

Technical summary - fraud_200k

fraud_200k.csv - 200,000 rows - run run_1037d177a96f - 2026-10-11 10:22 - the plain version is [REPORT.pdf](#)

In brief

Your question: What are the characteristics of fraudulent transactions?

The answer, in short (written by the AI; every number checked against the tool results)

- Fraud occurs only in transfers and cash-outs; payments, cash-ins and debits carry none.
- Fraudulent transactions are typically much larger, and the fraud rate climbs steeply with amount.
- Large amounts mark fraud among cash-outs, but not among transfers.
- Among transfers and cash-outs, type, timing and amount each stay associated with fraud when the others are held fixed.
- The fraud rate rises late in the month mainly because legitimate volume collapses while fraud keeps a steady pace.

Status: Gate 4 raised a warning; 9 thing(s) worth a look in section 4.


1 What we found

Answers and findings are written by the AI from the tool results; Gate 4 checked that every number in them is one a tool reported.

How common is each transaction type, and what share of each type is fraudulent?

answered. CASH_OUT is 70,084 rows and PAYMENT 67,660, with CASH_IN 43,991, TRANSFER 16,898 and DEBIT 1,367. Fraud occurs only in TRANSFER (0.7%, interval 0.6% to 0.9%) and CASH_OUT (0.2%, interval 0.1% to 0.2%). PAYMENT, CASH_IN and DEBIT show 0.0%.


- Fraud occurs only in transfers and cash-outs; payments, cash-ins and debits carry none. (strong)



Are fraudulent transactions larger than legitimate ones, and how does the fraud rate change across amount ranges?

answered. Yes. Fraudulent amounts have a median of 401,192 against 73,996 for legitimate ones. The fraud rate rises from 0.0% below 10,000 to 1.8% for 1,000,000 to below 5,000,000 and 4.7% at 5,000,000 and over. Fraud amounts top out at exactly 10,000,000.


- Fraudulent transactions are typically much larger, and the fraud rate climbs steeply with amount. (strong)
- Fraudulent amounts stop at exactly 10,000,000, a cap visible in the simulation. (moderate)



Does the fraud rate vary over the simulated month (by time step)?

answered. The fraud rate is about 0.1% in each of the first three simulated weeks. It rises to 0.8% for steps 504-672 and 1.0% from 672 on. Transaction volume is mostly between 10,188 and 18,027 rows per time range up to step 421, then drops to between 10 and 1,984 rows per range, while fraud counts stay roughly level. The rise in rate mostly reflects fewer legitimate transactions, not more fraud.


- The fraud rate rises late in the month mainly because legitimate volume collapses while fraud keeps a steady pace. (moderate)



Within the transaction types where fraud occurs (for example TRANSFER and CASH_OUT), do fraudulent amounts differ from legitimate ones?

answered. It depends on the type. Within CASH_OUT, fraud is much larger (median 455,710 against 146,796), and cash-outs of 1,000,000 and over are mostly fraudulent. Within TRANSFER, fraudulent amounts are not larger (median 371,771 against 485,695), and the fraud rate stays around 0.4% to 1.2% across amount bands below 5,000,000.


- Large amounts mark fraud among cash-outs, but not among transfers. (strong)



Taken together, which of type, amount and time step still go with fraud when the others are held fixed?

answered. Within the two types that carry fraud, all three features stay associated with fraud. TRANSFER goes with x3.67 the odds of CASH_OUT, each 100 steps later with x1.68, and each extra 1,000,000 with x1.12. The full model including the zero-fraud types did not settle on stable estimates, so those types get no usable estimate.

- Among transfers and cash-outs, type, timing and amount each stay associated with fraud when the others are held fixed. (strong)



More: [1_analysis.pdf](#) - every finding with all its charts and tool results

2 How good the data is

Rows	200,000 in the file, 200,000 read, 0 could not be read; 200,000 after cleaning (0 removed)
Problems found	10: values stored as the wrong type (8), personal data (2)
Values changed	416,320 value(s) changed; 8 column(s) converted to the right type; 0 value(s) turned into missing
Still missing	nothing: every value is present
Checks	loading (Gate 1): pass; cleaning (Gate 2): pass; review: accepted

What this means for the findings

Nothing in the data itself limits the findings.

