

Digital Finance — Practice Bank Solutions
96 Practice Questions with Model Answers — INSTRUCTOR ONLY

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

PRACTICE ONLY — NOT GRADED

- **This course has NO final written examination.** Graded assessment (per syllabus) is Project 50% + Presentation 20% + Quizzes 30%.
- **Purpose:** This instructor-only solutions deck mirrors the student-facing `exam_bank.pdf`. All 96 items are in MC, short-answer, or calculation format.
- **Distribution:** Internal reference only.
- **No grading:** Nothing in this bank counts toward the student's final grade — it exists purely as self-review scaffolding.

practice; no final exam. Instructor copy — model answers wrapped in `\ifsolutions` (revealed when `\AtBeginDocument{\solutionstrue}`).

- **96 model answers**, one per question in `exam_bank.pdf`, organised by module.
- **Three tiers:**
 - **Basic (QB):** MC — correct letter + 1-line rationale + brief misconception note for each distractor.
 - **Applied (QP):** MC, calc — for calc, full setup → formula → substitution → answer with units → interpretation; for MC, correct letter with rationale.
 - **Advanced (QA):** MC + Short — for short-answer, model 2–3 sentence response naming the specific mechanism / event the question expects.
- **Toggle:** Set `\solutionsfalse` to compile a question-only version for distribution.

key style follows `SHARED_ANSWER_KEY_TEMPLATE.md` (Plan 6).

Answer

QB1.1: (b) An entity facilitating payer-payee settlement (correspondent bank, card network, clearinghouse). (a) is too narrow; (c) confuses regulator with intermediary; (d) is a vendor, not an intermediary.

QB1.2: (a) ACH batch cycles vs. Bitcoin peer-to-peer + 10-min block. (b) both are encrypted; (c) is irrelevant; (d) wrong — ACH is electronic.

QB1.3: (b) Acquirer pays issuer, passed to merchant via MDR. (a) reverses the flow; (c) interchange is set by the network, not regulators; (d) is a separate fee category.

QB1.4: (b) Processors do messaging/clearing for merchants; networks set rules and connect issuers and acquirers.

index: (b) of QB1.3 is the most-confused; many students invert the cash flow.

Misc

QP1.1 [Calc]: (i) FX markup = $\text{EUR } 100 \times 2.5\% = \text{EUR } 2.50$; (ii) interchange = $\text{EUR } 100 \times 1.5\% = \text{EUR } 1.50$; (iii) total cardholder-borne cost = $\text{EUR } 2.50$ (FX is on the cardholder); merchant settlement deduction = $\text{EUR } 1.50$ interchange (the merchant pays the MDR, of which interchange is the dominant share). Cardholder pays $\text{EUR } 102.50$ net; merchant receives $\text{EUR } 98.50$.

QP1.2: (b) Neobanks earn interchange + subscription + FX (low NIM). (a) is the traditional bank model.

QP1.3: (a) Two-sided platform, chargeback risk, instant authorisation; ACH is a utility batch rail.

QP1.4: (b) Merchants and low-income payees benefit; card networks and short-term lenders lose float revenue.

must show units throughout. The cardholder/merchant split is the load-bearing distinction. QP1.

QA1.1 [Short]: *Counter-argument:* Regulation, AML, KYC, consumer-protection, deposit insurance, and dispute-resolution remain bound to licensed intermediaries; even if rails go on-chain, the legal layer (and therefore the intermediary function) survives. ACH cost \$0.25 per transaction for decades; the bottleneck is not technology but distribution and trust.

QA1.2: (b) Mobile-money corridors and licensed stablecoin rails routinely deliver $\leq 3\%$ end-to-end (Wave-Stellar, Strike, MoneyGram-USDC, M-Pesa-GCash).

QA1.3: (c) Late fees, basket inflation, and stacking on revolving credit are the empirical drivers of higher consumer cost (CFPB report 2022; UK FCA review 2023).

QA1.4 [Calc]: (a) Card MDR: $\$100 \times 2.5\% = \2.50 ; (b) ACH flat: $\$0.25$; (c) USDC on Polygon: $\$0.01 + \$100 \times 0.1\% = \$0.11$. Ranking cheapest→most expensive: USDC ($\$0.11$) < ACH ($\0.25) < Card ($\$2.50$). *Regulatory catch on USDC:* on/off-ramp providers must hold MTL or e-money licences and apply KYC/AML/Travel Rule, which adds friction not captured in the per-transaction fee.

student error on QA1.4: forgetting that the on-ramp fee is the regulatory chokepoint, not the on-chain fee.

QB2.1: (b) Lacking access to (or priced out of) basic financial services. (a) is too narrow (focuses on credit cards alone); (c) conflates choice with exclusion; (d) is the opposite of exclusion.

QB2.2: (b) Thin-file populations are systematically denied; historical data encodes prior unequal access.

QB2.3: (b) Systematic error disadvantaging a protected group from biased data, labels, or features.

QB2.4: (b) SMS-based airtime balances cashed in/out via human-agent network; no bank account.

M-

Pesa is the canonical leapfrog case; do not confuse with mobile banking apps.

QP2.1: (b) ZIP code is highly correlated with race in segregated US cities; disparate impact actionable under ECOA + EU AI Act (Art. 10 high-risk system data-governance duties). Lesson 2.2 (Credit Scoring) covers proxy discrimination.

QP2.2: (b) Alternative data widens access but raises consent, drift, and proxy-discrimination risks; FICO is narrower but better regulated.

