

The Economic Theory of Digital Markets

Why Digital Finance Looks the Way It Does

Prof. Dr. Joerg Osterrieder

Digital Finance – Intensive Course: Day 7A

Every product feature in digital finance is an answer to a market failure. Today we name the failures.

The Stripe Puzzle

Student A: "I just paid for lunch with my credit card. Zero fee. But Stripe charged the restaurant 2.9% plus 30 cents. That seems backwards."

Student A: "But WHY is the restaurant charged and not me? Who decided that?"

Student B: "Maybe the restaurant just absorbs the cost?"

Prof: "A 2003 economics paper. Two authors. One Nobel Prize. And it explains every pricing decision in digital finance."

Today we build the economic toolkit that makes the Stripe puzzle, the lemons problem, and the innovator's dilemma all obvious.

After today you will explain:

1. Why Stripe charges merchants, not users (Two-sided market theory – Rochet-Tirole 2003)
2. Why banks reject good borrowers even when rates are high (Information Asymmetry – Akerlof 1970)
3. Why JPMorgan cannot simply build Revolut (The innovator's dilemma – Christensen 1997)

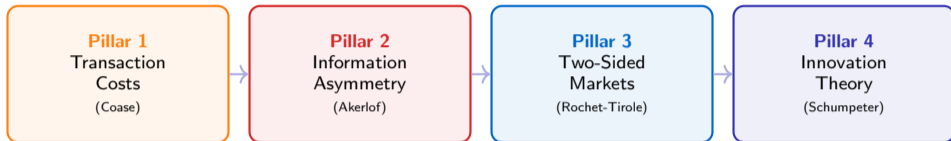
The common thread:

All three are market failures.
Each has a well-understood economic name.

Digital finance exists to *solve* these failures.

Every fintech company is an engineering solution to an economic problem. Identify the problem and you understand the strategy.

Today: Four Pillars of Economic Theory



Each pillar = one market failure = one fintech category

No prior economics required. Every concept is built from an example you already understand. The math follows the intuition.

These four pillars are complementary, not competing. A single platform like Ant Group exploits all four simultaneously.

The textbook assumption:

- Buyers and sellers know prices instantly
- Contracts are costless to write and enforce
- Both sides know the quality of goods
- New competitors can enter freely

None of these hold in finance:

- Finding a trustworthy counterparty is expensive
- Verifying creditworthiness requires deep information
- Contracts require lawyers, courts, collateral
- Regulation and capital requirements block entry

Result:

Finance is the industry where market failures are most severe – and most profitable to solve.

Every bank, every platform, every DeFi protocol exists because a textbook market would fail here.

The four pillars explain WHY these failures occur. Fintech is applied economics: use technology to make the failure smaller or cheaper.

Pillar 1: Transaction Costs (Coase, 1937)

Ronald Coase asked: Why do firms exist? If markets allocate resources efficiently, why hire employees rather than contracting every task on the open market?

His answer: Because using markets has costs. Finding suppliers, negotiating prices, writing contracts, enforcing delivery – these are *transaction costs* (TC).

The Coase result:

- When $TC = 0$: use markets (they are costless)
- When $TC > 0$: firms or intermediaries appear to reduce TC

Applied to finance:

A bank is not a store of money.
A bank is a *TC-reduction* machine.

It eliminates the cost of finding counterparties, verifying creditworthiness, and enforcing repayment.

Coase insight: identify the TC a company eliminates and you understand why it exists. Every fintech has a Coasian purpose.

The Four Types of Transaction Costs in Finance



Each row of the fintech landscape maps to one TC type. Every new product eliminates or reduces one or more of these four costs.

Diagnostic tool: ask which TC a new fintech reduces. If you cannot answer, either the business model is unclear or the category does not exist yet.

Why financial TC are especially large:

- **Trust** must be manufactured from scratch for every new relationship
- **Time** separates payment from delivery (repayment risk)
- **Asymmetric information**: borrowers know more than lenders
- **Jurisdiction**: enforcement crosses legal systems

The result:

- Traditional banks charge 3–5% spread on loans
- International wire transfers cost \$25–45 and take 2–3 days
- Cross-border equity settlement: 3 business days, multiple intermediaries

Fintech thesis:

Software can automate the TC-reduction that banks currently do with branches, lawyers, and loan officers.

