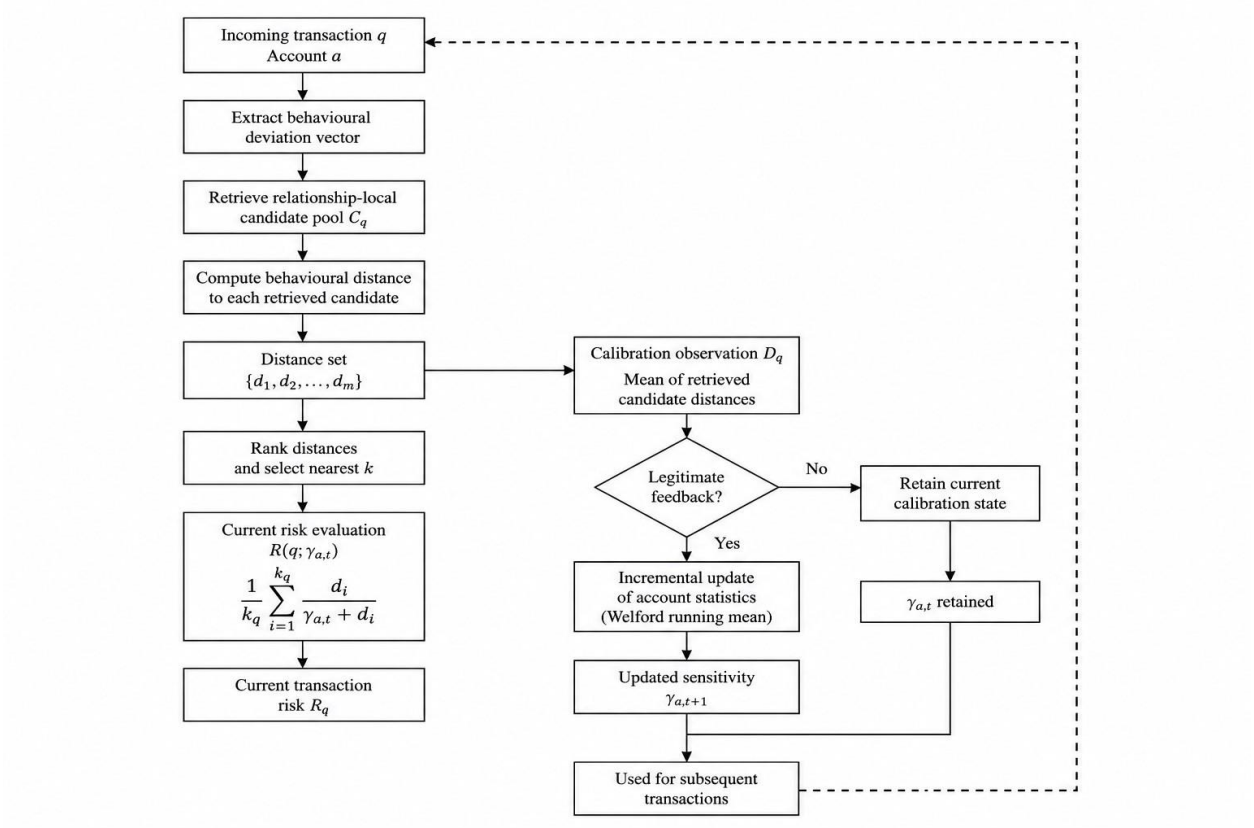
The workflow can be represented as:

$$q \rightarrow \mathrm{C}_q \rightarrow \{d_{q1}, d_{q2}, \ldots, d_{qm}\} \rightarrow R_q \rightarrow D_q \rightarrow \gamma_{a,t+1}. \quad (1)$$

Here, C_q denotes the retrieved candidate set, d_{qi} denotes the behavioural distance between the current transaction and candidate observation i, R_q denotes the current transaction risk, and D_q denotes the local-distance observation used for sensitivity calibration.

Figure 1

Online Calibration and Risk Evaluation Workflow



Note. The incoming transaction q is transformed into a behavioural deviation vector $\mathbf{d}_q$ and evaluated against a local candidate pool C_q. Nearest-neighbour distances drive immediate risk scoring (R_q) alongside calibration vector

computation (D_q). Confirmed legitimate feedback triggers an incremental Welford update to adjust account sensitivity $\gamma_{a,t+1}$ for subsequent risk evaluations.