QP2.3 [Calc]: Ratio = $\$3,300 / \$1 = 3,300\times$ reduction. As a percent, the new minimum is $1/3,300 = 0.030\%$ of the original; equivalently, the minimum dropped 99.97%.

QP2.4: (b) Low minimums + digital onboarding + automated rebalancing reduce time and friction barriers — the binding constraint for first-time investors, not just fee level.

1 verified the lesson citation: proxy discrimination is in L2.2, not L2.4.

QA2.1 [Short]: *When true:* Easy mobile credit + behavioural-design dark patterns combined with no affordability check (e.g., Kenya M-Shwari over-borrowing, UK BNPL revolving stack). *Mitigating lever:* Affordability rules + APR cap + cooling-off windows + open-banking-driven liability for credit decisions.

QA2.2: (b) Rental, utility, cash-flow data — causal link to repayment, contestable. (a)/(c)/(d) are proxy-discrimination minefields.

QA2.3 [Short]: *Net exploitative if* the operator monetises order flow or volume (PFOF → payment for order flow conflict, Robinhood SEC \$65M settlement Dec 2020); *net democratising if* fractional shares + zero commission + plain-language risk warnings. The Alex Kearns case (June 2020) is a concrete harm: confusing UI on options assignment.

QA2.4: (b) Expanded access via TPP licences but raises consent-fatigue, phishing-vector, and big-tech-data-monopoly risks (e.g., Apple Pay aggregation; FCA scam-loss reports).

The

Alex Kearns and DWS-greenwashing cases are recurring across the bank — they ground abstract claims.

QB3.1: (b) Without a central ledger, a digital coin can be spent twice; Bitcoin globally orders transactions via PoW + longest-chain rule.

QB3.2: (b) PoW = computation-as-sybil-cost (energy); PoS = stake-as-sybil-cost (capital, with slashing). Different attack surfaces.

QB3.3: (b) Code on-chain; deterministic execution; Uniswap AMM example.

QB3.4: (b) No legal enforcement of uncollateralised claims + volatile collateral \Rightarrow excess collateral provides liquidation cushion.

Note:

“DeFi has no enforcement” is the load-bearing reason for over-collateralisation, not regulation.

QP3.1 [Calc]: Liquidation when $\frac{\text{Collateral Value}}{\text{Debt}} = 150\%$. With debt = \$20,000, liquidation collateral value = $\$20,000 \times 1.50 = \$30,000$. Per-ETH liquidation price = $\$30,000 / 10 \text{ ETH} = \mathbf{\$3,000/ETH}$. Initial collateral was \$30,000, so the position is liquidated immediately at deposit if the user borrows the full \$20,000 (the example illustrates that with a 150% threshold, the safe borrow is much less than collateral value / 1.50). Margin call at 175% triggers earlier at \$3,500/ETH per the same formula. *Pedagogical note: in practice, protocols set the threshold at 110–125% with margin-call buffers higher; the 150% number used here is the textbook teaching value.*

QP3.2: (b) Coinbase = custodial + KYC; Uniswap = non-custodial + no protocol-level KYC (front-ends may geoblock).

QP3.3: (b) Trilemma: any design typically sacrifices one of decentralisation, security, scalability; Bitcoin chose decentralisation + security \Rightarrow low TPS; L2 (Lightning, rollups) addresses scale off-chain.

QP3.4: (b) Reflexive death spiral; no exogenous collateral floor; LUNA mint mechanism amplified the depeg.

QP3.

is fully gradeable now (Plan 1): liquidation threshold = 150%, margin call = 175%, so the answer is unique.

QA3.1 [Short]: *Code \neq law when* (i) oracles supply ambiguous or manipulated input (Mango Markets Oct 2022, \$117M oracle-manipulation theft prosecuted as commodities fraud); (ii) courts override on-chain state (Tornado Cash OFAC sanctions, August 2022); or (iii) DAO members exit via fork (DAO 2016).

QA3.2: (b) Quorum + timelock + quadratic/reputation voting + dispute channel (Kleros) + off-chain fallback.

QA3.3 [Short]: *Both partly true.* Innovation: 24/7 settlement, programmable money, lower remittance cost. Banking-risk replication: USDC depeg to \$0.87 during SVB collapse (Mar 2023); Tether \$41M CFTC settlement (Oct 2021) for misleading reserve disclosure — structurally similar to bank-run dynamics.

QA3.4: (b) Social override of code-is-law; preserved funds; set debated precedent (Ethereum Classic remains the un-forked chain).

Cite

specific events on the Advanced short-answer items: USDC-SVB, Tether-CFTC, Mango Markets, Tornado Cash sanctions.

QB4.1: (b) Conditional distributional claim: 1-day loss $>$ \$1M on $\sim 5\%$ of days; says nothing about magnitude beyond.

QB4.2: (b) Right (not obligation) to buy at strike; XYZ \$50 call exercised if $XYZ > \$50$.

QB4.3: (b) Counterparty fails to deliver; mitigated by central clearing, margin; materialised in OTC (Archegos 2021).

QB4.4: (b) Market, credit, operational — the Basel three.

distractor (a) is the most-common student misread (probability vs. frequency of exceedance).

QB4.

QP4.1 [Calc]: $\text{VaR}_{0.95,1d} = (z\sigma - \mu) \cdot V = (1.645 \times 0.012 - 0.0005) \times \$10,000,000 =$

$(0.01974 - 0.0005) \times \$10,000,000 = 0.01924 \times \$10,000,000 = \boxed{\$192,400}$. *Why Gaussian under-estimates:* financial returns have fat tails (kurtosis > 3); historical data shows many “5-sigma” moves; Plan 1 fixed M4L1 to show the corrected $\sim 3.5\text{M}$ -year frequency under Gaussian. Use Cornish-Fisher, EVT, or Expected Shortfall for tail-aware metrics.

