Activity 7B: Mechanism Redesign – SOLUTIONS

Digital Finance Intensive Course — Prof. Dr. Joerg Osterrieder

INSTRUCTOR COPY – do not distribute to students

Grading note. Mechanism redesign admits multiple correct answers. The key is whether the student (a) correctly diagnoses the incentive failure, (b) proposes a redesign with an explicit mechanism change (not just "better rules"), and (c) identifies who gains, who loses, and what new failure mode the redesign introduces. A redesign without a failure mode analysis is incomplete.

Problem 1 – Constant Product AMM and Impermanent Loss

(a) Why does the constant product rule create impermanent loss?

When ETH price doubles externally (e.g., on Coinbase), the pool price has not yet updated. Arbitrageurs buy cheap ETH from the pool, selling the other token (USDC) into it. Result: the pool now holds more USDC and less ETH than before. LPs, who deposited 50/50, now hold a portfolio that is heavier in the *depreciating* asset (USDC) and lighter in the appreciating one (ETH). Had they simply held, they would own more ETH. The constant product rule enforces automatic portfolio rebalancing *against* the LP when prices trend. The mechanism forces LPs to "buy high and sell low" by rebalancing continuously at the wrong price.

(b) Model redesign: concentrated liquidity (Uniswap v3 style)

Mechanism change: LPs specify a price range $[p_a, p_b]$ within which their capital is active. Outside that range, the LP's capital is fully in one asset (not at risk of impermanent loss from further movement). LPs earn higher fee yield per unit of capital because fewer LPs compete in each narrow range.

Who gains: Sophisticated LPs who correctly forecast the price range earn 10-50x more fees than v2. Traders get lower slippage (deeper effective liquidity around the current price).

Who loses: Passive LPs who cannot actively manage ranges: they earn less than v2 as their capital is frequently out of range, earning nothing. Small liquidity providers are effectively priced out.

New failure mode: LP positions require active range management. If the price moves outside the range, the LP holds 100% of one asset and earns zero fees until they reset. This introduces *active management risk*: a passive LP suffers worse outcomes in v3 than v2 unless they continuously monitor and rebalance.

(c) Can any mechanism eliminate IL entirely?

No – impermanent loss is an inescapable consequence of providing liquidity against informed arbitrageurs. Any mechanism that allows prices to adjust via arbitrage will result in LPs buying the falling asset and selling the rising one. Dynamic fees can *compensate* LPs for IL (higher fees during high volatility), but compensation is not elimination. Oracle-guided AMMs (e.g., Gyroscope) can reduce IL by pricing trades closer to the oracle price, but this trades the IL problem for oracle manipulation risk and reduced composability. The only way to fully eliminate IL is to remove the LP from bearing inventory risk – which means the AMM stops functioning as a decentralised market maker and becomes a centralised orderbook.

Problem 2 – Credit Score Misalignment

(a) Principal-agent map

Principal: Ostensibly the consumer (whose creditworthiness is scored). **Economic principal:** The lender (who pays for the score and relies on its accuracy to price loans).

Agent: The credit bureau (Equifax, Experian, TransUnion).

What does the agent optimise when paid by lenders? Bureaux are paid *per query* by lenders, not on accuracy. Their incentive is to maximise data volume (more attributes = higher perceived value) and lender queries, regardless of accuracy. Error correction is costly; disputes are handled reactively only when legally required (FCRA, GDPR). The consumer bears the cost of errors (loan rejection, higher rate) but cannot pay for correction; bureaux internalise no penalty for inaccuracy.

(b) Model redesign: accuracy-based bureau compensation

Mechanism change: Bureaux earn a fee only when their predicted score correctly forecasts loan repayment. A delayed payment is made 12-24 months after loan origination, conditioned on whether the borrower repaid as predicted. Incorrect predictions (defaults predicted as good, or good borrowers rejected) reduce the fee or trigger clawbacks.

Who benefits: Consumers (fewer errors; bureau now has incentive to include exculpatory data), lenders (more accurate risk pricing), regulators (reduced systemic credit mispricing).

Who loses: Bureaux (revenue delayed, volatile, requires accurate forecasting infrastructure). Lenders lose ability to blame the bureau for mis-pricing.

Barriers to adoption: Attribution problem – was the default caused by bad scoring or by an external shock (pandemic, job loss)? Requires time-delayed compensation, which creates cashflow risk for bureaux. Dominant incumbents have no incentive to voluntarily adopt a model that makes their revenue contingent on performance.

(c) Sesame Credit (Ant Group) – better or worse from consumer welfare perspective?

Akerlof framework applied: Traditional credit bureaux have a "lemons" problem in data quality (low-quality data crowds out high-quality signals; bureaux cannot credibly commit to accuracy). Sesame Credit resolves *one* adverse selection problem by using richer, higher-quality signals (payment behaviour, Alipay network connections, purchase patterns).

But introduces a new problem: Ant owns *both* the scoring model *and* the primary payment data. There is no external check on model accuracy; Ant can adjust scores to serve its credit product revenue. The consumer has no exit: switching away from Alipay reduces the Sesame score. This is the *worse* outcome from a consumer welfare perspective: a well-functioning lemons equilibrium has been replaced by a monopoly equilibrium with even less consumer recourse. **Verdict:** Better signal quality, worse governance and accountability structure.

Problem 3 – MYbank Loan-to-Deposit Ratio

(a) Why does the 310 model create systemic risk even if the ML model is excellent?

Tail risk: The model is trained on historical data. A black-swan event (COVID, credit freeze) creates defaults in a region of the feature space the model has never seen. The model's confidence scores are meaningless outside the training distribution.

