Supplementary materials for "FraudBench-NG: A NIBSS-Calibrated Synthetic Dataset and Generator for Nigerian Inter-Bank Fraud Detection"

Otu Jacob, Colin Kuka, Adam Knowles, Cyrus Ebere Orji

This file contains Appendices A–F of the manuscript. Section, table and figure numbers without a letter prefix refer to the main paper.

Appendix A Register of design choices

Table A1 lists every parameter of the generator with its status and justification. *NIBSS* means taken from the published 2023 tables; *assumed* means a modelling judgement with no source; *design* means a choice of scale or protocol. An assumed value is not claimed to be realistic, only plausible, and every one can be changed in `generate_dataset.py`.

Table A1: Register of design choices

Choice	Value	Status	Justification and implication
Fraud channel shares	Mobile 49.8%, Web 22.9%, POS 18.4%, IB 5.6%, ECOM 2.6%, ATM 0.8%	NIBSS	Channel counts divided by their own sum (95,525).
Monthly intensity	$q_m/(\text{days}_m/365)$	NIBSS	Monthly case shares; per-day scaling so short months are not under-weighted.
Technique shares	65.8, 10.6, 7.1, 5.5, 4.7, 6.3%	NIBSS	Five largest techniques plus pooled "other".
Fraud quotas	Incidents, not rows	Design	NIBSS counts and losses are per case, so channel quotas count incidents.

Choice	Value	Status	Justification and implication
Mean incident loss	Table D6	NIBSS	Channel loss divided by channel cases; actual not attempted loss. One loss per incident, split over its rows.
Loss split α	Dirichlet(2)	Assumed	Moderately uneven shares of an incident's loss over its transactions.
Fraud dispersion s	0.9	Assumed	NIBSS gives only means; s spreads fraud over about a decade and overlaps legitimate amounts.
Prevalence	0.1%	Design	Enough fraud rows ($\geq$2,000) for supervised learning; not comparable with the NIBSS case-to-transaction ratio (about 0.00085%); same order as ULB (0.17%) and PaySim (0.13%).
Takeover continuation ϕ	0.40	Assumed	Memoryless continue-or-stop; mean 1.67 transactions per incident (1.63 realised). No NIBSS statistic exists; chosen for a moderate burst effect; sensitivity at 0, 0.4, 0.7.
Takeover delay	Exp, mean 300 s	Assumed	"Within minutes"; shorter than the 15-minute legitimate burst scale.
Customers	25,000	Design	Gives about 2.1 million rows; scale only.
Age shares	18, 30, 24, 15, 8, 5%	Assumed	Young digital-banking population; matches no published distribution. Inert.
Location shares	Lagos 28, Abuja 20, Oyo 14, Ogun 12, Rivers 11, other 15%	Assumed	Assumed urban concentration; Lagos is 6.2% of the population [1] against 28% here, and EFInA reports higher exclusion in rural areas [2]. Inert; Lagos share of NIBSS 2025 fraud (63%) is not reproduced.
Banks	10, uniform	Assumed	Anonymised; no market shares. Inert.
Preferred channel $\mathbf{p}_c$	ATM 5, ECOM 11, IB 11, Mobile 42, POS 16, Web 15%	Assumed	Mobile largest, as in NIBSS fraud; NIBSS publishes no legitimate channel mix.
Channel stickiness	0.70	Assumed	Most payments through one main channel while still giving channel diversity.
Spend level μ_i	$\mathcal{N}(10.8, 0.6^2)$	Assumed	Median about NGN 49,000 per transaction; heterogeneity between customers.
Spend dispersion σ_i	$U(0.4, 0.9)$	Assumed	90th/10th percentile ratio between 2.8 and 10 within a customer.
Activity n_i	LogNormal $(4.2, 0.7^2)$, clipped 5–1000	Assumed	Median 67 per year; right-skewed [3].

