Cash-Flow Underwriting & Fraud Gate — Results & Recommendations

github.com/guptagaurav-info/cashflow-underwriting-fraud · public data: 68K bank-transaction descriptors (MIT), PaySim 6.36M · simulated SMB book: 6,000 applicants, 1.72M transactions · **out-of-time test Jul–Dec 2025 (n = 1,578)**

0.844
Cash-flow PD AUC vs 0.659 bureau-only (non-fraud test)

97%
Fraud caught at 5% review rate · false positives 0.3%

47 → 4
Fraudulent applications funded: no gate vs gate

66% → 100%
Loan-proceeds recall: classifier alone vs with agent routing

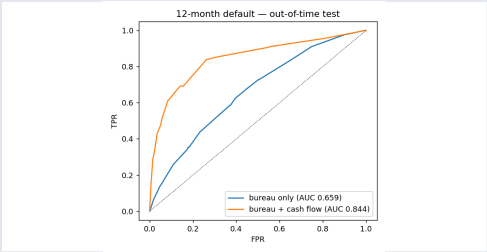
FINAL OUTPUTS

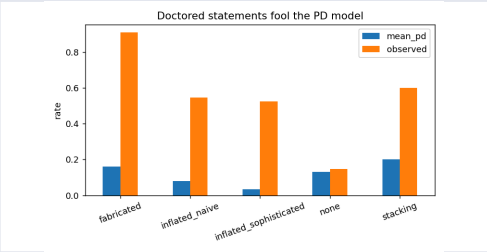
Policy (test window)	Funded	Bad rate	Fraud funded	P&L
Bureau-only PD	790	7.5%	25	$4.5M
Cash-flow PD, no fraud gate	1,129	7.1%	47	$6.1M
Fraud gate → cash-flow PD → offer	**1,083**	**5.3%**	**4**	**$7.3M**

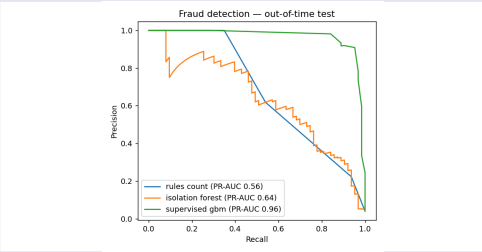
Fraud detector	PR-AUC	Recall @5%
9 explainable rules	0.56	57%
Isolation Forest	0.64	68%
Supervised GBM	**0.96**	**97%**

Fraud type	PD assigned	Actual default	Caught by gate
Inflated (sophisticated)	3.5%	52%	95%
Inflated (naive)	8.1%	55%	100%
Fabricated statement	16.1%	91%	100%
Stacking	20.1%	60%	75%

Classifier test	Result
Random split / unseen merchants	99.9% / 82.5%
Auto-accept lane (unseen merchants)	53% of rows at 98.6%
Accounts with >5% true-revenue error	1.9% → 0.4%
Entity resolution (pairwise)	precision 1.00 · recall 0.986







DATA & METHODOLOGY

SOURCES

- Synthetic: 68,000 US bank-transaction descriptors, 17 categories (DoDataThings, MIT); PaySim 6.36M mobile-money transactions with fraud labels. Both sha256-verified.
- Simulated (no public data links SMB transactions to outcomes): 6,000 businesses, 8 industries, 1.72M transactions, 8,237 applications across 40 lenders.

METHOD

- Simulator: default = cash runs out in a forward 12-month simulation with the new payment, existing MCA/term debt and unobserved shocks; fraud edits only the document the lender sees.
- Classifier: TF-IDF + logistic regression trained on 1,000 labelled accounts; routing to an agent (rules → account context → retrieval → optional LLM → human review).
- Features in SQL as of application date; entity resolution by blocking + fuzzy names + union-find.
- Out-of-time split (train ≤ Jun-2025); fraud excluded from PD training; isotonic calibration on out-of-fold predictions.