More: [2_data_quality.pdf](#) - the data as received, every problem, change and check

Column by column in Excel: [2_data_columns.xlsx](#)

3 Decisions made for you

	Question	Decided	By	How sure
J3	Should isFlaggedFraud be used as a feature, or excluded as a rule-based flag closely tied to the target?	Keep isFlaggedFraud as a feature	the AI	low: check it

	Question	Decided	By	How sure
J1	Should the 2,799 far outliers in amount (beyond 3x IQR, max 38,056,471.01) be kept or clipped?	Keep all amounts as they are	the AI	medium
J2	Should the four balance columns (oldbalanceOrig, newbalanceOrig, oldbalanceDest, newbalanceDest) be dropped from the table, or kept but excluded from analysis?	Keep the columns in the table, marked excluded	the AI	medium

To change one: edit edit/decisions.yaml, then run the cleaning and the analysis again.

More: [3_decisions.pdf](#) - every decision with its reasons, alternatives and effect

4 Wrap-up: what might be wrong, and what we recommend

Gathered from the checks, the decisions and the review, plus the analysis's own doubts (marked AI). It is what to keep in mind, and where to look.

What might be wrong

- J3 was decided by the AI with low confidence: 'Keep isFlaggedFraud as a feature'. Should isFlaggedFraud be used as a feature, or excluded as a rule-based flag closely tied to the target?
- AI The reviewer asked (isFlaggedFraud): Should isFlaggedFraud stay in the analysis as a feature? It is almost constant in this sample, so it adds little for describing fraud.
- AI The reviewer asked (isFraud): How should the very low fraud rate (0.12%) be handled in the analysis? The cleaning left it as is, which is correct, but any later comparison or model needs to account for it.
- Gate 4: T8 logistic_regression: The model did not settle on stable estimates, often because a category has only one outcome. Effects with very wide intervals are unreliable.
- AI F4: the late-month fraud rate is driven by a collapse in total volume, which looks like a feature of how the simulator spreads transactions. It should not be read as fraudsters preferring the end of the month.
- AI F3: the striking 77.8% fraud rate among large cash-outs rests on 54 rows. The jump also partly follows from legitimate cash-outs of that size being very rare in the simulation.
- AI F5: the 10,000,000 cap is probably a hard limit in the simulator. A check at one exact value with 92 rows is thin evidence.
- AI F6: treating amount as a straight-line effect hides that amount separates fraud strongly within CASH_OUT but hardly within TRANSFER, so the x1.12 per 1,000,000 averages two different patterns.
- AI F2: the overall amount gap is partly type mix, because the types with no fraud (PAYMENT, CASH_IN) are mostly small amounts.

What we recommend

- Check J3; if it is wrong, change it in edit/decisions.yaml, then clean and analyse again.
- AI Answer it in the brief (guidance or column notes) and run again.
- See details/1_analysis.pdf.
- AI Treat transfers and cash-outs as the only relevant population for any fraud description or model of this data; the other types carry no fraud here.
- AI Describe amount separately for each type. For cash-outs, 1,000,000 and over is a strong marker; for transfers, amount alone tells little.
- AI Do not use step as a fraud signal without first adjusting for transaction volume per step, since the rise in rate reflects the simulator's volume pattern.
- AI Check the PaySim documentation for a 10,000,000 cap on fraudulent amounts before using it as a marker.
- AI Confirm the plan's exclusion of isFlaggedFraud over the cleaning step's low-confidence choice (J3) to keep it; with a single flagged row it cannot be evaluated.
- AI Worth asking next: Within transfers, what does distinguish fraud from legitimate transfers, if not the amount, using only the allowed columns?
- AI Worth asking next: Are fraudulent transfers and cash-outs paired in time (a transfer followed by a cash-out), as the PaySim fraud scenario suggests?

More detail

details/1_analysis.pdf	every finding with all its charts and tool results; the plan
details/2_data_quality.pdf	the data as received, every problem, every change, the review
details/2_data_columns.xlsx	the data column by column, in Excel
details/3_decisions.pdf	every decision with its reasons, alternatives and measured effect

To change the result

- edit/decisions.yaml: the cleaning decisions, then run `pipeline.clean`
- edit/analysis_plan.yaml: the analysis questions, then run `pipeline.analyze`

The data

- data/clean_data.csv and .parquet: the cleaned data

logs/ holds the machine-readable audit trail.