Choice	Value	Status	Justification and implication
Channel shifts δ_c	ATM −0.30, ECOM −0.10, POS −0.10, Mobile 0, Web +0.10, IB +0.50	Assumed	Ordinal: IB highest (NIBSS mean fraud loss is highest there), ATM lowest (cash limits). Sizes are guesses.
Hour weights	Table E7	Assumed	Generic: night trough, daytime plateau, evening peak. Same for all customers and for fraud.
Weekday weights	Table E7	Assumed	Friday peak, Sunday low. No month-end or holiday effects.
Burst parameters	$\rho = 0.25$, mean 900 s	Assumed	Chosen to give inter-arrival CV well above 1 [4].
Merchant categories	13, uniform within the channel-compatible set	Assumed	No source for merchant mix; Table E9 only removes impossible combinations.
Merchant fraud factors	0.3–1.4	Assumed	Ordinal judgement (Table E8).
Burn-in	90 days	Design	Reduces cold-start bias; residual drift remains (Section 6).
Windows	24 h, 7 d, exact seconds	Design	Standard aggregation windows [5].

Appendix B What the generator does not model

The generator samples age group and location from fixed probabilities, assigns a bank and a preferred channel, and samples merchant categories uniformly within a channel-compatibility table. Table B2 lists structure that real transaction data contain and that this generator does not reproduce. The first five rows are the most important for studies of customer behaviour.

Table B2: Structure not modelled, what the generator does instead, and the consequence

Not modelled	What the generator does	Consequence
Merchant preferences of customer groups	Merchant drawn uniformly from the categories allowed for the channel (Table E9), independent of age, location and customer	No customer has favourite merchants; age or location cannot explain what a customer buys.
Relationships among location, bank and channel	Age, location, bank and preferred channel drawn independently; channel does not depend on location, bank or age	No regional or bank-specific channel patterns; subgroup analyses by location or bank are meaningless.

Not modelled	What the generator does	Consequence
Relationship between amount and merchant category	Amount depends only on customer and channel	A grocery and an electronics payment have the same amount distribution; amount cannot reveal merchant type.
Regular payments	Every transaction is independent apart from bursts; no repeat payee or merchant, no salary, rent, subscription, standing order or loan repayment	Customers have no routine; "unusual payee" or "missed regular payment" features cannot be built.
Individuals versus businesses	All customers share one family of distributions; no business accounts, no merchants as accounts	No high-volume business behaviour or settlement flows; heterogeneity comes only from the spend and activity distributions.
Demographic and geographic fraud patterns	Fraud independent of age, location and bank	Reported concentration in Lagos (63%) and among customers aged 40+ is absent.
Time of day of fraud	Fraud inherits the timing of the legitimate transaction it replaces	Hour and weekday carry no fraud signal.
Seasonality and trend	Same weekly and daily profile all year; no month-end, salary-day, holiday or trend	No concept drift in legitimate behaviour.
Customer life cycle	All customers active throughout; none join or leave	No new-account fraud, no dormancy.
Devices, locations and beneficiaries	Not generated	No device, IP, geolocation, new-beneficiary or mule-account features.
Detection and labelling process	Every fraud row is labelled instantly and correctly	No label delay, noise, undetected fraud or feedback from a fraud team.
Fraud rings and repeat victims	Incidents are independent; victims are chosen at random in proportion to activity	No shared mule accounts or repeated targeting.

Appendix C Variable reference

Table C3 defines every column. *Customer-level* variables are fixed once per customer; *transaction-level* variables are drawn per transaction or computed from the customer's causal history $\mathcal{H}_{i,k}$ (Section 3.7).