QP4.2: (b) TRS gives prime broker legal ownership \rightarrow Archegos avoided 13F triggers; correlated unwind by all brokers caused the fire sale.

QP4.3 [Calc]: (i) Intrinsic at expiry $= \max(S - K, 0) = \max(\$56 - \$50, 0) = \6 . (ii) Long P&L per share $= \text{intrinsic} - \text{premium} = \$6 - \$3 = +\3 . (iii) Breakeven $= \text{strike} + \text{premium} = \$50 + \$3 = \53 .

QP4.4: (b) DeFi yield farm has the thickest tail (smart-contract + oracle + depeg + protocol risk); savings account thinnest (deposit insurance); stablecoin between (USDC 2023, UST 2022).

QP4.

must show units (\$192,400 for a \$10M book, 1 day, 95%); under Cornish-Fisher / ES the number is materially higher.

QA4.1 [Short]: *Critique has merit:* VaR is silent on the conditional tail $E[L|L > VaR_\alpha]$; a portfolio with bounded VaR can have unbounded ES. *Alternative:* Expected Shortfall (CVaR), now mandated by Basel FRTB; or stress-test suite anchored on historical episodes (2008, March 2020, March 2023 SVB).

QA4.2: (b) Multi-layer limits: gross/net notional, stress-loss ceilings, concentration per counterparty + exchange, liquidity-bucket caps, automated circuit breakers.

QA4.3 [Short]: *Specific gap (SVB):* Held-To-Maturity (HTM) accounting masked unrealised loss on long-duration Treasuries; uninsured-deposit concentration (>90% above FDIC limit) made the bank vulnerable to a coordinated run; social-media-accelerated withdrawals (~\$42B/day at peak) outpaced any classical liquidity ratio.

QA4.4: (b) Transformed not eliminated: smart-contract bugs (e.g., bZx, Yam), oracle manipulation (Mango), governance attacks (Beanstalk), bridge hacks (Wormhole \$320M, Ronin \$625M).

hack losses 2022–2023 exceed \$2B cumulatively; this is the canonical example of “DeFi transforms risk.”

Bridg

QB5.1: (b) Fits noise; high training accuracy, poor OOS; backtest does not survive live.

QB5.2: (b) Supervised = labelled (x, y) (credit default); unsupervised = no labels (clustering, anomaly).

QB5.3: (b) Self-attention removed RNN sequential bottleneck; pre-training on huge corpora; reusable embeddings (BERT, GPT) transfer.

QB5.4: (b) Hallucination, prompt injection, drift/auditability — the three concrete categories cited in the EU AI Act, NIST AI RMF, and FSB GenAI 2024 reports.

(a) on QB5.4 is too narrow (those are infra concerns, not the load-bearing risks).

Distr

QP5.1 [Calc]: 100 frauds \times 80% recall = 80 TP, 20 FN. Precision = $TP / (TP + FP) = 80\% \Rightarrow FP = TP \cdot (1/0.80 - 1) = 80 \times 0.25 = 20$ FP. Sanity check: precision = $80 / (80 + 20) = 80\%$. Out of 10,000 transactions: 80 TP + 20 FN + 20 FP + 9,880 TN = 10,000.

QP5.2: (b) Stable regime + interpretability requirement \rightarrow rules win; non-stationary regime + labels available \rightarrow ML wins.

QP5.3: (b) GDPR Art. 22 + EU AI Act high-risk obligations; rejection must be explainable; automated decisions need human-oversight option.

QP5.4: (b) Hallucination is distributional; mitigations are RAG grounding, numeric verification, “don’t know” tuning. Not a code bug.

confusion-matrix arithmetic must close to 10,000; otherwise the student has confused recall and precision.

QP5.

QA5.1 [Short]: *Most automatable:* research summarisation, data cleansing, draft compliance attestations, first-pass equity-research note generation. *Least automatable:* client-relationship management, deal negotiation, judgement calls under genuine uncertainty (M&A defence, distressed-debt restructuring), regulator-facing escalations. Goldman Sachs / JPM analyst-task augmentation surveys 2024 show 30–40% time-savings, not headcount cuts.

QA5.2: (b) Versioned data lineage + reproducible training + hold-out validation + fairness metrics + canary + drift monitor + retraining triggers + incident playbook.

QA5.3 [Short]: *Knight Capital:* A deprecated SMARS “Power Peg” code path was reactivated by an erroneous deploy script on only 7 of 8 servers; orders fired with no risk gate. *Safeguard:* Pre-deploy code-review + automated kill switch on cumulative loss + position-limit enforcement at the broker-dealer (15c3-5).

QA5.4: (b) Non-trivial but proportionate; Annex III credit-scoring obligations phase in; CE marking + post-market monitoring.

Capital loss timeline (45 minutes, \$440M) is the canonical case for trading kill switches.

Knigh

QB6.1: (b) Messaging network (MT, ISO-20022); settlement is via correspondent banking or RTGS rails.

QB6.2: (b) RTGS = real-time gross in central-bank money (Fedwire, TARGET2); net = batched + offset (CHIPS, ACH, EBA STEP2). Tradeoff: liquidity vs. inter-cycle risk.

QB6.3: (b) Regulatory mandate (PSD2 EU, OBIE UK); APIs to TPPs; PSD3/PSR extending.

QB6.4: (b) Ledger (book of record), product engine, integration layer (ISO-20022, KYC, reporting).

confusion: SWIFT is messaging, not settlement — distractor (a) catches this.