Figure 1 shows the separation between current risk evaluation and subsequent calibration. The current transaction is scored using the sensitivity available at time t. Where the transaction is subsequently established as legitimate, its observed local-distance information is incorporated into the account's running sensitivity estimate. The updated sensitivity is therefore available for subsequent transactions rather.

2.3 Account Qualification

Because it is expected that an account possesses the behavioural context to be derived from identity dataset, our work does not consider all account as suitable for behavioural calibration. The account population was first restricted to transactions present in both IEEE-CIS transaction.csv and identity.csv. We merged the two files on TransactionID using inner join, so accounts without a matching identity record were excluded before qualification. Accounts were then screened according to minimum transaction-history requirements. An account was considered experimentally eligible when it contained at least 20 transactions and at least 30 days of observed transaction history. Of the 8,499 accounts in the merged dataset, 608 (7.15%) satisfied the qualification criteria. The stratified evaluation sample was drawn from these 608 accounts. Some strata contain fewer than 50 accounts because the qualified population did not supply enough accounts in those categories.The qualification rule was expressed as:

$$\text{Eligible}(a) = (|H_a| \geq N_{\min}) \wedge (T_a \geq T_{\min}), \tag{2}$$

where H_a represents the transaction history of account a, $N_{\min} = 20$ is the minimum number of transactions, and $T_{\min} = 30$ days is the minimum observation period.

Table 1 summarises the resulting account groups.

Table 1

Account Qualification and Eligibility Characteristics

Group	Accounts	Percentage	Mean Transactions	Mean History (days)
Experiment Eligible	608	7.15	187.57	169.57
Not Eligible	7,891	92.85	3.83	47.44
Total	8,499	100.00	16.98	56.17

The qualification step is important because an account with very little historical information cannot provide a sufficiently meaningful empirical reference for online behavioural sensitivity calibration.

2.4 Local Behavioural Reference Pool

For each eligible account (a), the local behavioural reference information represents previously observed transaction contexts and their associated behavioural characteristics. The candidate retrieval procedure is relationship-local rather than a global search over all transactions.

To select candidates, our model strictly follows a well-defined hierarchical relationship structure within the Behavioural Relationship Index (BRI). That is, the model must first check for exact matches across location, device, and recipient relationships. If this exact triad fails, it then considers any existing pairs. Lastly, it falls back on single matches or blank fields. This behavioural hierarchy is imposed to ensure that only contexts relevant to the current query context make the final list. The candidate retrieval process can therefore be conceptualised as an ordered cascading search function $f(L_q, D_q, Z_q)$, mapping from the most restrictive triad to relaxed attributes:

$$C_q = f(L_q, D_q, Z_q), \tag{3}$$

where L_q, D_q, and Z_q denote the relevant location, device, and recipient relationships of transaction q.

The number of retrieved candidates represents candidate-pool quantity:

$$M_q = |C_q|. \tag{4}$$

Candidate quality is subsequently determined through behavioural distance. Relationship locality determines which observations are eligible for comparison, while behavioural distance determines their relative closeness to the current transaction.

2.5 Behavioural Deviation and Local Distance

Each transaction is represented through four behavioural deviation dimensions:

$$\mathbf{d}_q = [d_{\text{rec}}, d_{\text{temp}}, d_{\text{amt}}, d_{\text{struct}}]. \tag{5}$$

The dimensions represent deviations in recency, temporal behaviour, transaction amount, and structural relationship behaviour, respectively.

For a current transaction q and stored behavioural observation j, the local behavioural distance is calculated as a weighted Euclidean distance:

$$D(q,j) = \sqrt{\sum_{r=1}^{4} w_r \left(d_{qr} - d_{jr}\right)^2} \tag{6}$$

where w_r represents the weight assigned to behavioural dimension r. In the present experiment, the four dimensions were assigned equal weights:

$$w_{\text{rec}} = w_{\text{temp}} = w_{\text{amt}} = w_{\text{struct}} = 0.25. \tag{7}$$

The resulting distance provides a local behavioural measure rather than a direct probability of fraud.

2.6 Nearest-Neighbour Risk Evaluation

For a transaction with retrieved distances $D_q = \{d_{q1}, d_{q2}, \ldots, d_{qm}\}$, the nearest k distances are selected for transaction-level risk evaluation. Let $k_q = \min(k,m)$ denote the number of available neighbours used for scoring.

The risk function is defined as:

$$R_q(\gamma_a) = \frac{1}{k_q} \sum_{i=1}^{k_q} \frac{d_{qt}}{\gamma_a + d_{qi}}, \tag{8}$$

The transformation maps the individual local distances to bounded risk contributions $R_q \in [0,1]$. Larger behavioural distances increase the resulting risk, while the sensitivity parameter γ_a controls how strongly the distance contributes to the resulting risk.