Table C3: Complete variable reference

Variable	Level	Definition
`customer_id`	Customer	Synthetic customer identifier.

age_group	Customer	Age bracket, fixed per customer (Section 3.2).
location	Customer	The customer's registered location. No per-transaction physical location is modelled.
bank	Customer	The customer's single primary originating bank.
transaction_id	Transaction	Unique identifier, ordered by time.
timestamp	Transaction	Transaction time (Section 3.3).
channel	Transaction	Payment channel of this transaction (Eq. 5).
merchant_category	Transaction	Merchant category, drawn per transaction.
amount	Transaction	Amount in NGN (Eqs. 6–7).
is_fraud	Transaction	Binary label (Section 3.6).
fraud_technique	Transaction	Technique, defined only when is_fraud=1.
Features computed from the prior history $\mathcal{H}_{i,k}$ (strictly causal, current transaction excluded):		
amount_log	Transaction	$\ln(1 + \texttt{amount})$.
amount_vs_mean_ratio	Transaction	$R_{i,k}$, Eq. 14.
amount_sum_24h, tx_count_24h	Transaction	$\Sigma_{i,k}^{24h}$, $n_{i,k}^{24h}$ (prior transactions only), Eq. 12.
amount_mean_7d, amount_std_7d	Transaction	Eq. 13.
amount_mean_total, amount_std_total	Transaction	Eqs. 10–11.
tx_count_total	Transaction	$k - 1$, the customer's number of prior transactions.
channel_diversity	Transaction	$D_{i,k}^{\text{ch}}$, Eq. 16.
online_channel_ratio	Transaction	$O_{i,k}$, Eq. 17.
velocity_score	Transaction	$V_{i,k}$, Eq. 15.
merchant_risk_score	Transaction	Eq. 18.
hour_sin/cos, day_sin/cos, month_sin/cos, is_peak_hour	Transaction	Section 3.3.

Appendix D NIBSS calibration data

This appendix collects the calibration tables referred to in Section 3.1.

Table D4: Aggregate NIBSS fraud figures by year. Only 2023 comes with the breakdown by channel, month and technique used by the generator.

Year	Cases	Actual loss (NGN bn)	Source and remark
2023	95,620	17.67	[6]; full tables
2024	70,111	52.26	[7]; one incident of NGN 31.1 bn
			[8]; no breakdown found
2025	67,518	25.85	[8]; no breakdown found

Table D5: Totals implied by the individual tables of the NIBSS 2023 report

Table (rows)	Cases	Actual loss (NGN m)
By channel (6)	95,525	16,636.1
By month (12)	95,620	17,774.6
By technique (14)	95,620	17,666.9
Stated total	95,620	17,666.8
Channel minus stated	−95 (−0.10%)	−1,030.7 (−5.8%)
Month minus stated	0	+107.8 (+0.6%)

Table D6: Mean fraud loss per case by channel, NIBSS 2023 (loss ÷ cases, NGN)

Channel	Mobile	Web	POS	IB	ECOM	ATM
$\bar{A}_c^{\text{NIBSS}}$	119,760	101,595	251,469	761,438	58,840	88,089

Appendix E Assumed activity weights and merchant factors

Hour and weekday weights are assumptions (Section 3.3); merchant factors are the authors' ordinal judgements (Section 3.8).

Table E7: Assumed activity weights

Day	Mon	Tue	Wed	Thu	Fri	Sat	Sun
w^{dow}	1.00	1.00	1.00	1.05	1.25	1.10	0.80
Hour	0–5	6	7	8	9–11	12–14	15
w^{hour}	0.3–0.5	0.8	1.2	1.6	2.0	2.2	2.0
Hour	16	17	18	19	20	21	22–23
w^{hour}	1.9	2.1	2.2	2.0	1.6	1.2	0.8, 0.5

Table E8: Merchant-category fraud factors (assumed), the published `merchant_risk_score`, and the realised fraud rows in the released dataset. "Share" is the category's share of all fraud rows and "Rate" is fraud rows per 1,000 transactions of that category.