KEY CALCULATIONS

- True revenue = deposits excluding transfers, loan proceeds and refunds.
- Pro-forma DSCR = (revenue − opex − tax − owner draws) ÷ (existing debt service + new payment).
- Runway days = average balance ÷ average daily outflow.
- Offer = min(request, (free cash flow ÷ 1.25) × 12 ÷ 1.25 factor), floored to $500.
- P&L: good deal earns 17% of advance (factor − 1 − 8% cost); default loses 65%. Recall @5% = fraud caught in top 5% of scores.

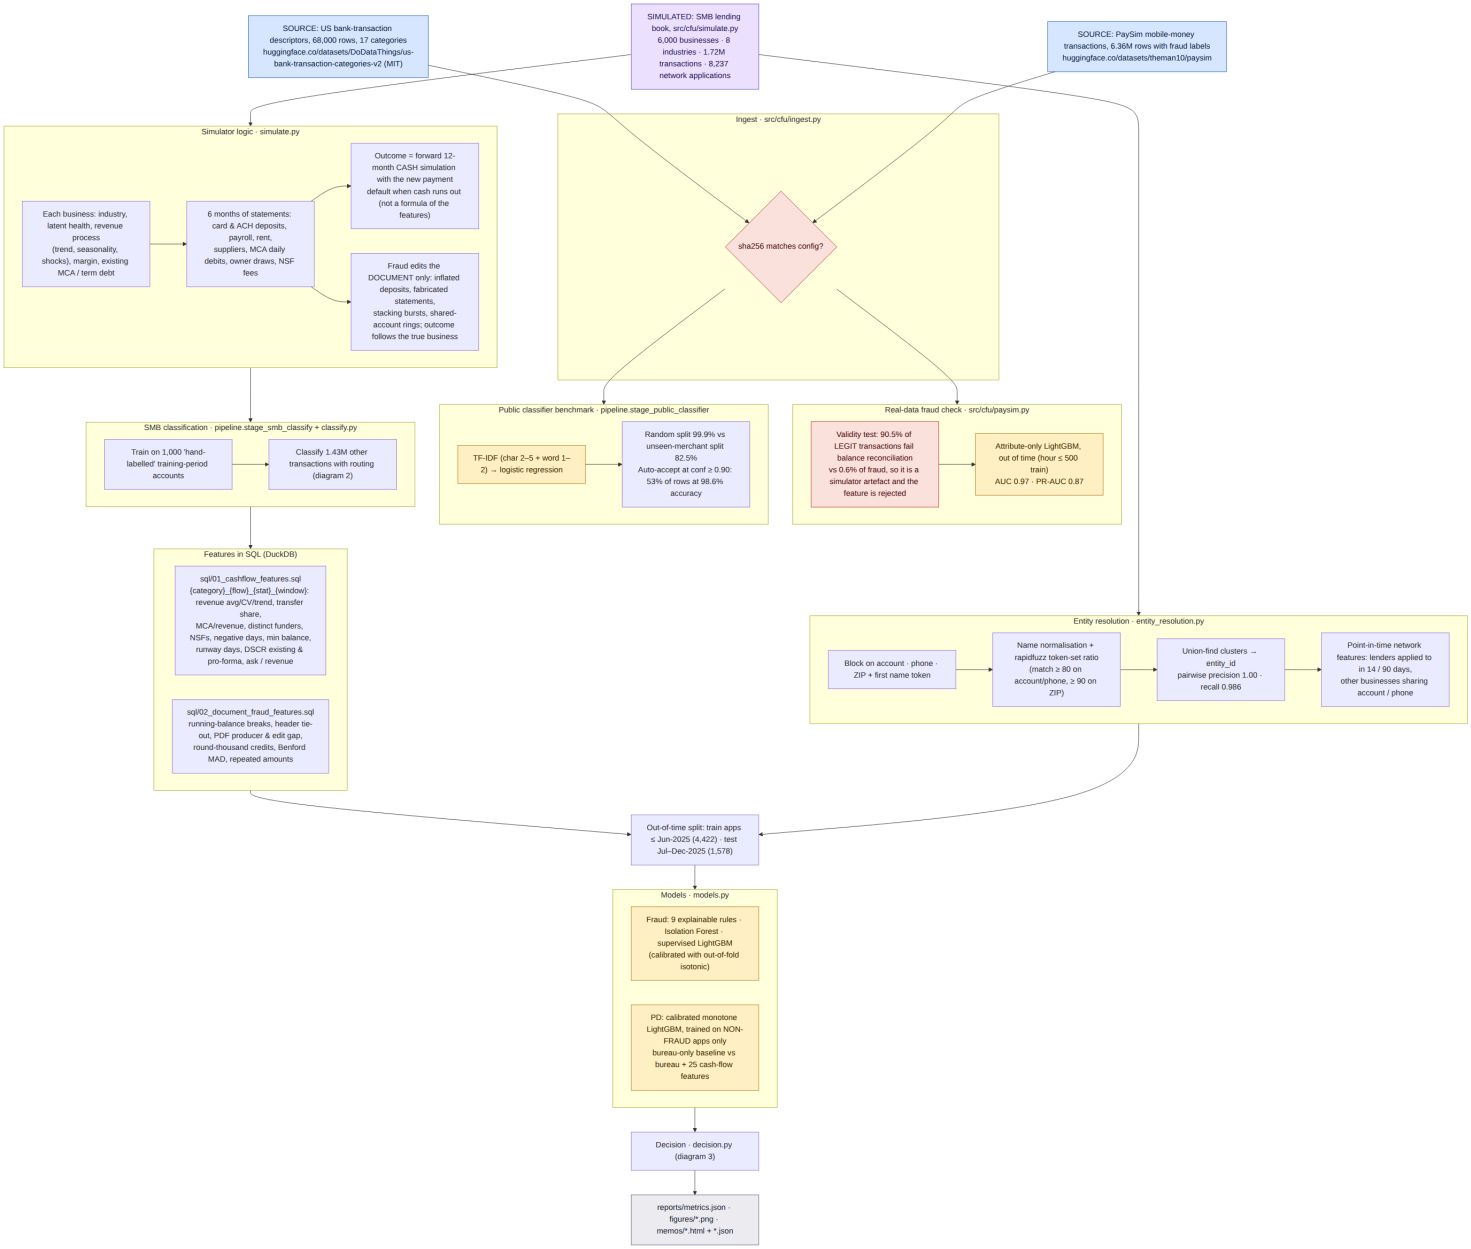
RECOMMENDATIONS

1. **Run fraud as a gate before credit, not as a PD input.** Doctored statements look healthy to the credit model; refer at fraud score ≥ 0.30 (0.3% false positives) and exclude confirmed fraud from PD training.
2. **Underwrite on cash flow.** Adding cash-flow features lifts AUC from 0.659 to 0.844; top drivers are liquidity runway, ask ÷ monthly revenue and pro-forma DSCR. Approve PD ≤ 12%, decline ≥ 30%, refer between.
3. **Size every offer from debt capacity** (free cash flow after existing debt ÷ 1.25× DSCR). It cuts exposure on stretched borrowers instead of only accepting or rejecting them.
4. **Route by confidence, generic text and dollar materiality.** Send only ~16% of transactions to the agent; it fixes the categories that move true revenue (loan proceeds) at a 0.05% human-review rate.
5. **Add post-funding monitoring for stacking.** Only 75% is visible at application time; ongoing bank data closes the gap.
6. **Before production: replace the simulated book with real network outcomes** and validate on each client's funded book (swap-set at their approval rate, their default definition).

The SMB book is simulated (mechanism-driven: default from forward cash, fraud edits documents only) — trust the method and relative lift, not the levels. Illustrative economics. Portfolio project, not a production model.