Com

QP6.1 [Calc]: (i) 99.99% uptime \Rightarrow 0.01% downtime = $0.0001 \times 365 \text{ days} \times 24 \text{ hours} \times 60 \text{ minutes} = 52.56$ minutes/year. (ii) Failed transactions = $0.0001 \times 1,000,000 \text{ transactions/day} \times 365 \text{ days} = 36,500$ failed transactions/year (assuming uniform flow + 100% loss during outage). *Caveat:* Real failures cluster (long outages cost more than the average); SLA reporting often excludes “planned” downtime, distorting the headline number.

QP6.2: (b) SWIFT 1–3 days probabilistic; Ripple seconds + cryptographic finality, but minor-corridor liquidity uneven.

QP6.3: (b) BaaS rental: BaaS partner holds licence + ledger; neobank owns CX + marketing. Solaris, ClearBank, Synapse (RIP, Apr 2024 collapse).

QP6.4: (b) Single-vendor concentration: a routine endpoint-security update bricked 8.5M Windows machines; not a breach but a supply-chain monoculture failure.

collapse (Apr 2024) is the canonical case of BaaS failure; cite alongside QP6.3.

Synap

QA6.1 [Short]: *True:* multi-AZ + auto-failover + cross-region replication remove most physical single-points-of-failure. *False:* control-plane (us-east-1 IAM event 2017, Dec 2021), provider-global account-management failures, and routing-misconfiguration-induced cascade (Cloudflare 2022) remain new-shaped concentration risks.

QA6.2: (b) Full stack: core-banking + payments + KYC/AML + card-issuing + fraud + front-end + SRE + reg-reporting + DR.

QA6.3 [Short]: *Capability gap:* programmable money (conditional payments, smart-contract automation), offline payment (smart-card-based Sand Dollar Bahamas), bearer-like privacy (cash-equivalent), central-bank-issued finality (no commercial-bank credit risk). FedNow / SEPA-Instant don't offer any of these.

QA6.4: (b) Live systemic concern; DORA EU (Jan 2025), FSB Critical Third-Party regime, US FFIEC Critical Service Provider concept; mandatory multi-cloud / regulated exit plans / op-resilience tests.

DOR

(in force Jan 17, 2025) introduces direct supervision of CTPPs — AWS, Azure, GCP all in scope.

QB7.1: (b) Customer Identification + Customer Due Diligence + ongoing monitoring. The three pillars match FATF Recommendation 10.

QB7.2: (b) Rule-based fires on static thresholds (high false-positive); risk-based scores per customer behaviour and typology (lower-but-better-quality alerts).

QB7.3: (b) Filed when activity reasonably suggests proceeds of crime, terrorism financing, evasion; FinCEN US 30-day rule; continuing-SAR for ongoing activity.

QB7.4: (b) NLP on regs + rules-engines on onboarding + ML on transactions + graph analytics on UBO + workflow on SAR/STR pipelines.

Recommendation 10 is the international standard underpinning the three KYC pillars.

FATF

QP7.1 [Calc]: (i) FP rate of flagged alerts = $495 / 500 = 99.0\%$. (ii) True-positive recall: confusion matrix on 1,000,000 transactions/day. True positives = 5 (5 of 500 flagged are genuine); False negatives = 2 (cases system misses); Recall = $TP/(TP+FN) = 5/(5+2) = 71.4\%$. *Interpretation:* Even with seemingly catchy headlines, 99% FP rate plus 71% recall means the system both burdens analysts (495 false alerts/day) AND misses real cases — a typical legacy AML problem driving the ML/risk-based shift.

QP7.2: (b) Bank: more statutory duties + more actual safeguards. DeFi: fewer statutory duties + fewer safeguards; user bears risk directly.

QP7.3: (b) EMT (one-fiat-pegged), ART (basket-backed), Other (utility/speculative); each tier with prospectus, capital, reserve duties.

QP7.4: (b) Inadequate CDD on non-resident portfolio + weak monitoring thresholds + ignored whistleblowers + slow regulator cooperation. All Module 7 failure modes.

Dans

\$2B-plus settlement (Dec 2022 DOJ + Estonian / Danish fines) is the largest EU AML penalty to date.

QA7.1 [Short]: *Tension:* pseudonymous-by-design vs. VASP-must-share-originator-and-beneficiary identity. *Mitigation:* Selective disclosure with zero-knowledge proofs (e.g., zk-SNARK attestation that originator passed KYC without revealing identity to counterparty); TRP (Travel Rule Protocol) interop standards; OpenVASP.

QA7.2: (b) Independent validation, model card, lineage, fairness + drift monitors, retraining triggers, human-override workflow, incident playbook, periodic audit.

QA7.3 [Short]: *Works:* ERC-3643 transfer-restriction tokens enforce whitelist-only transfers automatically; tokenised securities under SDX use code-enforced settlement. *Fails:* AML typology detection (structuring, layering, smurfing) requires context-dependent judgement; smart contracts can detect threshold patterns but not the narrative behind a transaction.

QA7.4: (b) EU MiCA = ex-ante clarity (predictable but potentially over-inclusive); US enforcement = case-by-case (target-specific harms but chilling uncertainty). Neither is dominant.

3643 is the canonical permissioned-token standard, used by Polymath and Tokeny.

ERC-

QB8.1: (b) Verifiable credentials in user-held wallet; W3C DID/VC standards; minimum-disclosure proofs.

QB8.2: (b) Shor's algorithm breaks RSA + ECC on FT quantum hardware; symmetric crypto only \sqrt{N} speedup (Grover); double key length for AES.