2.7 Account-Specific Online Risk Sensitivity

The account-specific sensitivity γ_a is maintained as a running empirical estimate of local behavioural distance. For a calibration observation D_q, the account-level mean is updated incrementally.

Let $\mu_{a,t}$ denote the account-specific running mean after t calibration observations.

The update is:

$$\mu_{a,t+1} + \frac{Dq + \mu_{a,t}}{n_{a,t}+1}, \tag{9}$$

where $n_{a,t}$ is the number of calibration observations accumulated for account a at time t.

The sensitivity γ_a is then obtained from the account-level empirical mean, subject to the minimum sensitivity floor γ_{floor} used by the implementation:

$$\gamma_{a,t} = \max(\mu_{a,t}, \gamma_{\text{floor}}). \tag{10}$$

To avoid computationally intensive re-computation of an account's sensitivity across its entire historical sequence, we implement Welford's method (Welford, 1962). This allows the running mean and corrected sum of squares to be updated incrementally in a single pass.

2.8 Calibration Observation

In the implementation used in this study, the mean local distance used to update sensitivity is calculated from the candidate-distance set:

$$D_q = \frac{1}{m}\sum_{j=1}^{m} d_{qj}, \quad (11)$$

where $m = |C_q|$ is the number of retrieved candidate observations with valid behavioural distances.

This calibration observation remains distinct from the nearest-neighbour subset defined in Equation 8. While the calibration observation evaluates the full candidate set to estimate the account's typical local distance, the risk evaluation isolates the *k*-nearest neighbours to evaluate the transaction's specific risk profile. Decoupling these subsets allows the model to address two distinct analytical objectives. Ultimately, the resulting observation is incorporated into the account's running sensitivity state only if the transaction is verified as legitimate; otherwise, it is discarded.

2.9 Adaptive and Fixed Reference Conditions

To examine the effect of sensitivity calibration, the adaptive risk score is compared with a fixed mathematical reference condition. The adaptive condition is:

$$R_q^A = \frac{1}{k_q}\sum_{i=1}^{k_q} \frac{d_{qt}}{\gamma_a + d_{qi}}, \quad (12)$$

$$R_q^F = \frac{1}{k_q}\sum_{i=1}^{k_q} \frac{d_{qt}}{1 + d_{qi}}, \quad (13)$$

$$\Delta R_q = R_q^A - R_q^F. \quad (14)$$

Where :

R_q^F is the fixed risk to computed using $\gamma = 1$

The fixed $\gamma = 1$ condition is not treated as a competing fraud model. It is a mathematical reference that permits the effect of adaptive sensitivity to be isolated while keeping the underlying local behavioural distances unchanged.

For an individual positive distance *d*, the corresponding difference is:

$$\frac{d}{\gamma + d} - \frac{d}{1 + d} = \frac{d(1-\gamma)}{(\gamma + d)(1 + d)}. \quad (15)$$

Consequently, $\gamma < 1 \Rightarrow \Delta R > 0$, while $\gamma > 1 \Rightarrow \Delta R < 0$. When $\gamma = 1$, the adaptive and fixed reference conditions are identical.

2.10 Online Calibration Algorithm

Algorithm 1 summarises the operational procedure used in the experiment.

Algorithm 1 Online Calibration of Account-Specific Risk Sensitivity

---

INPUT: Current transaction q for account a; local behavioural reference pool $B_{a,t}$; current sensitivity $\gamma_{a,t}$; neighbour count k; sensitivity floor γ_{floor}
OUTPUT: Current risk R_q and updated calibration state
Construct the behavioural deviation vector for q
Retrieve relationship-local candidate observations from $B_{a,t}$
Compute the behavioural distance between q and each valid candidate
Sort the valid distances and retain the nearest $k_q = \min(k,m)$ observations for risk evaluation
Compute current risk using Equation 8 and $\gamma_{a,t}$
Compute the calibration observation D_q from the retrieved valid candidate distances
Obtain the transaction validation outcome
if q is validated as legitimate and D_q is valid **then**
    Update the account running mean using D_q
    Set $\gamma_{a,t+1} \leftarrow \max(\mu_{a,t+1}, \gamma_{\text{floor}})$
else
    Retain the current calibration state
end if
return R_q and the resulting calibration state

During the experimental test sequence, benchmark ground-truth legitimacy was used as the feedback signal for determining whether a transaction could contribute to subsequent sensitivity calibration. The test procedure therefore represents an offline simulation of sequential online adaptation rather than a live banking deployment.