Lower TC \Rightarrow lower prices
 \Rightarrow larger addressable market.

The spread between the rate a bank charges borrowers and pays depositors is a TC proxy: it measures what trust and verification cost in that market.

Context (Kenya, 2007):

- Sending money home from Nairobi: give cash to a bus driver (\$2–5 fee, unsafe, 1–2 days)
- Banking penetration < 15% of adults
- No branch infrastructure in rural areas

M-Pesa's TC attack:

- **Search:** agent network replaced branch search cost
- **Verification:** mobile number as identity (SIM-based KYC)
- **Enforcement:** Safaricom's network as reputational collateral
- **Result:** fee fell to <\$0.10, instant settlement

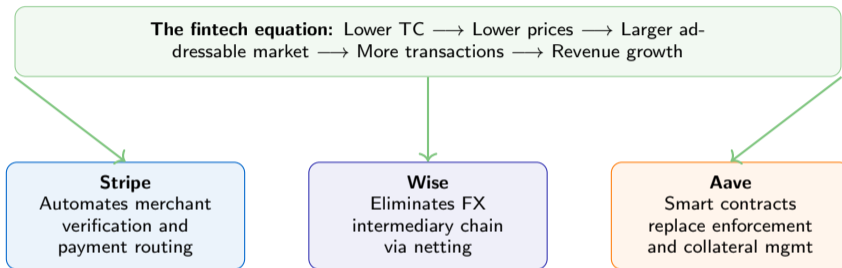
Why it worked:

Mobile penetration was higher than bank penetration. M-Pesa used an existing trust infrastructure (Safaricom) to substitute for a missing bank infrastructure.

TC collapsed from \$3 to \$0.10.

M-Pesa's moat: it owns the cheapest verification infrastructure in its market. Any competitor must either build a rival network or pay M-Pesa's verification cost.

Digital Finance: Software That Attacks Transaction Costs



Three different TC types, three different solutions. The Coasian lens predicts what problem each company must solve to grow: shrink its TC category further.

Pillar 2: The Lemons Problem (Akerlof, 1970)

George Akerlof's question (1970): Why does the used car market malfunction?

The setup:

- Good cars (“peaches”): worth \$10,000
- Bad cars (“lemons”): worth \$4,000
- *Only the seller knows which it is*

The logic:

- 1 Buyer cannot tell peach from lemon
- 2 Buyer offers *average* price (\$7,000)
- 3 Peach owners refuse: $\$7,000 < \$10,000$
- 4 Only lemons remain in market
- 5 Buyer lowers offer to \$4,000
- 6 Market collapses into all-lemon equilibrium

This Nobel-winning insight:

Asymmetric information does not just reduce market efficiency.

It can *destroy* markets entirely.

Financial markets are full of lemons problems.

Akerlof (1970) won the Nobel Prize in Economics in 2001 (shared with Spence and Stiglitz) for showing that information asymmetry is a fundamental market failure, not a minor friction.

Replace “used cars” with “loans”:

- Good borrowers: low default risk, willing to pay 6%
- Bad borrowers: high default risk, willing to pay 12%
- *Only the borrower knows which type they are*

The Akerlof spiral in lending:

- 1 Bank cannot distinguish types
- 2 Bank charges average rate (e.g. 9%)
- 3 Good borrowers leave (9% > their fair price of 6%)
- 4 Remaining pool is mostly bad borrowers
- 5 Bank raises rate – good borrowers leave further
- 6 *Credit market thins or collapses*

Adverse selection:

Raising the interest rate selects *against* good borrowers.

Unlike normal goods, higher price makes the average quality of buyers *worse*.

This is why credit markets need more than just price signals.

Adverse selection is why banks cannot solve the credit problem by simply raising interest rates: higher rates drive out the best borrowers, making the problem worse.

Market Response 1: Signaling (Spence, 1973)

How can good borrowers communicate quality?

Michael Spence's insight: A signal works only if it is *too costly* for bad types to fake.

Financial signals:

- **Collateral:** only borrowers confident of repayment pledge assets
- **Credit history:** years of timely repayment are costly to fake (they require actually repaying)
- **Down payment:** large deposits signal commitment and liquidity
- **Co-signers:** personal relationships stake reputation

Signals must have *differential cost*: cheap for good types, expensive for bad types.