QB8.3: (b) Green bonds, sustainability-linked loans, carbon credits, blended finance.

QB8.4: (b) Tokenised use-of-proceeds + on-chain disbursement + MRV oracles + third-party attestation + public dashboards.

NIST

PQC standards finalised Aug 2024: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA).

QP8.1: (b) ML-KEM (Kyber, FIPS 203), ML-DSA (Dilithium, FIPS 204), SLH-DSA (SPHINCS+, FIPS 205); RSA / ECC vulnerable to FT quantum.

QP8.2: (b) Estonia: strong legal recognition + centralised state trust; SSI: distributed trust + multi-issuer attestations + young regulatory recognition.

QP8.3 [Calc]: (i) Annual cost = 10,000 tonnes \times \$50/tonne = \$500,000/year. (ii) Pre-markup cost per tonne = \$50/1.30 = \$38.46; new cost with 5% markup = \$38.46 \times 1.05 = \$40.38; new annual cost = \$403,800. *Saving* = \$500,000 – \$403,800 = \$96,200/year (about 19%). *Caveat:* Tokenisation reduces intermediary markup but adds gas + minting + verification overhead, and only works if buyers trust the on-chain attestation.

QP8.4: (b) No segregation between customer funds and Alameda; no prudential regulation; opaque related-party loans; offshore jurisdiction arbitrage.

FTX SBF prosecution (verdict Nov 2023, 25-year sentence Mar 2024) is the canonical post-crash case.

The

QA8.1 [Short]: *Pessimistic:* Harvest-now-decrypt-later on dormant P2PK UTXOs once $\sim 2,000$ logical qubits are achievable. *Optimistic:* IBM Heron 2024 (5 GHz, error rates $\sim 10^{-3}$) and Google Willow Dec 2024 (error-corrected at scale) indicate progress, but FT at scale ($\geq 10,000+$ logical qubits) is plausibly \geq a decade away; Bitcoin can soft-fork to PQC signatures (Taproot-like upgrade path). Net: not 10 years, more likely 15–25, but crypto-agility planning should start now.

QA8.2: (b) DIDs + VCs + selective-disclosure (zk-proofs) + interoperable W3C / EBSI + regulated issuer + verifier lists + revocation infrastructure + cross-jurisdiction recognition.

QA8.3 [Short]: *Real change:* TCFD-driven divestment (e.g., NZAM net-zero pledges \rightarrow portfolio-emissions reporting); Science-Based Targets initiative (SBTi) committed companies cut emissions $\sim 8\%$ /year on average. *Data theatre:* DWS \$19M SEC settlement (Sep 2023) for ESG-overstatement; EU Taxonomy gas + nuclear inclusion (2022) showed political compromise over scientific consistency.

QA8.4: (b) Cover policy objective + design choices + disintermediation risk + cyber + legal-tender + phased pilot path (e-krona, digital euro, Helvetia III, mBridge).

III (SNB + BIS, 2023–2024) is the canonical wholesale-CBDC-on-DLT pilot; mBridge is the canonical multi-jurisdiction one.

QB7A.1 [Day 7A, Basic]: (c) Policing-and-enforcement costs.

Rationale. Coase (1937, “The Nature of the Firm”) named four canonical transaction-cost categories: search, bargaining, contracting, and policing-and-enforcement. Stripe’s economic value is concentrated in the fourth category. The API absorbs chargeback handling, fraud screening, dispute mediation, settlement guarantees, and PCI-DSS compliance overhead that no individual small merchant could economically absorb on its own.

Why the distractors fail. (a) Stripe does not match payer with payee identities; users find merchants through other channels. (b) Stripe’s pricing is fixed, not negotiated per transaction. (d) Cash transport is unrelated to a card-rails platform.

Common mistake. Students conflate “Stripe is a payments processor” with “Stripe reduces all transaction costs uniformly.” The Coasean lens is sharper: identify the *specific* TC category and the firm’s strategic justification becomes legible.

transaction-cost taxonomy: search, bargaining, contracting, policing-and-enforcement. Each fintech belongs to one or two categories; identify which to understand why the firm exists.

QP7A.2 [Day 7A, Intermediate, Calc]: Loss-aversion calculation under prospect theory.

Setup. $\lambda = 2.25$ (Tversky and Kahneman behavioral estimate). Utility model: $V = v(\text{gain}) - \lambda \cdot v(\text{loss})$ with linear $v(x) = x$.

Frame 1 (full \$120 weighted as loss):

$$V_1 = v(30) - 2.25 \cdot v(120) = 30 - 2.25 \cdot 120 = 30 - 270 = -240.$$

Frame 2 (only first \$30 weighted as the salient loss under behavioral re-framing):

$$V_2 = v(30) - 2.25 \cdot v(30) = 30 - 67.5 = -37.5.$$

Reframing gain: $V_2 - V_1 = -37.5 - (-240) = +202.5$ dollars of perceived utility.

Behavioral interpretation. The reframing changes nothing about the underlying cash flow (the user still spends \$120 over four periods). But loss aversion at $\lambda = 2.25$ means the framing change is roughly six times more impactful than the underlying price change. This is the behavioral mechanism behind BNPL conversion: “4 payments of \$30” is a UX choice, not a financial-engineering choice.

Common mistake. Students drop the λ multiplier and compute $30 - 120 = -90$ for Frame 1; they then under-state the reframing gain by a factor of more than 2.

QP7A.2 model answer must compute both frames and report the reframing-gain delta of \$202.50. Partial credit for correct λ application without the second frame.