Category	f_{cat}	Score	Share	Rate	Reasoning (judgement)
Transfer	1.4	0.900	8.9		
Electronics	1.3	0.827	9.8		
Bill payment	1.2	0.755	10.8		
Airtime	1.1	0.682	11.1		
Entertainment	1.1	0.682	10.5		
Fashion	1.0	0.609	7.8		
Transport	0.9	0.536	8.6		
Restaurant	0.8	0.464	7.5		
Retail	0.7	0.391	6.6		
Grocery	0.5	0.245	5.5		
ATM withdrawal	0.4	0.173	2.1		
Education	0.4	0.173	6.4		
Medical	0.3	0.100	4.5		

Table E9: Channel-to-merchant compatibility table (assumed). Within the allowed set, merchant category is drawn uniformly, for legitimate and follow-on rows; the first fraud row of an incident is selected with the weights of Eq. 8.

Channel	Allowed merchant categories
ATM	ATM Withdrawal only
POS	All 13 categories except Transfer (so ATM Withdrawal is 8.3% of POS fraud rows: cash-out at a terminal)
ECOM	All except ATM Withdrawal and Transfer
Mobile, Web	All except ATM Withdrawal
IB (internet banking)	Transfer, Bill Payment, Airtime, Education, Medical, Entertainment

Appendix F Additional validation

Tests and figures that support Section 4. Table F10 lists the internal-validity tests; Table F12 reports the merchant sensitivity analysis (mean ± s.d. over ten seeds of 8,000 customers; the PR-AUC values are lower than in Table 5 because these runs use a smaller dataset and fewer trees).

Table F10: Internal validity tests on the released dataset

Test	Method	Result
Causality of features	Recompute $R_{i,k}$, $\Sigma_{i,k}^{24h}$, $n_{i,k}^{24h}$, $\bar{A}_{i,k}^{7d}$ row by row from the strictly prior transactions (including burn-in) of 40 random customers	Max relative error 9.3×10^{-6} (rounding)
Velocity is not a copy of amount	Spearman correlation of $V_{i,k}$ with $R_{i,k}$	0.001
Sparse windows	Share of rows with no prior transaction in 24 h	57.8%
No "everyone is new" artefact	Share of rows with no prior transaction at all	0.002%
Prevalence	Count	2,059 / 2,102,151 = 0.098% (1,261 incidents)
Missing values, duplicate IDs	Count	0, 0
Label permutation	ROC-AUC of $R_{i,k}$ against shuffled labels	0.499 (chance)
Stability over seeds	Ten seeds, 8,000 customers each	See below
Channel quota	χ^2 of incident counts against NIBSS shares	$p \approx 1$
Incident loss	Mean loss per incident against NIBSS loss per case	Within 8% for four large channels (Section 4.1)

Table F11: Stability over ten seeds (8,000 customers each): rows, fraud rows, prevalence (%), maximum deviation of the incident channel mix from NIBSS (percentage points), inter-arrival CV, follow-on share of fraud rows, and single-feature ROC-AUC of the velocity score and of the amount.

Seed	Rows	Fraud	Prev.	Dev.	CV	Foll.	AUC vel.	AUC amt.
100	661,128	655	0.099	0.12	1.76	0.39	0.619	0.459
101	673,265	673	0.100	0.11	1.78	0.40	0.650	0.500
102	685,367	680	0.099	0.12	1.78	0.40	0.640	0.456
103	674,864	680	0.101	0.11	1.78	0.41	0.639	0.454
104	670,711	681	0.102	0.09	1.77	0.41	0.656	0.460
105	664,253	619	0.093	0.12	1.77	0.36	0.608	0.508
106	671,801	670	0.100	0.12	1.79	0.40	0.645	0.494
107	673,870	714	0.106	0.11	1.78	0.43	0.631	0.450
108	678,073	650	0.096	0.10	1.79	0.38	0.628	0.482
109	682,675	638	0.093	0.13	1.78	0.36	0.610	0.506
Mean	–	–	0.099	0.11	1.78	0.39	0.633	0.477
s.d.	–	–	0.004	0.01	0.01	0.02	0.017	0.023

Table F12: Sensitivity to the merchant factors (mean ± s.d. over ten seeds of 8,000 customers). The expected share of fraud rows in the three riskiest categories under no effect is 3/13 = 23.1%. RF PR-AUC is for the default features (no merchant columns) and with merchant columns added; "pooled" is the PR-AUC over the test rows of all ten seeds together.