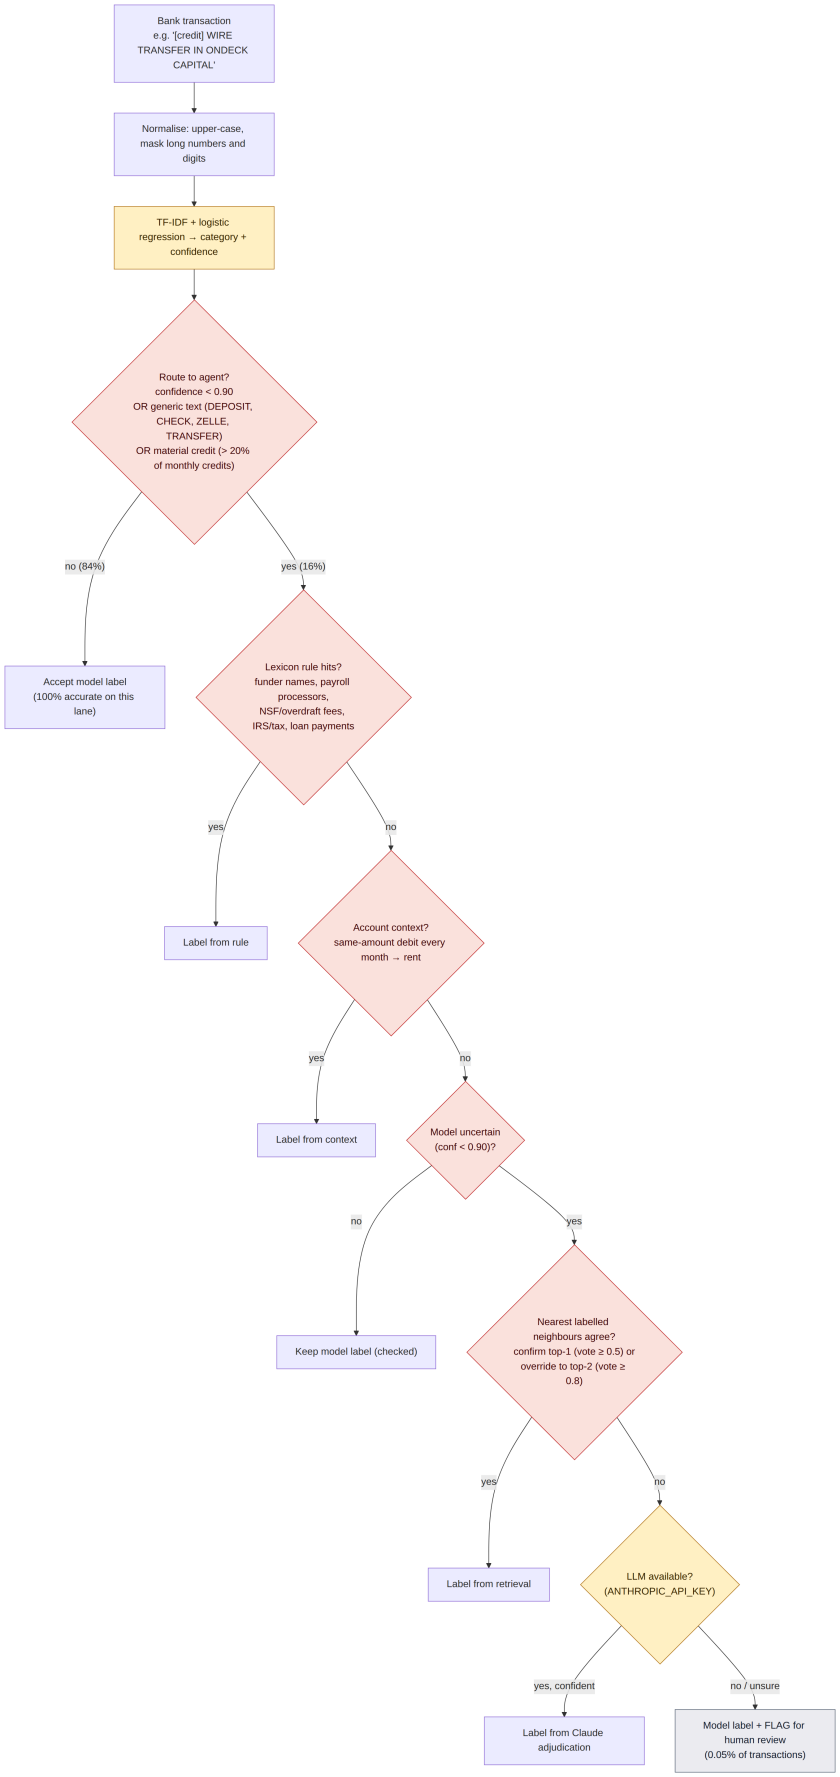
cashflow-underwriting-fraud · Pipeline: sources, steps and logic

Blue = external data source · purple = simulated data · lavender = code step (file named) · yellow = model · red = check or gate · grey = output.



