Digital Finance — Exam Preparation Guide

40 Practice Questions Across All 8 Modules

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

How to Use This Guide

- **40 questions:** 5 per module (2 multiple choice, 2 short answer, 1 calculation/analysis)
- **Answers** follow each module's question set on a separate slide
- **Difficulty:** Questions target BSc-level understanding — not memorization but comprehension and application
- **Study tip:** Try each question *before* looking at the answer. Write your answer, then compare.

These are practice questions only. The actual exam may differ in format and difficulty.

Module 1: The Cost Problem — Questions

Q1.1 [Multiple Choice] Which component of the Merchant Discount Rate (MDR) is typically the largest?

a. Gateway fee
b. Network assessment fee
c. Interchange fee
d. Processor fee

Q1.2 [Multiple Choice] The EU Interchange Fee Regulation caps consumer debit interchange at:

a. 0.1%
b. 0.2%
c. 0.5%
d. 1.0%

Q1.3 [Short Answer] Explain why BNPL providers charge merchants a higher MDR (3–8%) than traditional card networks (1.5–3%). What does the merchant get in return? **Q1.4 [Short Answer]** Define Customer Acquisition Cost (CAC) and Lifetime Value (LTV). What LTV/CAC ratio is generally considered necessary for a viable fintech? **Q1.5 [Calculation]** A merchant processes $500,000 in monthly card sales. The MDR is 2.60%, broken down as: interchange 1.80%, network fee 0.13%, processor fee 0.25%, acquirer markup 0.30%, gateway fee 0.12%. Calculate: (a) total monthly fees, (b) the interchange share in dollars, (c) what the merchant actually receives.

Answers — Module 1

A1.1: (c) Interchange fee. It typically accounts for 60–70% of the total MDR.

A1.2: (b) 0.2%. Consumer credit interchange is capped at 0.3%.

A1.3: BNPL providers advance the full purchase amount to the merchant immediately (minus MDR) and bear the credit risk for 6 weeks while the consumer pays in installments. The merchant pays the higher MDR in exchange for: (1) higher conversion rates (consumers more likely to complete purchase), (2) larger basket sizes, and (3) immediate settlement. The higher MDR is justified if conversion uplift $\times$ basket size increase $>$ incremental MDR cost.

A1.4: CAC = Total Acquisition Spend / Number of New Customers. LTV = total net revenue a customer generates over the entire relationship. A viable fintech generally needs LTV/CAC $> 3.0\times$, meaning each customer must generate at least 3 times the cost of acquiring them.

A1.5: (a) Total fees: \$500,000 $\times$ 2.60% = \$13,000. (b) Interchange: \$500,000 $\times$ 1.80% = \$9,000 (69% of total fees). (c) Merchant receives: \$500,000 $-$ \$13,000 = \$487,000.

Module 2: The Access Problem — Questions

Q2.1 [Multiple Choice] What is the key difference between "unbanked" and "underbanked"?

a) Unbanked have accounts but do not use them; underbanked have no accounts
b) Unbanked have no formal account; underbanked have an account but rely on alternative financial services
c) Both terms mean the same thing
d) Underbanked refers only to businesses, not individuals

Q2.2 [Multiple Choice] Chouldechova's impossibility theorem shows that when base rates differ across groups, it is impossible to simultaneously satisfy:

a) Accuracy, precision, and recall
b) Equal FPR, equal FNR, and equal predictive values
c) Demographic parity and individual fairness
d) Privacy, accuracy, and fairness

Q2.3 [Short Answer] Explain how M-Pesa is a "leapfrog technology" for financial inclusion. What did it bypass? **Q2.4 [Short Answer]** Define demographic parity and equalized odds. Give one reason why a regulator might prefer one over the other. **Q2.5 [Analysis]** A bank uses an ML model for loan approvals. The approval rate is 70% for Group A and 50% for Group B. The default rate among approved borrowers is 5% for both groups. Does this model satisfy demographic parity? Does it satisfy equalized odds (specifically, equal FPR)? Explain.

Answers — Module 2

A2.1: (b) Unbanked = no account at a formal financial institution. Underbanked = has an account but still uses alternative services (payday lenders, check cashers) for some needs.

A2.2: (b) Equal false positive rates, equal false negative rates, and equal predictive values cannot all hold simultaneously when base rates differ.