3 Results and Discussion

3.1 Heterogeneity in Account-Specific Sensitivity

The calibrated sensitivity values differed substantially across behavioural strata.

Table 2 summarises the account-level sensitivity estimates used in the experiment.

Table 2

Account-Specific Sensitivity by Behavioural Stratum

Behavioural Stratum	*n*	Mean γ	SD	Mean Shift
High Activity	50	0.7113	0.0837	-0.0089
Long History	15	0.8953	0.0972	0.0015
Random Baseline	50	1.0040	0.1540	0.0001
Short History	30	1.1295	0.1261	0.0079
Low Activity	27	1.2534	0.1746	0.0070

Note. The 172 accounts shown represent the stratified evaluation sample drawn from the 608 eligible accounts.

Sensitivity was not concentrated around the fixed reference value of one. The high-activity group had the lowest mean sensitivity ($M = 0.7113$, $SD = 0.0837$), whereas the low-activity group had the highest ($M = 1.2534$, $SD = 0.1746$). The random baseline was close to the reference condition ($M = 1.0040$, $SD = 0.1540$).

These observed variations directly validate our first research question, demonstrating that eligible accounts exhibit highly heterogeneous, account-specific sensitivity. This cross-account heterogeneity underscores that a uniform, static sensitivity threshold cannot adequately model empirical local-distance behaviours. Such a finding aligns with the central premise of our broader framework: behavioural identity is fundamentally account-dependent. Consequently, risk evaluation architectures must reflect an account's idiosyncratic regularity rather than relying on global, population-level norms. This extension bridges traditional behavioural profiling with transaction monitoring; while Yampolskiy and Govindaraju (2008) established recurring behavioural characteristics as robust modalities for user discrimination, our empirical results confirm that this foundational principle extends directly to modulating sensitivity metrics against localised deviations.

3.2 Transaction-Level Effect of Adaptive Sensitivity

The second aspect of the analysis examined how account-specific sensitivity affected risk for individual transactions while retaining the same local behavioural distances.

Table 3

Representative Transaction-Level Effect of Adaptive Sensitivity

Account	Transaction	Stratum	D_q	γ_a	R_A	R_F
5365	3496675	High Activity	2.5371	0.7668	0.7679	0.7173
9043	3500073	Long History	2.0128	0.7773	0.6327	0.5775
12866	3510120	Long History	2.2358	0.8619	0.5400	0.5030
15497	3502961	High Activity	2.6825	0.6819	0.7973	0.7284
14641	3538818	Low Activity	2.2584	1.2316	0.6435	0.6897
15093	3449755	Short History	0.8522	1.2031	0.4145	0.4600
18112	3483473	Short History	0.9750	1.0154	0.3637	0.3672

Note. R^A denotes adaptive risk and R^F denotes the fixed $\gamma = 1$ reference risk.

For the first four representative observations, $\gamma_a < 1$, and the adaptive risk was greater than the fixed reference risk. For example, account 15497 had $\gamma_a = 0.6819$ and produced an adaptive risk of 0.7973 compared with 0.7284 under the fixed reference.

Accounts 14641 and 15093 had sensitivity values greater than one. Their adaptive risks were lower than the corresponding fixed-reference values. Account 14641, for example, had $\gamma_a = 1.2316$, producing an adaptive risk of 0.6435 compared with 0.6897 under $\gamma = 1$.

These observations align precisely with the mathematical relationships established in Equation 15. Specifically, lower sensitivity amplifies the risk response to a given positive distance, whereas higher sensitivity attenuates that response. Crucially, this variance does not stem from a shift in the underlying behavioural distance measurement itself. Rather, it represents a change in the algorithmic sensitivity with which that distance is interpreted.

3.3 Risk Difference Across the Sensitivity Range

The relationship between adaptive sensitivity and risk difference was examined using

$\Delta R_q = R_q^A - R_q^F$. The direction of the risk adjustment changes at the reference value $\gamma = 1$.

