Exercises: Lesson 7.2 – RegTech — Automating Compliance at Scale

Module 7: The Compliance Problem

Prof. Dr. Joerg Osterrieder

Exercise 1: False Positive Cost Analysis

A mid-sized bank's transaction monitoring system generates the following monthly statistics:

Metric	Value
Total alerts generated	45,000
False positive rate	96%
Average analyst time per alert	40 minutes
Fully loaded analyst cost	\$55 per hour
SARs filed (from true positives)	720

Tasks:

a) Calculate the number of true positives and false positives per month.

b) Calculate the total monthly cost of investigating all alerts.

c) What percentage of total investigation cost is spent on false positives?

d) An ML system reduces the false positive rate to 55% while maintaining the same 1,800 true positives. Calculate the new total alert volume and the monthly cost savings.

e) If the ML system costs \$1.2 million per year to license and maintain, what is the annual net savings (or cost)?

Exercise 2: ML Model Threshold Tuning

An ML transaction monitoring model scores each transaction from 0 (not suspicious) to 1 (highly suspicious). The compliance team tests three score thresholds:

Threshold	Alerts Generated	True Positives Captured	False Positives
0.4	25,000	1,900	23,100
0.6	12,000	1,750	10,250
0.8	5,000	1,400	3,600

Tasks:

a. Calculate the precision (TP / total alerts) and false positive rate for each threshold.

b. If the bank's analysts can review at most 15,000 alerts per month, which threshold(s) are operationally feasible?

c. Calculate the "cost of missed true positives" for threshold 0.8 vs. 0.4, assuming each missed SAR has an expected regulatory penalty of $50,000.

d. Recommend a threshold and justify your choice, balancing operational capacity, detection rate, and regulatory risk.

e. If the team uses a two-tier system (ML auto-clears scores below 0.4 and prioritizes scores above 0.8 for immediate review), what is the remaining middle-tier volume?

Exercise 3: Regulatory Reporting Error Analysis

A bank submits 120 regulatory reports per quarter. The table below shows error rates before and after implementing an XBRL-based automated reporting system:

Metric	Manual Process	XBRL Automated
Reports requiring restatement	12%	1.5%
Average restatement cost	\$18,000	\$4,000
Staff hours per report	45	8
Staff cost per hour	\$65	\$65
System license (per quarter)	—	\$95,000

Tasks:

a. Calculate the quarterly restatement cost under each approach.

b. Calculate the quarterly staff cost for report preparation under each approach.

c. What is the total quarterly cost (staff + restatement + license) for each approach?

d. Calculate the annual net savings from automation.

e. Beyond cost savings, list two non-financial benefits of XBRL automation.

Exercise 4: Biometric Verification System Evaluation

A neobank is evaluating two identity verification vendors. Each vendor processes 50,000 onboarding attempts per month:

Metric	Vendor A	Vendor B
Document verification accuracy	98.5%	99.2%
Facial match accuracy	99.0%	98.5%
Liveness detection accuracy	97.0%	99.5%
Average verification time	45 seconds	90 seconds
Cost per verification	\$0.80	\$1.40
False rejection rate (legitimate users)	3.5%	1.8%

Tasks:

a. Calculate the combined pass-through probability for a legitimate user under each vendor (multiply accuracies of all three stages).

b. If each false rejection costs the neobank \$25 in customer acquisition loss, calculate the monthly false rejection cost for each vendor.

c. Calculate the total monthly cost (verification fees + false rejection costs) for each vendor.

d. Which vendor would you recommend? Consider cost, security (liveness accuracy), and user experience (speed).

e. A deepfake attack targets 500 onboarding attempts per month. How many would each vendor's liveness detection catch?

Exercise 5: Differential Privacy Budget Allocation

A regulator asks a bank to report the average transaction amount per customer segment. The bank wants to use differential privacy to protect individual records.

Given:

- Number of customer segments: 4
- True average transaction amounts: Retail \$85, SME \$2,400, Corporate \$45,000, High-Net-Worth \$125,000
- Total privacy budget: $\varepsilon = 2.0$
- Noise mechanism: Laplace noise with scale $b = \frac{\Delta f}{\varepsilon}$
- Sensitivity Δf: the maximum change from adding/removing one record (assume \$500 for retail, \$5,000 for SME, \$100,000 for corporate, \$250,000 for HNW)

Tasks:

a. If the bank splits the budget equally across 4 queries ($\varepsilon_i = 0.5$ each), calculate the Laplace noise scale b for each segment.

b. Calculate the expected noise as a percentage of the true average for each segment. Which segment has the worst signal-to-noise ratio?

c. Suggest a non-uniform budget allocation that reduces noise for the most sensitive segment. Justify your allocation.

Exercise 6: Federated Learning vs. Centralized Training

A consortium of three banks considers two approaches to training a shared AML detection model:

Bank	Transactions/month	SARs filed/month	SAR rate
Bank A	12,000,000	3,600	0.030%
Bank B	3,000,000	1,200	0.040%
Bank C	800,000	480	0.060%

Tasks:

a. If all data were pooled (centralized), calculate the combined SAR rate.

b. In federated learning, each bank trains locally and sends gradients. If Bank C's local model achieves 85% recall (true positive rate) alone but the federated model achieves 92%, how many additional true SARs does Bank C detect per month?

c. What privacy risk exists in centralized training that federated learning eliminates?

d. Identify one weakness of federated learning when banks have very different transaction profiles.

e. If Bank B leaves the consortium, estimate the impact on the federated model's ability to detect cross-bank layering schemes.

Exercise 7: RegTech Return on Investment

A bank is evaluating a RegTech platform that consolidates transaction monitoring, regulatory reporting, and KYC into a single system. Current and projected costs:

Cost Category	Current Annual	With RegTech
Compliance staff (FTE $\times$ salary)	\$8,500,000	\$4,200,000
Legacy system licenses	\$1,800,000	\$0
RegTech platform license	\$0	\$2,400,000
Implementation (Year 1 only)	\$0	\$1,500,000
Regulatory fines (3-year avg.)	\$3,200,000	\$800,000
Training and change management	\$0	\$350,000

Tasks:

a. Calculate the total cost of ownership (TCO) for Year 1 under each scenario.

b. Calculate the TCO for Years 2–3 (assume no implementation or training costs).

c. Calculate the 3-year cumulative savings.

d. Calculate the simple payback period (months until cumulative savings exceed cumulative implementation and transition costs).

e. List two risks that could reduce the actual ROI below the projected figures.

Exercise 8: Secure MPC for Sanctions Screening

Two banks want to check whether any of their customers appear on the other bank's internal watchlist, without revealing their full customer lists.

Setup:

- Bank X has 2,000,000 customers
- Bank Y has 1,500,000 customers
- They use a Private Set Intersection (PSI) protocol (a form of secure MPC)
- PSI reveals only the intersection (common entries), not the non-matching records
- Computational cost: 0.002 seconds per comparison pair
- The naive approach compares every customer in X against every customer in Y

Tasks:

a. How many comparison pairs would the naive approach require? Estimate the total computation time.

b. A hash-based PSI optimization reduces comparisons to $O(n + m)$ rather than $O(n \times m)$. Calculate the new number of operations and computation time.

c. If the match rate is 0.05%, how many shared customers are expected?

d. What information does each bank learn from the PSI result? What information remains hidden?