Digital disruption:

Fintechs create new signals that are cheaper for honest borrowers:

Rent payment history, utility bills, e-commerce transaction patterns, social connections.

More signals \Rightarrow thinner adverse selection.

Signaling theory explains why fintech credit scores using alternative data are economically meaningful, not just technically clever: they lower the cost of a good signal for underserved borrowers.

Market Response 2: Screening (Stiglitz, 1975)

If borrowers cannot credibly signal, lenders can screen.

Joseph Stiglitz's insight: Design a *menu of contracts* where each type self-selects into the option that reveals their type.

Classic screening menu in lending:

- Contract A: 6% rate, 20% collateral required
- Contract B: 10% rate, 0% collateral required

Good borrowers (confident of repayment) choose A. Bad borrowers (avoiding collateral risk) choose B.

The menu *separates* the types without the lender knowing them directly.

Screening in practice:

Credit scoring: menu of products calibrated to risk buckets.

Insurance deductibles: high-deductible plans attract low-risk policyholders.

Vesting schedules: long vesting periods screen out short-tenure employees.

Every pricing tier, every product variant, every feature bundle is potentially a screening device. Good product designers are implicitly doing mechanism design.

The Digital Twist: Data Abundance Flips the Asymmetry

Traditional finance: borrower knows more than lender. This causes adverse selection. Lenders respond with signaling requirements and screening menus.

Digital finance platforms: the platform may know *more than you know about yourself*.

- Your Alipay transaction history reveals income, habits, and social network
- Your Google Pay behaviour predicts creditworthiness better than your stated salary
- Your ride-sharing data reveals employment stability

The asymmetry reverses: Platform knows more \Rightarrow platform can price risk more accurately \Rightarrow lower rates for good borrowers **and** higher profit.

But: reversed asymmetry creates new problems.

Discrimination risk: data-rich models may encode historical biases.

Exploitation risk: platform uses info advantage to extract more surplus.

Opacity risk: borrower cannot understand their own credit decision.

The information advantage that enables platform credit is the same advantage that enables price discrimination. The same mechanism that benefits borrowers can also harm them.

Credit Scoring: A Screening Mechanism in Disguise

FICO score (US): invented 1956, still dominant.

- Payment history (35%): did you repay on time?
- Amounts owed (30%): credit utilisation ratio
- Length of history (15%): years of observable behaviour
- New credit (10%): recent hard enquiries
- Credit mix (10%): variety of credit types

What it is economically: A continuous-valued screening score that collapses the full distribution of borrower types into a single number. Lenders use it to choose which contract to offer.

A credit score is not a measurement. It is a *contract-assigning device*.

Score 720+: 3.5% mortgage

Score 680-719: 4.0%

Score 620-679: 5.5%

Score below 620: denied

Each threshold is a *designed* screening boundary, not a natural feature of the data.

Credit scores are economic mechanisms: they reveal information while also shaping behaviour. A borrower who optimises for score is responding to the mechanism's incentive structure.

Ant Group's Zhima Credit ("Sesame Credit", 2015):

- Inputs: payment history, purchase patterns, social connections, registered assets, contractual fulfilment
- Score range: 350–950
- Over 600M users scored; 45M SME credit decisions

What Zhima could predict that FICO cannot:

- Whether a borrower paid rent (via Alipay rent transfer data)
- Whether a borrower's social connections had defaults
- Whether spending patterns matched stated income
- Whether deliveries were accepted or refused (punctuality signal)

The information reversal:

Traditional bank:
borrower knows more.

Ant Group:
platform knows more.

Result:

MYbank (Ant's lending arm)
approved SME loans
in 3 minutes with
0 human underwriters,
<1% default rate.

Traditional bank:
2 weeks, documents,
collateral required.

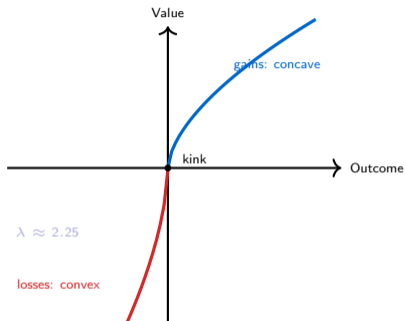
Zhima Credit is the most powerful credit scoring system ever built. Its advantage is not the algorithm – it is the data moat. Removing the data makes it an ordinary scorecard.