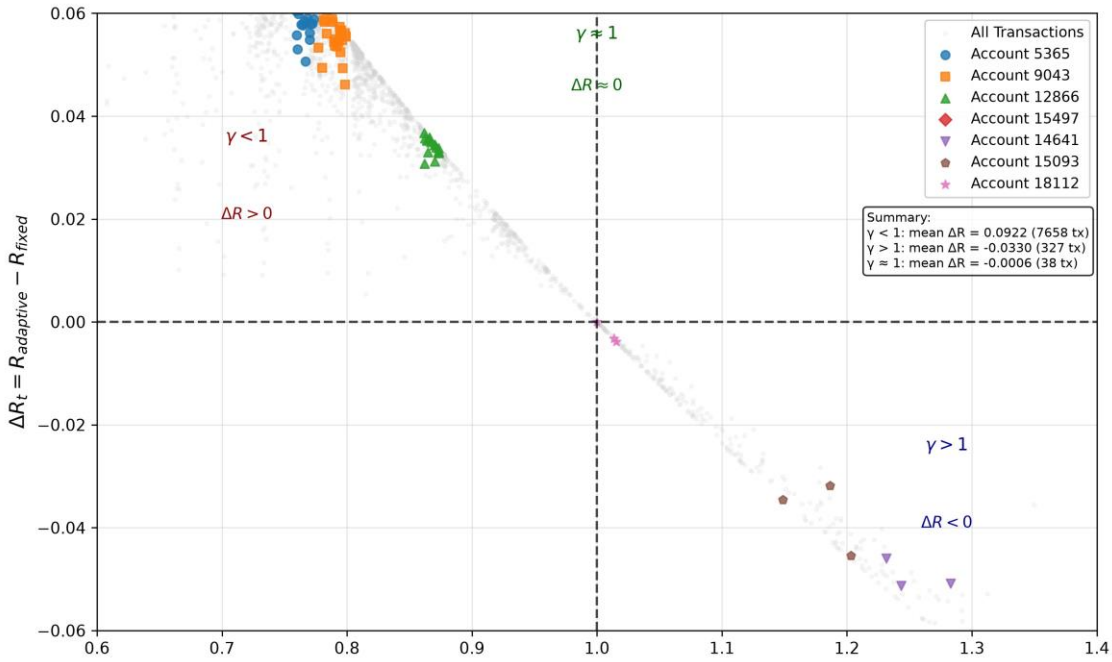


Figure 2

Transaction-Level Risk Difference Between Adaptive and Fixed Sensitivity

Figure 2 illustrates the expected transition around $\gamma = 1$. Sensitivity values below one produce positive risk differences, while sensitivity values above one produce negative risk differences. The relationship represents continuous risk modulation rather than a discrete change in classification.

The distribution is highly imbalanced: 7,958 transactions were associated with $\gamma_a < 1$, while only 38 were associated with $\gamma_a > 1$. The direction of the risk adjustment is consistent in both directions, but the magnitude of the negative adjustment is based on a small number of observations. This imbalance reflects the dataset's account composition rather than the mechanism itself. The contribution of the calibration procedure is not to introduce another binary fraud decision; rather it changes the degree to which a given local behavioural distance contributes to continuous risk.

3.4 Sequential Online Evolution

The calibration procedure was also examined sequentially. For a selected account, the local-distance observation, sensitivity, and resulting risk were tracked over successive observations.

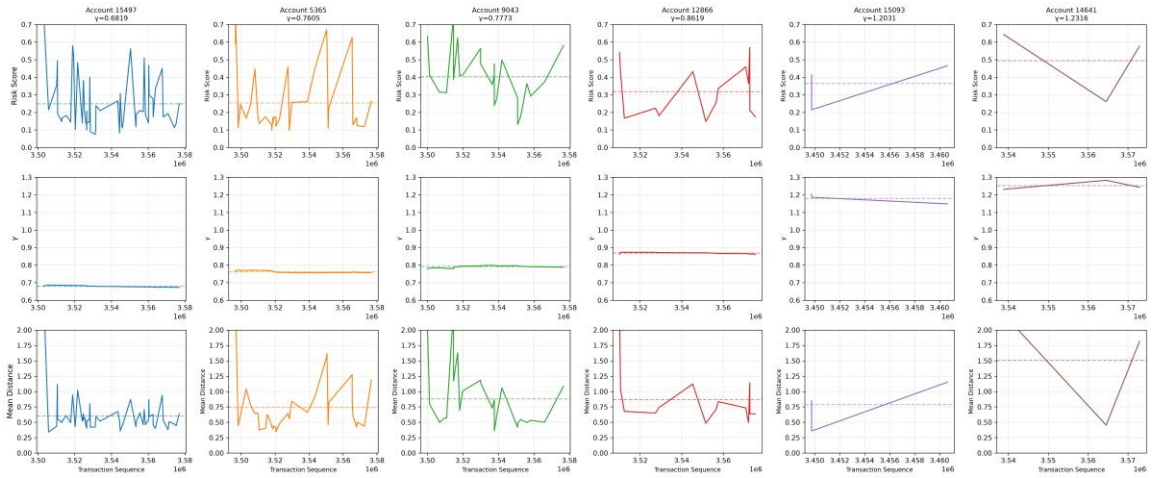


Figure 3

Sequential Evolution of Mean Local Distance, Risk Sensitivity and Transaction Risk