A2.3: M-Pesa (launched in Kenya, 2007) enabled financial transactions via SMS on basic mobile phones. It leapfrogged the entire traditional banking infrastructure — no branches, no bank accounts, no internet required. It bypassed the need for physical banking infrastructure and formal identity documentation, reaching populations that banks never served.

A2.4: Demographic parity requires equal approval rates across groups: $P(\hat{Y} = 1|A = 0) = P(\hat{Y} = 1|A = 1)$. Equalized odds requires equal TPR and FPR across groups. A regulator focused on *equal opportunity* might prefer equalized odds because it conditions on actual outcome — equally qualified applicants from both groups should have the same chance of approval. A regulator focused on *equal representation* might prefer demographic parity.

A2.5: Demographic parity: **No** — approval rates differ (70% vs. 50%). Equalized odds (equal FPR): We cannot determine this from the given information alone. The 5% default rate among *approved* borrowers is the same, but equalized odds requires comparing rates among both approved *and rejected* applicants. We would need to know the default rate of rejected applicants (counterfactual).

Module 3: The Trust Problem — Questions

Q3.1 [Multiple Choice] A Merkle tree allows verification of a single transaction in:

a. $O(1)$ hash computations
b. $O(\log n)$ hash computations
c. $O(n)$ hash computations
d. $O(n^2)$ hash computations

Q3.2 [Multiple Choice] In a constant product AMM ($x \cdot y = k$), what happens to slippage as the trade size increases?

a. Slippage decreases
b. Slippage stays constant
c. Slippage increases
d. Slippage depends only on k, not trade size

Q3.3 [Short Answer] Explain the Byzantine Generals Problem and why the threshold $n \geq 3f + 1$ is necessary. What does f represent? **Q3.4 [Short Answer]** Compare fiat-backed stablecoins and algorithmic stablecoins. Which type is riskier and why? **Q3.5 [Calculation]** An AMM pool contains 100 ETH (x) and 200,000 USDC (y), so $k = 20{,}000{,}000$. A trader buys 5 ETH. (a) What is the initial price per ETH? (b) How much USDC must the trader pay? (c) What is the effective price per ETH for this trade?

A3.1: (b) $O(\log n)$ — only the sibling hashes along the path from the leaf to the root are needed (the inclusion proof).

A3.2: (c) Slippage increases. The constant product curve means larger trades move the price ratio more, resulting in a worse effective price.

A3.3: The Byzantine Generals Problem models a situation where distributed participants ("generals") must agree on a common action, but some may send conflicting messages (act maliciously). The threshold $n \geq 3f + 1$ means the system needs at least $3f + 1$ total nodes to tolerate f faulty/malicious nodes. This is necessary because each honest node needs enough independent confirmations to distinguish genuine messages from malicious ones.

A3.4: Fiat-backed stablecoins (e.g., USDC) are backed 1:1 by reserves in bank accounts. Risk: custodial and counterparty risk (is the reserve actually there?). Algorithmic stablecoins maintain their peg through supply/demand mechanisms without full collateral. They are **riskier** because the stabilization mechanism can fail under extreme market conditions (e.g., Terra/UST collapse in 2022 — a "death spiral" where the algorithm cannot restore the peg).

A3.5: (a) Initial price: $P = y/x = 200{,}000/100 = \$2{,}000$ per ETH. (b) After buying 5 ETH: $x' = 95$, so $y' = k/x' = 20{,}000{,}000/95 = 210{,}526.32$. Trader pays: $210{,}526.32 - 200{,}000 = \$10{,}526.32$. (c) Effective price: $\$10{,}526.32/5 = \$2{,}105.26$ per ETH (vs. \$2,000 initial — 5.3% slippage).