Tversky and Kahneman (1979): the classical economic agent is too rational. Real people decide under bounded rationality, anchoring, and asymmetric attitudes to gain versus loss.

Prospect theory in one paragraph:

- Value function is *concave over gains* (diminishing marginal joy)
- Value function is *convex over losses* (diminishing marginal pain)
- A sharp *kink at zero*: losses hurt about 2.25 times as much as equivalent gains feel good
- Loss aversion ratio $\lambda \approx 2.25$ is the canonical empirical estimate

Why this matters for finance: risk decisions depend on the reference point (today's wealth), not on absolute outcomes. Frame the same lottery as a loss and people refuse it; frame it as a gain and they accept.



Prospect theory generalises information-asymmetry: not only do agents lack information, they also weight gains and losses asymmetrically. Every pricing decision in digital finance is read through a kinked value function, not a smooth one.

Every UX choice is an economic mechanism. Designers exploit (or fight) behavioral patterns by default.

Three concrete examples in 2024-2026 fintech:

- **Robinhood gamification:** confetti on first trade, streak counters, push notifications. About 78% of active retail users log in daily, well above passive-investing app norms.
- **BNPL default selection:** Klarna and Afterpay default to the 4-installment plan at checkout; opt-out friction lifts conversion to credit by an estimated 20-30%.
- **Mental accounting in payment apps:** “round-up to save” (Acorns, Revolut Vaults) reframes spending as saving without changing dollar flows.

Auto-enrollment lift: 401(k) participation jumps from about 60% (opt-in) to about 90% (opt-out). The default IS the mechanism.

Every UX choice in a fintech is an economic mechanism. Defaults, framing, and friction set the price of an action even when no dollar changes hands; designers are doing mechanism design whether they know it or not.

Reading the UX:

Behavioral lever =
economic lever.

Default option =
choice architecture.

Loss-framing prompt =
coordination device.

Push notification =
attention extraction.

*Every screen carries
a measurable behavioral
effect on the user-side
of the platform.*

System-1 thinking meets two-sided pricing.

Apple App Store charges *users* \$0 to install apps; charges *developers* 15-30% commission. The classical Rochet-Tirole answer is “charge the inelastic side.” The behavioral answer adds a deeper layer:

- **Loss aversion** on the user side: a 99 cent app fee feels like a loss, even when value is far greater
- **Sunk-cost trap**: once users have 200 apps installed, switching ecosystems is psychologically painful, not just technically expensive
- **Status-quo bias**: every iOS user defaults to App Store; deliberate alternative is friction
- **Mental accounting**: developer commission feels like a business cost; user friction feels like a personal cost

Result: the platform monetises the side whose behavioral biases are easier to harvest – developers, not users.

Behavioral asymmetries shift the platform pricing rule. The user side is monetised indirectly through attention and lock-in, not directly through fees, precisely because users are more loss-averse to dollar costs than the supply side.

Behavioral asymmetry shifts the pricing rule:

Classical: charge the inelastic side.

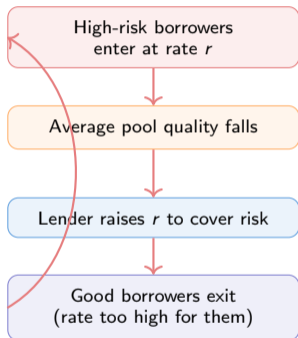
Behavioral: charge the side whose biases are *less* loss-averse toward dollar costs.

For digital platforms, that is almost always the supply side, not the user side.

Free for users is not generosity; it is loss-aversion management.

Formal Intuition: How Asymmetry Collapses Markets

The adverse selection spiral (diagram):



The formal condition:

- Let $\theta \in \{L, H\}$ be borrower type (Low/High risk)
- $r(\theta)$ = rate that type θ is willing to pay
- Lender observes r but not θ
- At pooling rate \bar{r} :

$$r(\bar{r}) < r(L) \Rightarrow \text{low-risk exits}$$

Remaining pool:

$$\Pr(H | \text{stays}) \rightarrow 1$$

Market unravels.