Factors	Score ROC-AUC	Fraud in top-3	RF PR-AUC	RF PR-AUC + merchant	Pooled PR-AUC	RF recall @1% FPR
Assumed	0.578±0.011	30.2±1.3%	0.023±0.017	0.022±0.013	0.014	0.208±0.079
Flat	0.500±0.000	23.4±2.2%	0.032±0.026	0.030±0.024	0.023	0.202±0.059
Reversed	0.555±0.009	15.8±1.4%	0.015±0.012	0.015±0.011	0.010	0.205±0.044

Table F13: Sensitivity to fraud-process assumptions (ten seeds of 8,000 customers each). Panel A: mean follow-on share of fraud rows and single-feature ROC-AUC of the velocity score and of the amount (values below 0.5 mean that fraud rows have *smaller* amounts, because incident losses are split). Panel B: random-forest metrics with default features (mean ± s.d.; "pooled" is over all ten test sets); PR-AUC for first events and follow-ons is computed against all legitimate test rows.

Panel A

Setting	Follow-on share	AUC velocity	AUC amount
Assumed ($\phi = 0.4$, $s = 0.9$)	0.40	0.64	0.46
No follow-on ($\phi = 0$)	0.00	0.50	0.67
More follow-on ($\phi = 0.7$)	0.69	0.76	0.28
Loss dispersion $s = 0.5$	0.40	0.65	0.54
Loss dispersion $s = 1.3$	0.40	0.63	0.37

Panel B

Setting	PR-AUC	Pooled PR-AUC	Recall @1% FPR	PR-AUC first	PR-AUC follow-on
Assumed	0.023±0.017	0.014	0.208±0.079	0.013±0.015	0.026±0.030
$\phi = 0$	0.041±0.019	0.031	0.241±0.047	0.041±0.019	–
$\phi = 0.7$	0.141±0.082	0.096	0.419±0.096	0.005±0.004	0.190±0.106
$s = 0.5$	0.020±0.013	0.013	0.138±0.060	0.020±0.016	0.009±0.012
$s = 1.3$	0.068±0.036	0.060	0.328±0.045	0.019±0.009	0.104±0.073

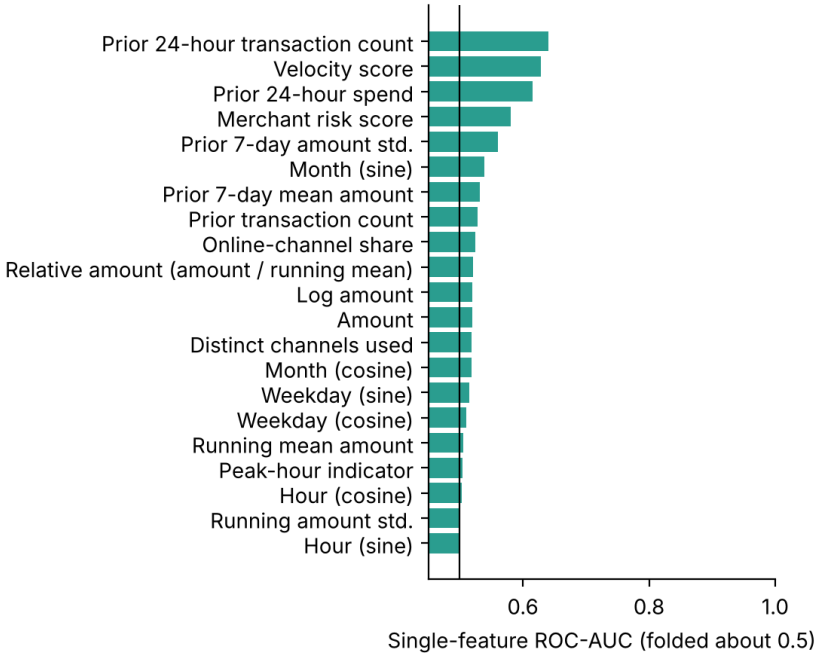


Fig. F1 Single-feature ROC-AUC (folded about 0.5) for all numeric features.