Model correlation: MYbank's 1.7tn RMB portfolio is largely SME loans to Taobao sellers. If Taobao GMV falls (e.g., regulatory crackdown, consumer sentiment shift), defaults are correlated across all borrowers simultaneously – the model cannot diversify because all signals come from the same platform. A single regulatory shock to Alibaba would impair the entire loan book.

Liquidity vs. solvency: MYbank's 10% reserve means it can handle small random fluctuations but cannot survive a confidence shock (bank run). Digital bank runs can propagate in hours, not days. Even a solvency-sound bank fails if it cannot meet immediate withdrawal demand. The mechanism (approve instantly, hold minimal reserves) works perfectly until a correlated shock causes simultaneous withdrawals and correlated defaults.

(b) PBOC's financial holding company requirement: right redesign?

What it fixes: Forces Ant to hold capital commensurate with its loan book (8-12% Tier 1 capital under

Basel III equivalents), eliminating the reserve mismatch. Regulatory perimeter now covers Ant's systemic risk; contagion risk to the broader financial system is reduced.

What it breaks: The 310 model's social value was precisely the absence of capital overhead. Traditional banks required capital because they bore credit risk on their balance sheet. Ant's value proposition was that algorithmic risk selection allowed it to price and extend credit to borrowers traditional banks ignored (SMEs, rural merchants), at lower cost, because it had better information. Forcing bank capital requirements onto Ant forces it to price loans like a bank, reducing access benefits for underserved borrowers. This is a *mechanism transplant* – applying a mechanism designed for deposit-funded, relationship-lending banks to a data-driven, capital-light platform. The fit is imperfect.

(c) Alternative redesign: tiered capital + model auditability

Mechanism: Capital requirement is a function of portfolio concentration, not total volume. Reserve ratio = base rate + concentration premium, where concentration is measured by the correlation of default probability across borrowers (e.g., all in the same supply chain). Additionally: the ML model must be submitted to PBOC for adversarial stress testing against specified black-swan scenarios before each quarterly approval cycle.

What it preserves: Capital-light operation for genuinely diversified portfolios; fast approval for small-ticket loans that do not create systemic concentration.

What it adds: Systemic risk disincentive (Ant internalises the capital cost of correlation); regulatory transparency without nationalising the model; dynamic requirement that adjusts as the portfolio evolves, rather than a static Basel-equivalent floor.

Problem 4 – Gas Auction Front-Running and MEV

(a) EIP-1559 fixed overbidding but not MEV. Why?

EIP-1559 fixed the *auction mechanism* (first-price gas auction replaced by base fee + optional tip). It did not touch the *mempool visibility* or the *validator's ordering power*.

The structural feature enabling MEV: **public mempool + validator discretion over ordering**. Transactions are broadcast publicly before inclusion. Any searcher can observe a large pending swap, predict its price impact, and insert a transaction before (buy) and after (sell) it. EIP-1559's base fee is burned – it doesn't compensate the victim of sandwich MEV. The validator still chooses which transactions to include and in what order; a searcher can still outbid the victim for queue position via the optional tip. The mechanism's flaw is architectural: separating transaction broadcast from transaction execution creates a front-running window regardless of fee structure.

(b) Model redesign: encrypted mempools

Mechanism change: Transactions are submitted encrypted. The transaction content (token amounts, slippage tolerance) is only revealed after the block is finalized. Validators include transactions without knowing their content; ordering is done on encrypted blobs. The encryption key is shared only after inclusion, making it impossible to front-run the transaction content.

What it sacrifices: (1) *Composability*: smart contracts that depend on reading other pending transactions (e.g., liquidation bots) cannot operate in encrypted mempools. (2) *Transaction speed*: threshold decryption requires coordination among validators before revealing content, adding latency. (3) *Decentralisation*: encrypted mempools typically require a decentralised threshold network (e.g., Shutter Network) – if the threshold network has 5 of 10 validators, compromising 5 nodes leaks all transaction content.

(c) Is some MEV socially beneficial? Should the redesign distinguish types?

Yes – arbitrage MEV is socially beneficial. When Uniswap ETH/USDC price deviates from Binance, arbitrageurs correct the discrepancy, providing accurate price discovery. Without this mechanism, AMMs would have persistently stale prices, making swaps worse for all users.

Sandwich MEV is pure extraction. No user benefits; value is transferred from the swapper to the searcher with zero social surplus.

How would a mechanism distinguish them? *Structural difference:* Arbitrage MEV requires executing a round-trip trade across two or more venues to close a price gap. Sandwich MEV requires inserting transactions immediately before and after a specific victim transaction. A fair-sequencing service can apply time-ordering (FCFS) within short time windows (e.g., 12 seconds): arbitrage bots get their slot fairly; sandwiching requires inserting *relative to a specific pending transaction*, which FCFS prevents. *In practice:* No mechanism perfectly separates them – MEV taxonomy is contested. Flashbots attempted a public good framing (redistribute MEV to validators via structured auctions) but shifted the centralisation to a new bottleneck (Flashbots builder market share >50% in 2023).

Cross-problem common thread (model answer): All four failures share the same structure – an intermediary (arbitrageur, bureau, bank, validator) can *observe a signal before the mechanism's intended beneficiary can respond.* The arbitrageur sees the LP's stale price. The bureau sees the lender's data dependency. The 310 model sees correlated defaults before the regulator does. The validator sees the pending transaction before the user's trade executes. Good mechanism design either removes the signal advantage, delays it until it is no longer actionable, or ensures the intermediary's payoff is aligned with the beneficiary's outcome.