Figure 3 illustrates the intended adaptive behaviour. As legitimate observations accumulate, the account's running sensitivity updates automatically rather than recalculated from scratch. The resulting sensitivity is subsequently used for later transactions. The sequential behaviour addresses the second research question by demonstrating that account-specific sensitivity can be maintained incrementally. The procedure does not require a complete recalculation of historical observations after every transaction. Welford's method permits incremental updating of the running statistics, so the calibration state can be represented using a small number of account-level statistics rather than requiring repeated processing of the entire historical sequence.

3.5 Discussion of Findings

The first expectation was that account-specific sensitivity would vary across accounts exhibiting different behavioural characteristics. The results support this expectation within the evaluated behavioural strata. Mean γ_a ranged from 0.7113 in the High Activity group to 1.2534 in the Low Activity group, with intermediate values of 0.8953, 1.0040, and 1.1295 for the Long History, Random Baseline, and Short History groups, respectively. The ratio between the highest and lowest mean sensitivity was approximately 1.76. This variation indicates that the calibration procedure did not assign a common sensitivity to all eligible accounts. Under a fixed sensitivity, such account-level differences would be removed because the same distance-to-risk transformation would be applied regardless of the calibrated characteristics of the account.

The direction of the observed differences is also relevant to the risk transformation. Since ($R = \frac{d}{d+\gamma_a}$), smaller values of (γ_a) make the transformation more sensitive to a given behavioural distance, whereas larger values produce a less sensitive transformation. Thus, the lower mean observed for the High Activity group corresponds to greater sensitivity than the higher mean observed for the Low Activity group. However, these differences should not be interpreted as evidence that activity level itself determines behavioural sensitivity. The strata were used to examine variation in the calibrated parameter, while (γ_a) was derived from local behavioural-distance observations. The present results therefore establish variation in the calibrated sensitivity

across the evaluated strata; they do not identify activity level, transaction frequency, or history length as causal determinants of that variation.

The second expectation concerned whether account-specific sensitivity could be maintained incrementally from transaction-level local behavioural observations. The streamed evaluation shows that (γ_a) changed as observations were incorporated. The use of Welford's running statistics allowed the empirical mean of the local distance observations to be updated without retaining the complete sequence of previous observations. This provides a constant-memory mechanism for maintaining the account-level parameter as new behavioural evidence becomes available. The result demonstrates the operational suitability of the update procedure for the streaming setting considered, but it does not establish that Welford's method is superior to alternative online estimators.

The third expectation concerned whether adaptive sensitivity would alter the resulting transaction risk when the underlying behavioural evidence remained unchanged. This comparison provides the clearest test of the role of (γ_a). Both the adaptive and fixed ($\gamma = 1$) conditions use the same underlying local behavioural distances; therefore, the behavioural representation and distance observations are held constant. Differences in the resulting risk values consequently arise from the sensitivity parameter in the distance-to-risk transformation. For a positive distance, an account with (γ_a<1) produces a higher risk value than the fixed-(γ =1) condition, whereas an account with (γ_a>1) produces a lower value. Adaptive calibration therefore changes how the same local behavioural deviation is translated into continuous risk according to the calibrated characteristics of the account.

The relationship-local retrieval procedure provides an additional contextual layer to this calibration. Candidate selection does not treat every historical transaction as equally relevant. Instead, the historical reference space is first restricted according to the behavioural relationships associated with the incoming transaction, after which local behavioural distance is calculated. The approach is consistent with broader graph-based fraud and anomaly-detection research in which local connectivity and relationships provide information beyond isolated transaction attributes (Pu et al., 2020).

The findings should nevertheless be interpreted within the scope of the component evaluated. The observed variation in (γ_a) demonstrates heterogeneity in the calibrated sensitivity and shows that this parameter changes the conversion of local distance into continuous risk. It does not, by itself, demonstrate improved fraud-detection performance, reduced false positives, increased fraud recall, or financial loss prevention. Those outcomes depend on the subsequent decision thresholds, policy actions, fraud outcomes, and operational costs. Similarly, the present study does not establish that relationship-local retrieval is responsible for the observed differences in sensitivity because the retrieval mechanism was not isolated through an ablation experiment.