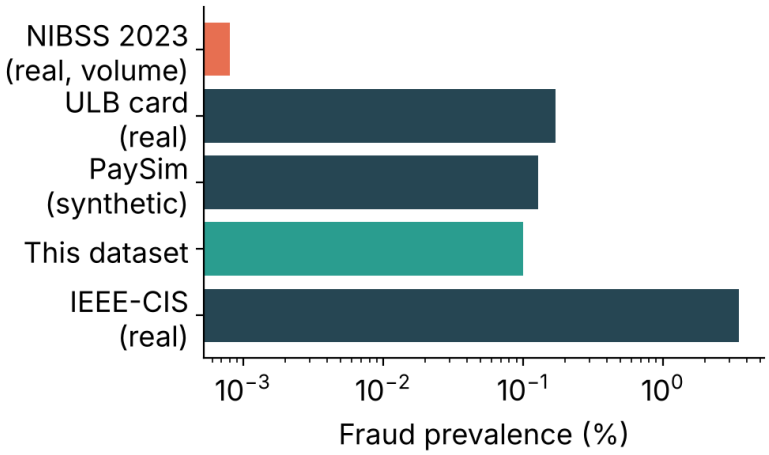


Fig. F2 Fraud prevalence of this dataset and of reference datasets (log scale). Values for ULB, PaySim and IEEE-CIS are taken from their documentation.

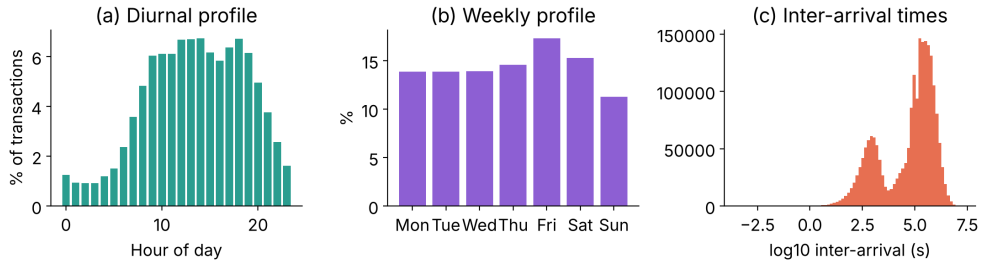


Fig. F3 Temporal structure: (a) hourly share of transactions; (b) weekday share; (c) distribution of within-customer inter-arrival times, log10 seconds.

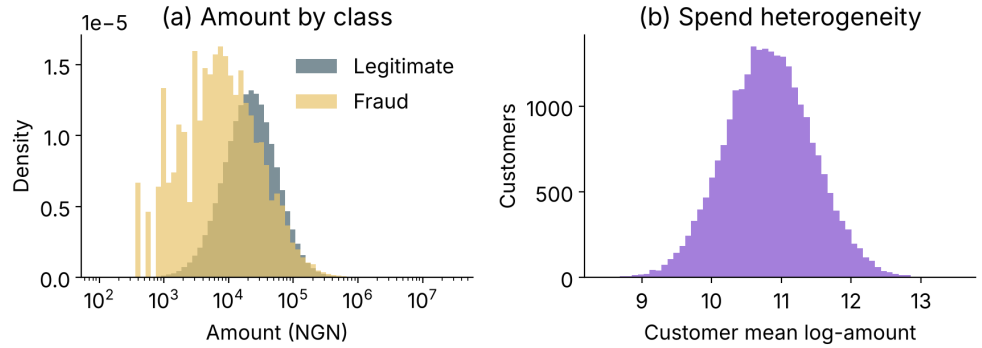


Fig. F4 Amounts: (a) legitimate versus fraudulent amounts (density, log axis); (b) distribution of customers' mean log amount.