Module 4: The Risk Problem — Questions

Q4.1 [Multiple Choice] Expected Shortfall (ES) differs from VaR because ES:

a) Always equals VaR
b) Measures the average loss in the entire distribution
c) Measures the average loss in the tail beyond VaR
d) Is always smaller than VaR

Q4.2 [Multiple Choice] Delta (Δ) of a call option measures:

a) How much the option price changes per \$1 move in the underlying
b) How fast Delta itself is changing
c) The option's sensitivity to implied volatility
d) The daily time decay of the option

Q4.3 [Short Answer] Explain the three pillars of the Basel III framework. Which pillar was added in Basel II that did not exist in Basel I? **Q4.4 [Short Answer]** Why does cyber risk violate the assumptions underlying traditional VaR models? Name two specific assumptions that are violated. **Q4.5 [Calculation]** A portfolio has 1,000 daily returns sorted from worst to best. The 50 worst returns (in %) are: $-4.2, -3.8, -3.5, -3.3, -3.1, -2.9, -2.7, -2.5, \ldots$ (a) What is the 95% VaR? (b) If the average of the 50 worst returns is -2.8%, what is the 95% ES? (c) Explain in one sentence why ES is preferred by Basel III.

Answers — Module 4

A4.1: (c) ES measures the average loss in the tail beyond VaR. It is always $\geq$ VaR. Basel III requires ES at the 97.5% level.

A4.2: (a) Delta measures how much the option price changes per \$1 move in the underlying stock. It also approximates the probability of finishing in-the-money.

A4.3: The three pillars are: **Pillar 1** — Minimum capital requirements (risk-weighted assets, CET1 $\geq$ 4.5%). **Pillar 2** — Supervisory review (regulators assess bank-specific risks beyond Pillar 1). **Pillar 3** — Market discipline (public disclosure of risk and capital information). Basel I had only a single capital standard. **Pillar 2 (supervisory review) and Pillar 3 (market discipline)** were added in Basel II.

A4.4: Traditional VaR assumes: (1) **stationarity** — risk factor distributions do not change over time. Cyber attacks are discrete, unpredictable events that violate this. (2) **Historical representativeness** — past returns predict future tail risk. Novel cyber threats have no historical precedent. Additionally, VaR assumes you can sell at current prices, but a major cyber incident can freeze liquidity entirely.

A4.5: (a) 95% VaR = the 50th worst return (5% of 1,000) = -2.5% (the boundary of the worst 50). More precisely, the 50th return from the bottom. (b) 95% ES = average of the 50 worst returns = -2.8%. (c) ES is preferred because it captures the *severity* of tail losses, not just the threshold — it tells you what happens when things go wrong, not just the boundary.

Module 5: The Automation Problem — Questions

Q5.1 [Multiple Choice] A model with high training accuracy but low test accuracy is exhibiting:

a) Underfitting
b) Overfitting
c) Data drift
d) Concept drift

Q5.2 [Multiple Choice] Retrieval-Augmented Generation (RAG) addresses which LLM limitation?

a) Slow inference speed
b) Hallucination of factually incorrect information
c) High training cost
d) Lack of multilingual support

Q5.3 [Short Answer] Explain the bias-variance trade-off. A credit scoring model approves almost everyone (low threshold). Is this model more likely underfitting or overfitting? Why? **Q5.4 [Short Answer]** Define concept drift, data drift, and covariate shift. Give one financial example of concept drift. **Q5.5 [Analysis]** A fraud detection model was trained on 2019–2022 transaction data and deployed in 2023. Performance drops sharply in Q2 2023. Identify three possible causes of the performance degradation, linking each to a specific ML concept from Module 5.

A5.1: (b) Overfitting. The model has memorized the training data (high training accuracy) but fails to generalize to unseen data (low test accuracy).

A5.2: (b) Hallucination. RAG retrieves relevant documents from a curated corpus and provides them as context to the LLM, grounding its responses in factual sources rather than relying solely on parametric memory.

A5.3: The bias-variance trade-off: simple models have high bias (systematic errors, underfitting) and low variance; complex models have low bias but high variance (overfitting). A model that approves almost everyone is likely **underfitting** — it has not learned meaningful patterns to distinguish good from bad credit risks. It has high bias (always predicts "approve") and low variance.

A5.4: Concept drift: $P(Y|X)$ changes — the relationship between features and target shifts. Data drift: $P(X)$ changes — the distribution of input features shifts. Covariate shift: $P(X)$ changes but $P(Y|X)$ stays the same. Financial example of concept drift: during COVID-19, previously reliable borrowers defaulted at much higher rates — the relationship between income/employment features and default probability fundamentally changed.