Credit markets do not clear by price alone. They require screening, signaling, or information technology to separate types. This is the economic foundation of every credit fintech.

Welfare Economics: Consumer Surplus, Producer Surplus, Deadweight Loss

Three welfare quantities measured on the demand/supply diagram:

- **Consumer surplus (CS):** value buyers receive above the price they paid (triangle above price, below demand curve)
- **Producer surplus (PS):** value sellers receive above their cost (triangle below price, above supply curve)
- **Deadweight loss (DWL):** value destroyed when trades that would benefit both sides do not happen (triangle between supply and demand, beyond traded quantity)

Why this matters for fintech: “the platform makes money” is a producer-surplus statement. “The platform is welfare-improving” is a CS + PS minus DWL statement. They are not the same.



Surplus accounting separates “the platform makes money” from “the platform is welfare-improving.” DWL is the value of trades that did NOT happen; markets with high DWL are exactly where fintech intermediation captures large profits by closing the wedge.

Welfare in Two-Sided Platforms: Cross-Subsidy as Hidden Redistribution

The Stripe 2.9% + \$0.30 fee is not just a price; it is a welfare transfer.

Worked example: a \$100 transaction.

- Merchant pays Stripe: $\$2.90 + \$0.30 = \$3.20$
- Consumer pays nothing; may receive 1-2% cashback
- Stripe's cost of processing: roughly \$0.40 (interchange + acquiring float)
- Stripe captures roughly \$2.80 of producer surplus

The hidden redistribution: Of the \$3.20 the merchant pays, roughly \$1.00-\$1.50 flows back to the consumer side as cashback, free fraud protection, and chargeback rights.

Merchant fees subsidise the consumer-side ecosystem.

The consumer-facing "free" product is funded by a wedge that lives entirely on the merchant side of the welfare triangle.

Welfare audit of a two-sided platform:

Add CS gained by side B
(cardholders: cashback, fraud protection)

Subtract PS extracted from side A
(merchants: fee burden)

Subtract DWL from merchants priced out of card acceptance

Net welfare = sum, and it is not always positive.

Cross-subsidisation is hidden welfare redistribution, not free service. Every "user pays zero" platform routes value from one side to the other; whether the net is welfare-positive depends on which DWL it closes versus which one it opens.

Welfare Loss in Practice: A Preview of MEV (Day 7B Bridge)

MEV (Maximal Extractable Value) is welfare economics applied to blockchains.

A trader sends a swap on Uniswap. Validators and bots can reorder, insert, or front-run the transaction to extract profit. This profit is not earned through intermediation; it is *extracted from* the trader.

The welfare triangle on a Uniswap trade:

- **CS retained by trader:** the swap they wanted, minus slippage
- **PS captured by liquidity providers:** the 0.30% pool fee
- **DWL to gas + MEV:** value extracted by validator reordering and sandwich attacks

Order of magnitude (2024-2026): estimated \$300-500M per year extracted via Ethereum MEV. Every dollar is a wealth transfer, not a productivity gain.

MEV extraction is a wealth transfer from traders to validators and bots, not a productivity gain. This frame bridges to Pillar 3 (two-sided markets price the externality) and to Day 7B (AMM mechanism design tries to redesign the DWL away).

Welfare reading of MEV:

Validators and bots do not create the surplus; they redirect it.

From a public-finance standpoint, MEV is a regressive tax on retail traders.

From a mechanism-design standpoint, MEV is the DWL that AMMs leave on the table by design.

Day 7B Section 2 formalises this as a mechanism failure.

Pillar 3: Two-Sided Markets (Rochet-Tirole, 2003)

Back to the Stripe puzzle:

- Merchants pay 2.9% + \$0.30 per transaction
- Cardholders pay nothing, sometimes get 1–2% cashback
- Stripe subsidises the user side entirely


Why?

Because Stripe is a *two-sided platform*. It creates value only by connecting two groups who need each other:

- **Side A:** Merchants need customers with payment capability
- **Side B:** Cardholders need merchant acceptance

Each side's value depends on the size of the other side. This is a *network externality between groups*.

Stripe charges merchants because merchants benefit more from user adoption than users benefit from merchant adoption. The externality asymmetry determines who pays.



What Makes a Market “Two-Sided”?