The fixed (γ =1) should be interpreted as a controlled mathematical reference rather than as a competing model or baseline classifier. Its value is that it holds the local behavioural distances constant while removing account-specific sensitivity. The comparison consequently isolates the effect of the calibration parameter within the specified risk function. The findings support the role of (γ_a) as an account-specific risk sensitivity parameter.

3.6 Limitations

Several limitations should be considered when interpreting the findings.

The experiment used the IEEE-CIS fraud detection benchmark rather than live banking transaction data. The benchmark provides transaction observations suitable for empirical evaluation, but it does not reproduce all operational characteristics of a production banking environment. For example, the identity dataset contains substantial missing values in the fields designated for device and location; these fields are used (alongside recipient field) to compute structural deviation. Within an account, several transactions may share the same imputed context combination. This will be registered as a familiar relationship by Behavioural Relationship Index (BRI). This artificial familiarity would reduce the structural deviation assigned to those transactions. Because structural deviation is one of the four components from which the behavioural distance is computed, the reduction propagates into the distance and, through the calibration observation D_q computed as the mean over all retrieved candidate distances, into the account-specific sensitivity estimate.

Second, the streaming component was evaluated through chronological simulation. The transaction sequence was processed in temporal order $(t_1 < t_2 < \cdots < t_n)$, based on the recorded timestamp in the IEEE-CIS transaction dataset, so that the current sensitivity $(\gamma_{a,t-1})$ was used before subsequent legitimate feedback updated the account state including the new sensitivity $(\gamma_{a,t})$. However, benchmark ground-truth legitimacy was available to the experimental procedure. In an operational banking system, the corresponding feedback signal would need to arise from an appropriate validation, authentication, investigation, or post-authorisation mechanism.

Third, the present study focuses on the sensitivity parameter; does not independently optimise all components of the wider localised behavioural graph system. Candidate construction, behavioural representation, and distance weighting may each affect the resulting local-distance distribution and therefore the calibrated sensitivity.

Finally, the calibration procedure uses empirical local-distance observations and does not establish that a particular distributional form governs account behaviour. Further work could examine alternative online estimators, robust statistics, adaptive forgetting mechanisms, and sensitivity updates under behavioural drift.

4 Conclusion

This study set out to evaluate an online calibration procedure for account-specific risk sensitivity (γ_a) using local behavioural-distance observations in transaction streams. Empirical results confirmed that sensitivity varies systematically across behavioural strata, with mean γ_a values moving from 0.7113 (High Activity) to 1.2534 (Low Activity). Crucially, this demonstrates that local distance sensitivity does not have to be uniform across all accounts. Because the risk equation ($R = \frac{d}{d+\gamma_a}$) relies directly on γ_a, setting sensitivity below one amplifies the assigned risk for any given positive distance compared to a fixed $\gamma = 1$, whereas values above one suppresses it. Because both the adaptive and fixed setups processed the exact same underlying distances, this divergence stems entirely from the calibrated sensitivity rather than a shifted behavioural representation. Furthermore, the findings proved that γ_a can be maintained incrementally. Applying Welford's

running update provides a constant-memory mechanism that tracks the account-level empirical mean without recomputing historical sequences which is a vital requirement for real-time streaming environments. These results move behavioural fraud prevention beyond simple nearest-neighbour matching. Instead of merely checking which historical transaction is most similar, the framework adapts the interpretation of a local deviation to match the baseline profile of the individual account. This bridges the gap between rigid global thresholds and dynamic, account-tailored risk scoring while preserving nearest-neighbour locality.

5 List of Abbreviations

BRP	Behavioural Reference Pool
BRI	Behavioural Relationship Index
BSP	Behavioural Sensitivity Profile
IEEE-CIS	Institute of Electrical and Electronics Engineers–Computational Intelligence Society
k-NN	k-Nearest Neighbour
LEG-kNN	Localised Ego-Graph k-Nearest Neighbour

6 Declarations

6.1 Ethics Approval and Consent to Participate

This study used the publicly available IEEE-CIS fraud detection benchmark dataset. No human participants were directly recruited for the study, and no new personally identifiable information was collected from participants.

6.2 Consent for Publication

Not applicable. The study did not involve the collection of identifiable individual participant information.

6.3 Availability of Data and Materials

The datasets analysed during the current study are available in the IEEE-CIS Fraud Detection repository, https://www.kaggle.com/c/ieee-fraud-detection. The processed data and code used to generate the results are available from the corresponding author on reasonable request.

