Quiz: Lesson 7.2 – RegTech — Automating Compliance at Scale

Module 7: The Compliance Problem

Prof. Dr. Joerg Osterrieder

Question 1 (Understand)

What is the primary purpose of RegTech?

A. To replace financial regulators with automated systems
B. To help financial institutions avoid regulatory requirements
C. **To help financial institutions comply with regulations more efficiently using technology**
D. To create new financial regulations using artificial intelligence

Question 2 (Understand)

Which of the following is NOT one of the five functional categories of RegTech described in this lesson?

A. Transaction Monitoring
B. Regulatory Reporting
C. Identity Verification
D. **Algorithmic Trading**

Question 3 (Understand)

In transaction monitoring, what is a Suspicious Activity Report (SAR)?

A. An internal document summarizing all false positive alerts

B. A daily summary of all transactions above $1,000

C. **A formal report filed with regulators when a transaction or pattern is deemed genuinely suspicious after investigation**

D. An automated alert generated by a rule-based monitoring system

Question 4 (Understand)

What does XBRL stand for, and what is its primary function?

A. eXtended Banking Regulation Language; it defines compliance rules for banks

B. **eXtensible Business Reporting Language; it tags financial data with machine-readable labels for automated regulatory reporting**

C. eXecutable Business Risk Logic; it automates risk calculations

D. eXternal Bank Reconciliation Ledger; it reconciles interbank payments

Question 5 (Apply)

A bank's rule-based transaction monitoring system generates 80,000 alerts per month with a 96% false positive rate. How many alerts are true positives (genuinely suspicious)?

A. 800
B. 76,800
C. **3,200**
D. 8,000

Question 6 (Apply)

A customer makes five cash deposits of $9,500 each over three consecutive days, staying just below the $10,000 Bank Secrecy Act reporting threshold. Which suspicious pattern does this match?

A. Layering through shell companies

B. Rapid fund transfer (round-tripping)

C. **Structuring (also called "smurfing")**

D. Sanctions evasion

Question 7 (Apply)

An ML-based transaction monitoring model outputs a suspicion score between 0 and 1. The compliance team sets a threshold of 0.7 for generating alerts. If the threshold is lowered to 0.5, what is the expected effect?

A. Fewer alerts and fewer false positives

B. Fewer alerts but more false positives per alert

C. **More alerts, more false positives, but also more true positives captured**

D. No change, since the model's accuracy is fixed

Question 8 (Apply)

A digital identity verification system uses document AI to extract data from a passport, facial recognition to match a selfie, and liveness detection. A fraudster submits a stolen passport with a deepfake video of the passport holder. Which component is most likely to catch this fraud?

A. Document AI (it will detect the passport is stolen)

B. Facial recognition (the deepfake will not match)

C. **Liveness detection (it analyzes micro-textures and depth cues to distinguish real faces from synthetic ones)**

D. None — the system will be fooled

Question 9 (Apply)

A bank currently spends $35 per alert investigation and processes 60,000 alerts per month with a 95% false positive rate. An ML system reduces the false positive rate to 60% while keeping the same number of true positives. How much does the bank save per month on alert investigation costs?

- A $1,050,000
- B **$1,837,500**
- C $735,000
- D $2,100,000

Hint: TP = 60,000 × 0.05 = 3,000. New total alerts = 3,000 / 0.40. Savings = (old alerts − new alerts) × $35.

Question 10 (Apply)

In federated learning for AML detection, what is sent from each participating bank to the central aggregation server?

A. Encrypted copies of all customer transaction records
B. Anonymized customer profiles
C. **Model weight updates (gradients) computed on local data**
D. Suspicious Activity Reports for cross-referencing

Question 11 (Apply)

A compliance team is designing a differential privacy mechanism for sharing aggregate transaction statistics with a regulator. They must choose the privacy parameter ε. Which statement is correct?

A. Higher ε means stronger privacy and more noise

B. **Lower ε means stronger privacy but less accurate (noisier) results**

C. ε has no effect on the accuracy of the output

D. ε must always be set to exactly 1.0 for regulatory compliance

Question 12 (Apply)

A bank wants to perform sanctions screening against another bank's customer list without either bank revealing its customers to the other. Which privacy-preserving technique is most appropriate?