A platform is two-sided when:

- 1 There are *two distinct user groups*
- 2 Each group creates *value for the other group*
- 3 The value flows only *through the platform*, not directly

Classic examples:

- Credit cards: merchants and cardholders
- Uber: riders and drivers
- App stores: developers and users
- Insurance: the insured and (indirectly) capital providers
- Stock exchanges: buyers and sellers

The key test:

Can both sides transact without the platform?

If yes: the platform adds efficiency but is not truly two-sided.

If no: the platform creates value by coordinating the externality – and can capture part of it.

Two-sided market theory is the economics of intermediaries. Every exchange, every marketplace, every matchmaker is (at least partially) a two-sided platform.

Cross-Group Externalities: Why the Platform Has Pricing Power

Direct network effect (within one side): More users on the same side makes it more valuable. Example: more Visa cardholders \Rightarrow each cardholder has more people to pay.

Cross-group (indirect) network effect: More users on side A makes side B more valuable, and vice versa.

- More Visa merchants \Rightarrow Visa card is more useful
- More Visa cardholders \Rightarrow merchants want Visa terminal

The externality creates a subsidy logic: If adding one more merchant creates \$5 of value for cardholders, Stripe should be willing to pay up to \$5 to acquire that merchant.

Rochet-Tirole result:

The platform internalises the cross-group externality.

It does this by charging the side with lower externality more, and subsidising the side with higher externality.

“Charge the side that benefits most from adding the other side.”

Cross-group externalities are the source of the platform's pricing power: it can redistribute value between sides in ways that no single-product firm can replicate.

The Rochet-Tirole Pricing Rule

The key result (Rochet-Tirole, 2003; Tirole Nobel 2014):

$$p_A^* = c_A - \frac{\Delta v_{B \leftarrow A}}{\varepsilon_A}$$

Variable legend (read left to right):

where p_A^* is the platform's optimal price charged to side A,
 c_A is the marginal cost of serving one side-A user,
 $\Delta v_{B \leftarrow A} \equiv \partial V_B / \partial n_A$ is the cross-group externality
(the extra value one more side-A user creates for each side-B user),
and ε_A is the own-price elasticity of demand on side A.

Reading the rule:

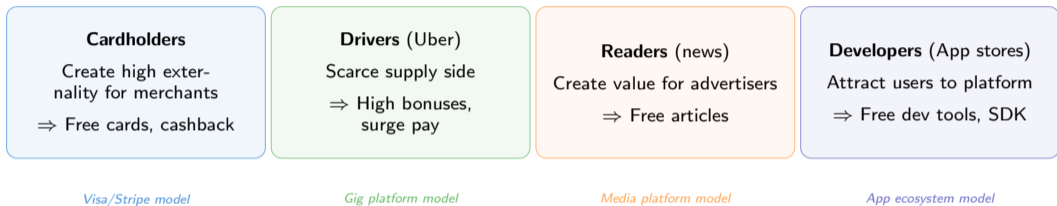
- If A creates *lots* of value for B ($\Delta v_{B \leftarrow A}$ large): subsidise A
- If A has *elastic* demand (ε_A large): subsidise A (they walk away easily)
- If A creates *little* value for B and has inelastic demand: charge A more

The pricing rule is counterintuitive: the side that benefits more from the platform pays less, not more. The platform charges the side that gains the most from keeping the other side on board.

Applied to Stripe:

- Merchants: inelastic (must accept cards), low externality per merchant for users
- Cardholders: elastic (many cards available), high externality per cardholder for merchants
- Result: charge merchants, subsidise cardholders

Who Gets Subsidised? The Four Cases



Every “free” product is a subsidy funded by the other side. When something is free, ask: which side is paying, and why did the platform choose this direction of subsidy?

The Chicken-and-Egg Bootstrap Problem

Every two-sided platform faces this at launch:

- Merchants need cardholders before accepting a new card
- Cardholders need merchant acceptance before wanting the card
- *Neither side joins without the other*

Four strategies to bootstrap:

- 1 **Subsidise both sides** initially (Uber: free rides + driver bonuses)
- 2 **Fake the supply side** (OpenTable: listed empty restaurants, gradually filled them)
- 3 **Solve for one side first** (Airbnb: listed on Craigslist to get renters first)
- 4 **Build a single-player tool first** (Excel \Rightarrow SharePoint; iMessage \Rightarrow FaceTime)

Alipay's bootstrap:

Taobao marketplace had distrust problem: buyers would not pay sellers they had never met.