The

QP7A.3 [Day 7A, Intermediate]: (b) Two pure-strategy Nash equilibria; (Adopt, Adopt) is payoff-dominant (+3, +3) but (Wait, Wait) is risk-dominant (0, 0); the chicken-and-egg problem is a stable low-coordination Nash.

Why both are Nash. At (Adopt, Adopt), neither side gains by deviating: $+3 \rightarrow -1$ for either deviator. At (Wait, Wait), neither side gains by deviating: $0 \rightarrow -1$ for either unilateral mover. Hence both pure-strategy profiles are Nash.

Why (Wait, Wait) is risk-dominant. The (Wait) action minimises maximum loss: a Wait-player loses at most 0; an Adopt-player risks -1 if the other side waits. Under Harsanyi-Selten risk dominance, the deeper-basin equilibrium is the one to which best-response dynamics converge from random starting points.

Game-theory interpretation for fintech. This is the canonical chicken-and-egg framing for two-sided platform launches. Subsidy launches (Uber driver bonuses, Stripe free integration credits) shift the payoff matrix temporarily so that adopting becomes a dominant strategy on at least one side, breaking the trap.

Distractors. (a) ignores the second Nash; (c) confuses coordination with Prisoner's Dilemma (here both Adopt is Pareto-superior, unlike PD where Defect is); (d) confuses positive payoff at (Adopt, Adopt) with dominant-strategy existence.

Coor
games have multiple Nash equilibria; identify both and rank by payoff-dominance vs risk-dominance. The chicken-and-egg trap is risk-dominance in action.

QA7A.4 [Day 7A, Advanced, Calc + Short]: Welfare-economics arithmetic on a two-sided platform.

Setup parameters (a 100-USD card transaction). Merchant fee 3.20, consumer cashback 1.20, cash counterfactual merchant handling 0.50, consumer inconvenience 0.20, platform marginal cost 0.40.

Part (i): Per-party welfare versus the cash counterfactual.

Merchant: receives 96.80, avoids 0.50 cash-handling cost. Versus cash baseline of 99.50, the merchant nets 96.80. Welfare change: minus 2.70.

Consumer: pays 100, receives 1.20 cashback, avoids 0.20 cash inconvenience. Versus cash baseline of minus 100.20, the consumer nets minus 98.80. Welfare change: plus 1.40.

Platform: revenue 3.20 minus MC 0.40 minus cashback 1.20 equals plus 1.60 producer surplus.

Part (ii): Total social welfare. Sum: minus 2.70 plus 1.40 plus 1.60 equals plus 0.30. Social welfare net of cash counterfactual is plus 0.30. No new deadweight loss is created; a small cash-friction DWL is closed.

Part (iii): Profitable versus welfare-improving. The platform earns plus 1.60 of producer surplus. Total social welfare gain is only plus 0.30. The platform extracts 2.70 of merchant welfare and recycles 1.40 to the consumer; net social welfare gain is barely positive. “Free for users” platforms route value across sides rather than creating it.

Common mistake. Students compute platform profit and stop, missing the cross-side redistribution which is the load-bearing welfare insight.

QA7A.4 welfare arithmetic must surface the redistribution insight; profitable and welfare-positive are different concepts; the platform sits inside a redistribution triangle on the merchant side.

QB7B.1 [Day 7B, Basic]: (b) $L = (P - MC)/P$ measures markup over marginal cost as a fraction of price.

Rationale. Abba Lerner (1934) defined this index as the operational measure of market power. $L = 0$ when $P = MC$ (perfect competition, zero markup). $L \rightarrow 1$ when $MC \ll P$ (monopoly limit, near-infinite markup). L is the elasticity-adjusted Cournot-style markup the firm extracts.

Application to fintech. Apple Pay charges banks roughly 0.15% per transaction on iPhone NFC payments. The marginal cost of facilitating one more NFC tap is essentially zero (the Secure Element is fixed infrastructure). Lerner index is therefore close to 1. Combined with the NFC restriction (an artificial barrier to entry under the dominance test), EU regulators concluded prima facie abuse of dominant position; this drove the July 2024 Apple Pay commitments.

Distractors. (a) confuses concentration with pricing power; a firm can have a high share without pricing power if entry is easy. (c) revenue is a scale measure, not a pricing-power measure. (d) NPS is satisfaction, not market structure.

index is the IO-text bridge between economic theory and competition law. EU enforcement against Apple Pay was built explicitly on the Lerner-index + entry-barrier two-part test.

QP7B.2 [Day 7B, Intermediate]: (b) EMT = single-fiat-referenced, redeemable at par, regulated as e-money under EMD2; ART = multi-asset or commodity-referenced with heavier prudential, reserve, custody, and disclosure rules, especially when the EBA classifies the token as “significant” under MiCA Article 43.

Specific MiCA references. MiCA Title III (Articles 16–47) covers ARTs. Title IV (Articles 48–58) covers EMTs. Title V (Articles 59–77) covers crypto-asset service providers. Title III is materially heavier on reserve composition, redemption rights, custody, and ongoing disclosure than Title IV; a “significant” ART under Article 43 triggers EBA-level supervision rather than national-competent-authority supervision.

Worked example. Circle’s USDC slots into the EMT regime in the EU because it references a single fiat currency (USD) 1:1 and is redeemable at par. A hypothetical multi-currency basket stablecoin (e.g., 40% EUR + 40% USD + 20% CHF) would fall under the ART regime; if it grew large enough to be “significant,” it would face EBA supervision, capital requirements on the issuer, and tighter reserve-composition rules. This is the regulatory architecture that contributed to Meta’s discontinuation of the Libra / Diem multi-currency stablecoin in 2022.