A5.5: (1) **Concept drift:** Fraudster tactics evolved (new attack patterns not in training data). (2) **Data drift:** Transaction patterns shifted (e.g., new payment channels, post-pandemic spending changes). (3) **Regime change / non-stationarity:** A structural break (new regulation, market event) changed the data-generating process, making the 2019–2022 training data non-representative of 2023 conditions.

Module 6: The Infrastructure Problem — Questions

Q6.1 [Multiple Choice] In a Real-Time Gross Settlement (RTGS) system, transactions are:

a) Batched and netted at end of day
b) Settled individually in real time
c) Queued until both parties confirm
d) Settled only during business hours

Q6.2 [Multiple Choice] SWIFT is best described as:

a) A settlement system that moves money between banks
b) A messaging network that sends payment instructions between banks
c) A central bank payment rail like Fedwire
d) A cryptocurrency network for cross-border transfers

Q6.3 [Short Answer] Compare RTGS and DNS in terms of: (a) counterparty risk, (b) liquidity requirements, (c) settlement speed. Which is used for large-value interbank payments and why? **Q6.4 [Short Answer]** What is Banking-as-a-Service (BaaS)? Give an example of how a non-bank company might use BaaS to offer financial products. **Q6.5 [Analysis]** A bank is considering migrating its core banking system from an on-premise COBOL mainframe to a cloud-native platform. Identify three benefits and three risks of this migration, referencing concepts from Module 6.

Answers — Module 6

A6.1: (b) RTGS settles each transaction individually in real time. This eliminates counterparty risk but requires high liquidity.

A6.2: (b) SWIFT is a messaging network — it transmits payment instructions but does not settle or move money itself. Settlement happens through correspondent banking or payment systems like TARGET2.

A6.3: (a) RTGS: zero counterparty risk (immediate finality). DNS: settlement risk during the netting window (if a participant defaults before settlement). (b) RTGS: high liquidity requirements (full amount per transaction). DNS: lower (only net positions settle). (c) RTGS: minutes. DNS: hours to end-of-day. RTGS is used for large-value interbank payments (e.g., Fedwire, TARGET2) because the cost of counterparty risk on large transactions outweighs the liquidity cost.

A6.4: BaaS is a model where licensed banks provide banking infrastructure (accounts, cards, payments, lending) to non-bank companies via APIs. Example: An e-commerce platform (like Shopify) partners with a BaaS provider to offer merchant bank accounts, instant payouts, and business loans — without becoming a bank itself. The BaaS provider holds the banking license; the e-commerce platform provides the customer relationship.

A6.5: Benefits: (1) Scalability and elasticity (cloud scales with demand). (2) Faster product development (API-first, microservices). (3) Lower long-term maintenance cost (no COBOL developer shortage). Risks: (1) Regulatory scrutiny (data sovereignty, cloud outsourcing requires supervisory approval). (2) Migration risk (data integrity during cutover, dual-running cost). (3) Vendor lock-in (dependence on a single cloud provider for critical infrastructure).

Module 7: The Compliance Problem — Questions

Q7.1 [Multiple Choice] Under MiCA, a stablecoin pegged to a single fiat currency (e.g., USDC) is classified as:

a) Asset-Referenced Token (ART)
b) E-Money Token (EMT)
c) Utility Token
d) Security Token

Q7.2 [Multiple Choice] SR 11-7 defines model risk as:

a) Risk that a model is too complex to validate
b) Potential for adverse consequences from incorrect or misused model outputs
c) Risk that a model takes too long to run in production
d) Risk that competitors develop better models

Q7.3 [Short Answer] Describe the three stages of money laundering (placement, layering, integration) and give a brief example of each. **Q7.4 [Short Answer]** Explain the "three lines of defense" model in the context of model risk governance. What is the role of each line? **Q7.5 [Analysis]** A fintech uses an AI model for automated AML transaction monitoring. The EU AI Act classifies this as "high risk." What three obligations does the EU AI Act impose on high-risk AI systems? How might these obligations conflict with the need for fast, scalable AML screening?

Answers — Module 7

A7.1: (b) E-Money Token (EMT). Under MiCA, tokens referencing a single fiat currency are classified as EMTs and regulated similarly to e-money. Tokens referencing multiple assets are ARTs.

A7.2: (b) SR 11-7 defines model risk as "the potential for adverse consequences from decisions based on incorrect or misused model outputs and reports." Note that *misuse* of correct outputs is also model risk.