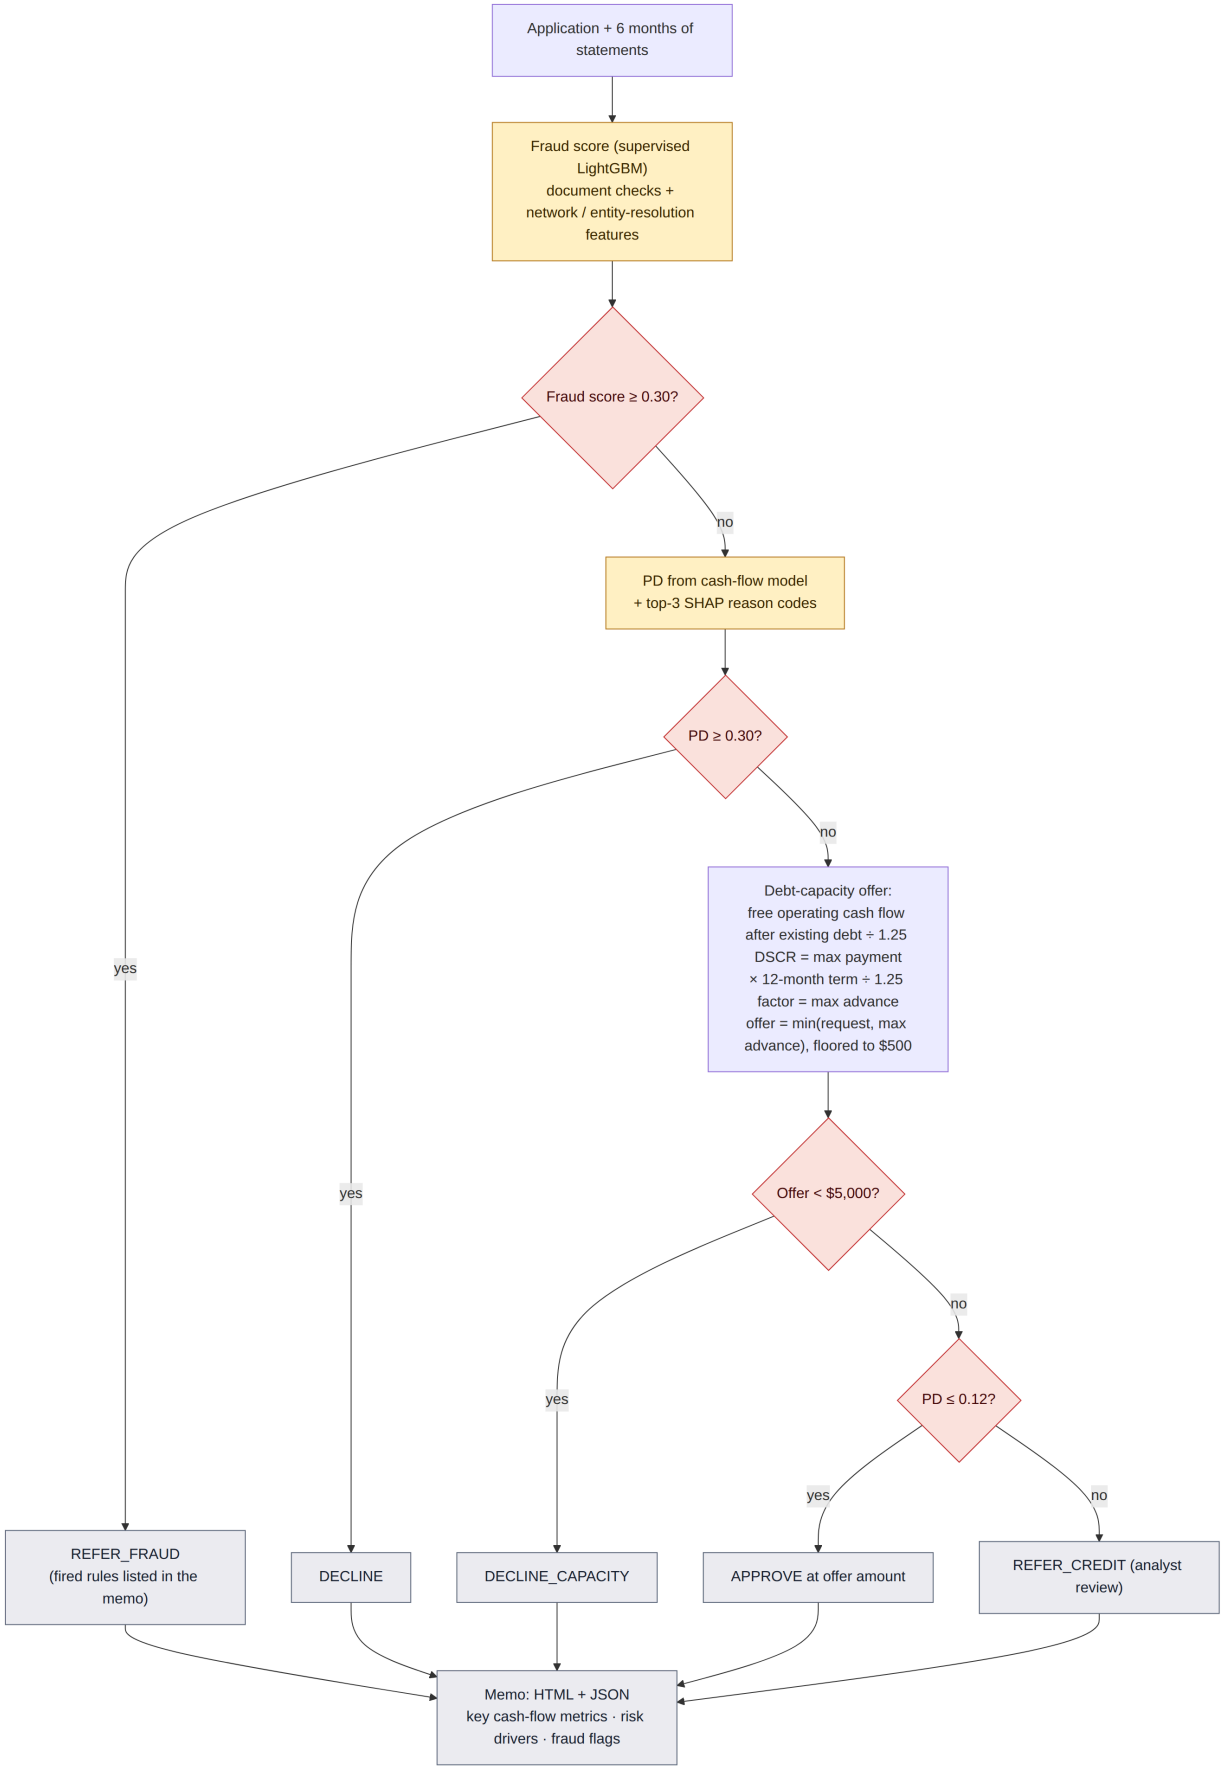
cashflow-underwriting-fraud · Transaction classification with agent routing

Blue = external data source · purple = simulated data · lavender = code step (file named) · yellow = model · red = check or gate · grey = output.



cashflow-underwriting-fraud · Decision logic for one application

Blue = external data source · purple = simulated data · lavender = code step (file named) · yellow = model · red = check or gate · grey = output.



cashflow-underwriting-fraud · Production design (proposed, not implemented)

Blue = external data source · purple = simulated data · lavender = code step (file named) · yellow = model · red = check or gate · grey = output.