Distractors. (a) and (d) deny the legal split that MiCA explicitly codifies. (c) inverts the regulatory perimeter; ARTs are the more heavily regulated class.

EMT (Title IV, single-fiat, EMD2-aligned) vs ART (Title III, multi-asset, heavier obligations, EBA supervision if “significant”). Most-tested MiCA structural fact.

QP7B.3 [Day 7B, Intermediate]: (b) Alipay payments survived; integrated credit (Huabei + Jiebei) was carved out into a separately capitalised consumer-finance company subject to bank-style prudential capital; the credit-scoring data was shared with PBOC's national credit bureau; the IPO remained blocked; an approximately RMB 7.13 bn (about USD 1 bn) fine was levied on Ant in July 2023; in January 2024 Jack Ma transferred his controlling voting rights to a group of Ant executives.

What survived. The Alipay payment app, the consumer interface, the QR-code merchant network, the basic e-money e-wallet function. The user-visible product on a 2024 iPhone in Shanghai still looks largely like the 2019 product.

What was dismantled. The integrated super-app credit business (Huabei consumer credit + Jiebei small-business credit), the unrestricted Sesame Credit data flywheel feeding credit decisions, the founder-voting-control structure that triggered the November 2020 IPO suspension. The credit business now sits in a separately capitalised consumer-finance company with PBOC-equivalent capital requirements; the data flows into the national credit bureau rather than into Ant's own scoring engine.

Two 2023–2024 finalising events. (i) July 2023: regulatory fine of approximately RMB 7.13 bn (USD 1 bn) on Ant for historical violations. (ii) January 2024: Jack Ma transferred his controlling voting rights to a group of Ant executives, dispersing founder control and breaking the original IPO-blocking structure. Together these signalled that the restructuring phase was substantially complete.

Distractors. (a) ignores the restructuring entirely. (c) is factually wrong (Alipay still operates). (d) ignores all regulatory events.

Ant

2024 canonical fact set: Alipay survives; integrated credit carved out; July 2023 RMB 7.13 bn fine; January 2024 voting-rights transfer. Half-credit if either 2023/2024 event is omitted.

QA7B.4 [Day 7B, Advanced, Short]: DMA enforcement applied to Apple Pay NFC access mechanism design.

Model answer (must name all 4 elements):

(i) **Technical access mechanism.** Apple provides a Host Card Emulation (HCE) API plus regulated access to the iPhone's Secure Element via a non-discriminatory API surface. Third-party wallet providers register and obtain Apple-issued entitlements without paying Apple per-transaction fees on the NFC interface (the NFC chip itself is treated as an essential facility).

(ii) **Conformance / certification process.** EMVCo-style certification path open to third parties on equal commercial terms with Apple Pay. Published technical documentation, public test suite, deterministic certification timelines (e.g., 90-day SLA). Failure modes (rejection, suspension) must be reasoned in writing and appealable to a neutral conformance body.

(iii) **Non-discrimination rule.** (a) Default-wallet selection exposed to the user at first NFC use (no Apple-Pay-as-default lock-in). (b) No preferential ordering of Apple Pay in the payment sheet or in OS-level NFC handlers. (c) No per-transaction fees on the NFC interface that exceed Apple's own internal cost basis (fee parity at marginal cost).

(iv) **Monitoring / sanctioning mechanism.** European Commission monitoring trustee under the July 2024 commitments. Annual compliance reports filed with the Commission and ECN. Sanction architecture under DMA Article 30: fines up to 10% of global annual turnover for first infringement, 20% for repeat infringement; periodic penalty payments up to 5% of average daily worldwide turnover; structural remedies (divestiture) for systematic non-compliance.

Common mistake. Students propose “open NFC” without specifying HCE, certification, fee parity, or sanction. The answer key requires all four named elements; partial credit is proportional.

four-element DMA enforcement template (technical access + certification + non-discrimination + sanction) generalises to any DMA case (browser engines, app stores, default-search-engine choice).

The

QB7C.1 [Day 7C, Basic]: (b) Customer Segments is the Business Model Canvas (BMC) box that captures who the firm serves.

Rationale. Osterwalder and Pigneur's BMC names nine boxes; "Customer Segments" specifically holds the cluster(s) of users targeted. For Nubank, the segment choice ("underbanked Brazilians first, then upsell middle class") is foundational; it cascades into Channels (mobile-first, branchless), Value Propositions (no-fee credit card, transparent pricing), Revenue Streams (interchange + interest), and Cost Structure (no branch overhead).

Why the distractors fail. (a) Key Resources captures what the firm owns (transaction data, brand, software). (c) Cost Structure captures what the firm spends. (d) Channels captures how the firm reaches customers (the mobile app, the website), not who they are.

Common mistake. Students confuse Customer Segments with Channels. Test: "A 25-year-old Brazilian student using mobile data" = Customer Segment. "A mobile-first iOS / Android app delivered via the app store" = Channel.

Segments is the load-bearing BMC box for fintech because the segment choice cascades into the other eight boxes; mis-identifying the segment box leads to incoherent canvas analysis.

QP7C.2 [Day 7C, Intermediate]: (b) Counter-Positioning and Switching Costs are the two Helmer 7 Powers most clearly active for Stripe.