6.4 Competing Interests

The authors declare that there is no competing interests.

6.5 Funding

No external funding was received for this study.

6.6 Authors' Contributions

Musa Olowo Alabi conceived the study, developed the computational procedure, conducted the experiments, analysed the results, and prepared the manuscript. Morufat Damilola Gbolagade supervised the research, contributed to the methodological development, and critically reviewed the manuscript. Both authors approved the final manuscript.

6.7 Acknowledgements

The authors acknowledge the Department of Computer Science, Al-Hikmah University, Ilorin, Nigeria, for providing the academic environment in which this research was conducted. Additionally, the corresponding author explicitly thanks the management of Auchi Polytechnic for providing the institutional flexibility and support required to balance academic research with workplace demands.

References

Aha, D. W., Kibler, D., & Albert, M. K. (1991). Instance-based learning algorithms. *Machine Learning*, *6*(1), 37–66. https://doi.org/10.1023/A:1022689900470

Bello, O. A., Folorunso, A., Ejiofor, O. E., Budale, F. Z., Adebayo, K., & Babatunde, O. A. (2023). Machine learning approaches for enhancing fraud prevention in financial transactions. *International Journal of Management Technology*, *10*(1), 85–109. https://doi.org/10.37745/ijmt.2013/vol10n185109

Central Bank of Nigeria. (2022, November). *Nigeria payments system vision 2025 (PSV 2025): Ensuring a safe, efficient, and robust payments system* (tech. rep.) (Released 24 November 2022). Central Bank of Nigeria, Banking Services Department. Abuja, Nigeria. https://www.cbn.gov.ng/PaymentsSystem/PSV2025.html

Committee on Payments and Market Infrastructures. (2026, April). *Fast payments and their implications for central banks* (CPMI Briefs No. 12). Bank for International Settlements. https://www.bis.org/statistics/payment_stats/commentary2604.htm

Cover, T. M., & Hart, P. E. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, *13*(1), 21–27. https://doi.org/10.1109/TIT.1967.1053964

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591–635. https://doi.org/10.1207/s15516709cog2704_2

Halder, R. K., Uddin, M. N., Uddin, M. A., Aryal, S., & Khraisat, A. (2024). Enhancing k-nearest neighbor algorithm: A comprehensive review and performance analysis of modifications. *Journal of Big Data*, *11*(1), 113. https://doi.org/10.1186/s40537-024-00973-y

Huang, A., Wang, Y., Zhou, P., Zhang, X., Tsai, S., & Chen, L. (2025). Dynamic calibration of decision thresholds for financial anomaly detection: Verification with payment platform information and data. *Journal of Global Information Management, 33*(1), 1–XX. https://doi.org/10.4018/JGIM.395852

IEEE Computational Intelligence Society. (2019). Ieee-cis fraud detection [Kaggle competition dataset]. https://www.kaggle.com/c/ieee-fraud-detection

Khan, S., Devlen, C., Manno, M., & Hou, D. (2024). Mouse dynamics behavioral biometrics: A survey. *ACM Computing Surveys*, *56*(6), 1–33. https://doi.org/10.1145/3640311

Pu, J., Zhang, J., Shao, H., Zhang, T., & Rao, Y. (2020). Egodetect: Visual detection and exploration of anomaly in social communication network. *Sensors*, *20*(20), 5895. https://doi.org/10.3390/s20205895

Salami, I. A., Popoola, A. D., Gbadebo, M. O., Kolo, F. H. O., & Adesokan-Imran, T. O. (2025). Ai-powered behavioural biometrics for fraud detection in digital banking: A next-generation approach to financial cybersecurity. *Asian Journal of Research in Computer Science*, *18*(4), 473–494. https://doi.org/10.9734/ajrcos/2025/v18i4632

Welford, B. P. (1962). Note on a method for calculating corrected sums of squares and products. *Technometrics*, *4*(3), 419–420. https://doi.org/10.1080/00401706.1962.10490022

Yampolskiy, R. V., & Govindaraju, V. (2008). Behavioural biometrics: A survey and classification. *International Journal of Biometrics*, *1*(1), 81–113. https://doi.org/10.1504/IJBM.2008.018665

Zhang, Z., Wan, J., Zhou, M., Lai, Z., Tessone, C. J., Chen, G., & Liao, H. (2023). Temporal burstiness and collaborative camouflage aware fraud detection. *Information Processing & Management*, *60*(2), 103170. https://doi.org/10.1016/j.ipm.2022.103170