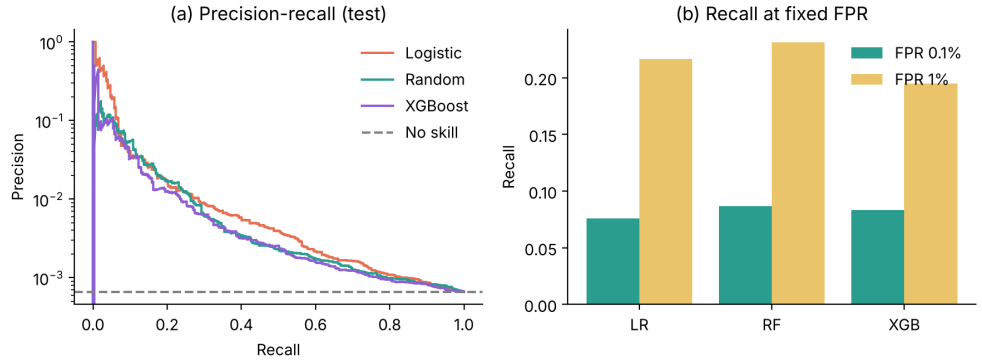


Fig. F5 Benchmark models on the test period: (a) precision–recall curves (log precision axis) with the no-skill line; (b) recall at fixed false-positive rates.

References

[1] National Bureau of Statistics: Demographic Bulletin 2023: projected population by state, as reported by 21st Century Chronicle. https://21stcenturychronicle.com/kano-records-15-7m-lagos-13-7m-as-nigerias-population-hits-221m-in-2023-nbs/. Accessed 13 October 2026 (2023)

[2] A2F Nigeria: EFInA report: Nigeria's formal financial inclusion grows to 64% in 2023. https://a2f.ng/efina-report-nigerias-formal-financial-inclusion-grows-to-64-in-2023/. Reports the EFInA Access to Finance Survey 2023: 52% of adults banked, 93% mobile-phone adoption, 37% of rural and 17% of urban adults financially excluded. Press report of the survey; accessed 16 October 2026 (2024)

[3] Limpert, E., Stahel, W.A., Abbt, M.: Log-normal distributions across the sciences: keys and clues. BioScience **51**(5), 341–352 (2001) https://doi.org/10.1641/0006-3568(2001)051[0341:LNDATS]2.0.CO;2

[4] Barabási, A.-L.: The origin of bursts and heavy tails in human dynamics. Nature **435**, 207–211 (2005)

[5] Bahnsen, A.C., Aouada, D., Stojanovic, A., Ottersten, B.: Feature engineering strategies for credit card fraud detection. Expert Systems with Applications **51**, 134–142 (2016) https://doi.org/10.1016/j.eswa.2015.12.030

[6] Nigeria Inter-Bank Settlement System (NIBSS): Annual fraud landscape: January–december 2023. Technical report, NIBSS, Abuja, Nigeria (Apr 2024). https://nibss-plc.com.ng/wp-content/uploads/2024/04/2023-Annual-Fraud-Landscape.pdf

[7] Ojoko, I.: Nigeria's financial sector suffers NGN 52.26 billion loss to fraud in 2024 – NIBSS report. Nairametrics, 26 February 2025. Accessed 9 October 2026 (2025). https://nairametrics.com/2025/02/26/nigerias-financial-sector-suffers-n52-26-billion-loss-to-fraud-in-2024-nibss-report/

[8] Nigeria Inter-Bank Settlement System (NIBSS): Digital payment fraud drops 51% to NGN 25.85b, Lagos accounts for 63%. NIBSS press release, 22 January 2026. Accessed 9 October 2026 (2026). https://nibss-plc.com.ng/digital-payment-fraud-drops-51-to-n25-85b-lagos-accounts-for-63/