Counter-Positioning (the dominant Helmer Power for Stripe). Stripe built a developer-API model that legacy acquirers (Worldpay, FIS, Chase Paymentech) cannot rationally copy because doing so would cannibalise their sales-led enterprise contracts. Worldpay's "sales engineer-led integration" is the strategic asset that makes Worldpay valuable to large merchants and unable to offer Stripe-style 30-minute self-serve onboarding. Helmer (2016): a Power is Counter-Positioning when the incumbent's optimal action is to *not* copy.

Switching Costs (second Helmer Power). Once a merchant integrates the Stripe API into checkout, payouts, subscriptions, Radar (fraud), Sigma (analytics), Connect (marketplaces), and Tax, the migration cost rises sharply. Re-coding eight integrations, retraining staff, re-certifying for PCI-DSS, and re-onboarding processors creates a high effective lock-in.

Why Network Economies do not dominate. Stripe rides on existing Visa / Mastercard rails; it does not build a two-sided network of users-and-acceptors itself. Network Economies are weaker than for Visa.

Distractors. (a) Branding and Process Power are present but secondary; (c) Cornered Resource and Scale Economies are weaker. (d) is dismissive and incorrect.

Positioning is the most-tested Helmer 7 Powers concept for fintech: it captures why incumbents cannot rationally respond to a disruptor without destroying their existing economics.

QP7C.3 [Day 7C, Intermediate]: (b) The Bass model rate form $\frac{dN}{dt} = (p + q \cdot N/m)(m - N)$ with $p \approx 0.004$ and $q \approx 0.55$ implies viral-dominated growth.

Step 1: ratio q/p . $q/p \approx 0.55/0.004 \approx 137.5$. Imitation pull dominates innovation push by two orders of magnitude.

Step 2: peak adoption time. Closed-form integration of the Bass model yields peak adoption rate at $t^* = \frac{1}{p+q} \ln\left(\frac{q}{p}\right)$.

Substituting: $t^* = \frac{1}{0.554} \ln(137.5) \approx 1.805 \cdot 4.924 \approx 8.9$ years from launch. Empirically Nubank's user-count trajectory peaks in adoption rate around 2022–2023, consistent with the 8.9-year estimate from the 2014 launch.

Step 3: peak adoption level. The Bass model peaks at $N \approx m/2$, which for Nubank's LatAm $m \approx 350\text{m}$ gives a peak rate at about 175m cumulative adopters.

Strategic interpretation of $p:q$. The ratio is the actionable strategic input: a fintech wanting to compress time-to-mass-market raises p (paid acquisition); a fintech accepting longer time but optimising for capital efficiency raises q (referral programmes, social proof, viral loops). Nubank chose the latter and won.

Distractors. (a) inverts the sign of dominance. (c) is mathematically wrong (the logistic-like Bass curve has an inflection point). (d) denies the well-established behavioural meaning of q .

p:q ratio is the load-bearing Bass parameter; peak time formula $t^ = \frac{1}{p+q} \ln(q/p)$ is the standard closed-form derivation. Students must reproduce the closed form for full credit.*

Day 7C — Advanced Answer (Three Horizons Portfolio)

QA7C.4 [Day 7C, Advanced, Short]: Three Horizons portfolio analysis applied to a retail bank.

Part (i): Classification with one-line justification.

Legacy mortgages and current accounts are H1: extend existing core, payback of 1 to 2 years, defended by regulation, brand, and switching costs; this is the bank's quarterly-earnings engine.

Robo-advisor pilot is H2: emerging business with revenue line of sight 2 to 5 years; fits within the existing banking licence but targets a new customer segment (digital-native investors) and a new tech stack (algorithmic portfolio construction).

DLT wholesale-settlement experiment is H3: option-creating, 5 to 10 year horizon, no near-term profit and loss; creates the right to participate if wholesale-CBDC and tokenisation become real settlement rails.

Part (ii): Portfolio shape comparison. The bank's 95/4/1 allocation is severely H1-skewed relative to the textbook 70/20/10 ideal mix (Baghai-Coley-White, 2000). This is Christensen's innovator's dilemma stated in budget form: the bank optimises for the present and underfunds the future. McKinsey research (2000 to present) shows most large banks sit between 90/8/2 and 96/3/1; the bank is at the under-investing tail.

Part (iii): Rebalancing recommendation. Ring-fence at least 15 percent of the innovation budget for H2 ventures with separate P+L lines and a 5-year patient-capital horizon (so the robo-advisor cannot be killed by quarterly earnings pressure). Allocate at least 8 percent to H3 via a corporate-VC vehicle with an independent investment committee structurally insulated from the operating bank (so DLT and tokenisation work cannot be defunded). The numerical reweighting (15 and 8) is necessary; the governance mechanism (separate P+L, patient capital, independent committee) is what makes the reweighting durable.

Common mistake. Students recommend "invest more in innovation" without a numerical target or a governance

108 model answers complete (96 modules + 12 Day 7).

- **Toggle:** Set `\solutionsfalse` (or comment the `\AtBeginDocument` line) to compile a question-only deck for distribution.
- **Format:** Aligned with `exam_bank.tex` and `exam_prep_guide.tex`; uses the `SHARED_ANSWER_KEY_TEMPLATE` structure (setup → formula → substitution → answer with units → interpretation → common mistake) for the 8 calculation items.
- **Common-mistake highlights:**
 - QB1.3: students often invert acquirer→issuer interchange flow
 - QB4.1: VaR confused with expected loss
 - QP3.1: liquidation threshold ambiguity (now fixed)
 - QP5.1: confusion-matrix arithmetic must close
 - QP7.1: 99% FP rate *and* 71% recall both must be reported
- **Reminder:** This bank is *optional practice*. The graded assessment is Project 50% + Presentation 20% + Quizzes 30%.