A7.3: (1) **Placement:** Introducing illicit cash into the financial system (e.g., depositing cash in small amounts below reporting thresholds — "structuring"). (2) **Layering:** Moving money through complex transactions to obscure its origin (e.g., multiple wire transfers between shell companies in different jurisdictions). (3) **Integration:** Reinvesting the laundered money into legitimate assets (e.g., purchasing real estate or a business).

A7.4: 1st Line: Business units (model owners/developers) — responsible for building and using models correctly. 2nd Line: Risk management/compliance (independent validation team) — responsible for challenging and validating models independently. 3rd Line: Internal audit — provides independent assurance that the first two lines are functioning effectively. Each line must be independent of the others.

A7.5: The EU AI Act requires high-risk systems to provide: (1) human oversight (human-in-the-loop), (2) transparency and explainability (users must understand how decisions are made), (3) data quality and bias testing. Conflicts: human oversight slows automated screening (millions of transactions/day cannot each have human review). Explainability may require simpler models, potentially reducing detection accuracy. Data quality audits add cost and delay to model updates needed to catch evolving threats.

Module 8: The Future — Questions

Q8.1 [Multiple Choice] Shor's algorithm threatens modern cryptography because it can:

a) Break symmetric ciphers like AES-256
b) Factor large integers in polynomial time, breaking RSA and ECC
c) Generate random numbers faster than classical computers
d) Decrypt any encrypted message without a key

Q8.2 [Multiple Choice] Self-Sovereign Identity (SSI) differs from federated identity because:

a) SSI requires a central authority; federated does not
b) SSI credentials are controlled by the individual; federated relies on a trusted third party
c) SSI is less secure than federated identity
d) SSI requires biometric authentication; federated does not

Q8.3 [Short Answer] Explain the difference between compliance carbon markets (cap-and-trade) and voluntary carbon markets. Who participates in each? **Q8.4 [Short Answer]** Describe the four-scenario framework from Lesson 8.4

(FinTech Supernova, Regulated Renaissance, Slow Burn, Digital Fortress). What two axes define the scenarios? **Q8.5**

[Analysis] A bank's CISO must prepare a post-quantum migration plan. Outline the four phases of migration from classical to post-quantum cryptography. For each phase, identify one specific action the bank should take.

Answers — Module 8

A8.1: (b) Shor's algorithm factors large integers in polynomial time, breaking RSA and elliptic-curve cryptography (ECC). It does *not* break symmetric ciphers (Grover's algorithm weakens those, but only quadratically).

A8.2: (b) In SSI, the individual holds their credentials in a digital wallet and selectively discloses attributes — no central authority is contacted during verification. In federated identity (e.g., "Sign in with Google"), a trusted third party vouches for the user.

A8.3: Compliance markets (cap-and-trade): Regulators set an emissions cap; covered entities (e.g., power plants, heavy industry) must hold allowances for each tonne emitted. They can buy/sell allowances on a regulated exchange. The EU Emissions Trading System (EU ETS) is the largest. Voluntary markets: Companies and individuals buy carbon credits voluntarily to offset emissions. Not regulated, credits issued by independent standards (e.g., Verra, Gold Standard). Participants include corporates pursuing net-zero pledges.

A8.4: Two axes: (1) Regulation (permissive ↔ restrictive), (2) Technology pace (rapid ↔ gradual). FinTech Supernova (permissive + rapid): AI-native banks, high innovation, systemic fragility. Regulated Renaissance (restrictive + rapid): strong regulation, CBDCs dominate, innovation within guardrails. Slow Burn (permissive + gradual): incremental change, banks still run COBOL. Digital Fortress (restrictive + gradual): regulation stifles innovation, compliance costs reach 15% of revenue.

A8.5: Phase 1 — **Inventory:** Audit all cryptographic assets (certificates, keys, protocols). Phase 2 — **Plan:** Prioritize systems by risk (public-facing TLS first, internal systems later). Phase 3 — **Test:** Deploy hybrid certificates (classical + PQC) in test environments. Phase 4 — **Deploy:** Migrate production systems to NIST-approved PQC algorithms (e.g., CRYSTALS-Kyber, CRYSTALS-Dilithium).