A. Differential privacy

B. Federated learning

C. **Secure multi-party computation (MPC)**

D. XBRL reporting

Question 13 (Analyze)

A bank switches from rule-based to ML-based transaction monitoring and reduces its false positive rate from 97% to 55%. However, the regulator expresses concern. What is the most likely reason?

A. The regulator prefers a higher false positive rate

B. ML is illegal for compliance purposes in most jurisdictions

C. **The regulator wants assurance that the true positive detection rate has not decreased, and requires the ML model to be explainable**

D. The bank should not have changed its monitoring system without permission

Question 14 (Analyze)

Why is a hybrid approach (rules + ML) more common in production transaction monitoring than a pure ML approach?

A. ML models are always less accurate than rules

B. **Some thresholds are legally mandated (e.g., the $10,000 reporting threshold) and must be implemented as hard rules, while ML handles pattern detection**

C. Regulators have banned the use of ML in all jurisdictions

D. Rules are cheaper to implement and maintain than ML models

Question 15 (Analyze)

A federated learning system for AML detection involves four banks. Bank A has 10 million transactions, Bank B has 500,000, Bank C has 8 million, and Bank D has 200,000. What problem is most likely to arise?

A. The model will only learn patterns from Bank A because it has the most data

B. **The global model may be biased toward the transaction patterns of larger banks, underrepresenting typologies common in smaller banks**

C. Banks B and D will be forced to share their raw data to compensate

D. The model will converge faster with more heterogeneous data

Question 16 (Analyze)

A RegTech vendor claims its system reduces compliance costs by 60%. A bank's Chief Compliance Officer (CCO) should be most skeptical about which aspect?

A. The technology behind the claim

B. The vendor's financial stability

C. **Whether the cost reduction accounts for implementation costs, ongoing model maintenance, and the risk of missed true positives leading to regulatory fines**

D. Whether the vendor has other bank clients

Question 17 (Analyze)

Active liveness detection requires a user to perform an action (e.g., turn their head), while passive liveness analyzes a single selfie. A neobank optimizing for conversion rate during onboarding would prefer:

- A. Active liveness for all users, because it is more secure
- B. **Passive liveness as the default (lower friction), with active liveness as step-up verification for flagged cases**
- C. No liveness detection, because it reduces conversion
- D. Active liveness only for users in high-risk jurisdictions

Question 18 (Analyze)

A bank uses network analysis as part of its transaction monitoring. It discovers that Customer X sends funds to Customer Y, who sends to Customer Z, who sends back to Customer X within 48 hours. What typology does this suggest?

- A. Structuring (smurfing)
- B. **Layering / round-tripping through intermediary accounts**
- C. Trade-based money laundering
- D. Tax evasion through offshore accounts

Question 19 (Evaluate)

A regulator proposes mandating that all banks in its jurisdiction use a single, government-operated transaction monitoring platform. What is the strongest argument *against* this proposal?

A. Banks would save money by not maintaining their own systems

B. It would improve detection by aggregating data across institutions

C. **A single centralized platform creates a massive privacy risk and single point of failure; it also removes competitive incentives for banks to innovate in compliance technology**

D. It would be faster to implement than bank-specific solutions

Question 20 (Evaluate)

A consortium of five banks is deciding between federated learning and secure MPC for a joint AML detection initiative. The consortium processes 50 million transactions per day and needs results within 24 hours. Which recommendation is most sound?

A. Secure MPC, because it provides stronger privacy guarantees

B. **Federated learning, because it scales better to high transaction volumes and can train models within the time constraint; MPC's computational overhead makes it impractical for this volume**

C. Neither — they should simply share anonymized data

D. Both simultaneously, using MPC for model training and federated learning for inference

Answer Key

1. (C) – Efficient compliance via tech
2. (D) – Algorithmic Trading
3. (C) – Formal report to regulators
4. (B) – XBRL tags financial data
5. (C) – 3,200 true positives
6. (C) – Structuring (smurfing)
7. (C) – More alerts, more FP, more TP
8. (C) – Liveness detection
9. (B) – \$1,837,500 monthly savings
10. (C) – Model weight updates
11. (B) – Lower ε = stronger privacy
12. (C) – Secure MPC
13. (C) – Explainability + TP rate
14. (B) – Legal mandates require rules
15. (B) – Bias toward larger banks
16. (C) – Hidden costs and TP risk
17. (B) – Passive default, active step-up
18. (B) – Layering / round-tripping
19. (C) – Privacy risk + no innovation
20. (B) – FL scales better for volume