Alipay launched as escrow: buyer pays Alipay, Alipay releases to seller after delivery.

Both sides joined immediately: problem was already there.

Successful platform bootstraps either already have one side (marketplace with an existing user base) or solve a single-player problem first. Cold-start chicken-and-egg is nearly always fatal.

Not All Network Effects Are Equal

Direct NE

More users of type A
⇒ more value for
other type-A users

WhatsApp, phone networks

$$V \sim n(n-1)$$

Indirect NE

More of side A
⇒ more value for
side B (two-sided)

Visa, Uber, App stores

$$V_A \sim n_B, V_B \sim n_A$$

Data NE

More users
⇒ better model
⇒ more users

Zhima Credit, Google, Ant

most durable

Key insight: Data NE are the most durable because they compound. A direct NE can be replicated by a competitor who pays for users. A data NE requires years of transactions to replicate.

When evaluating a platform's moat, identify which type of NE is operating. Data NEs create the widest competitive moats and are the hardest to regulate away.

Game Theory Primer: The Platform Launch Coordination Game

Setup: two players (a representative merchant and a representative user) must each decide whether to adopt a new payment platform. Both gain only if the other also adopts.

Payoffs:

- Both adopt: +3, +3 (high-coordination Nash)
- Neither adopts: 0, 0 (low-coordination Nash)
- One adopts, the other does not: -1, 0 (wasted effort)

Multiple Nash equilibria: both (adopt, adopt) and (skip, skip) are stable. Neither side wants to deviate alone.

Implication for fintech: this is why launches stall at zero, then suddenly tip. The equilibrium switch requires a coordinating push (subsidy, exclusivity, regulatory mandate).

	User: Adopt	User: Skip
M: Adopt	Nash 1 (+3, +3)	(-1, 0)
M: Skip	(0, -1)	Nash 2 (0, 0)

Two stable equilibria. Mechanism design picks which one obtains.

Two-sided platform launches are coordination games with multiple Nash equilibria. The chicken-and-egg problem is not irrationality; it is a stable low-coordination equilibrium that nobody can break alone.

Trust-mechanism failures are coordination games, not just information games.

March 2023 USDC depeg: Silicon Valley Bank collapsed; Circle held \$3.3B reserves there. USDC briefly traded as low as \$0.87.

- For an individual holder: redeem immediately (dominant strategy)
- For the collective: orderly redemption preserves the peg
- Individually rational behaviour triggers collective bank run

This is a textbook prisoner's dilemma. Each holder defects (redeems) because waiting is dominated. The Nash equilibrium is a run, even when no fundamentals justify one.

Schelling point in reserve design: 100% short T-bills + 24-hour redemption is a focal point that coordinates expectations. The mechanism is the reserve rule, not the dollar value.

Trust mechanisms must solve TWO games:

Information game:
prove the reserves are real (Akerlof).

Coordination game:
prove that other holders will not run (Schelling).

MiCA 2024 and the US GENIUS Act add coordination anchors (1:1 reserves, daily attestation) precisely to defeat the prisoner's dilemma.

Trust mechanisms must solve coordination games, not just information games. Reserve rules and daily attestation are Schelling points that prevent rational-individual behaviour from triggering collective runs.

A Schelling point is a focal solution that coordinates expectations without explicit communication. Protocol governance creates Schelling points by ritual: scheduled votes, public proposals, fixed quorums.

Two concrete examples:

- **MakerDAO MKR voting (2024-2026):** weekly executive votes are the Schelling point for stablecoin parameter changes. Holders know when to coordinate; opportunistic deviation is visible and punishable.
- **Uniswap fee switch (debated 2023-2025):** the UNI tokenholder vote on activating a protocol fee is itself the coordination device. The outcome (so far: deferred) is the equilibrium that holders coordinate on.

Why this generalises: every DAO governance vote is a public coordination signal. The vote count IS the equilibrium-selection mechanism.

Coordination economics is the bridge from mechanism design to platform competition. DAO governance, regulatory mandates, and exclusivity deals are different instruments serving the same function: select between multiple Nash equilibria.

Coordination economics bridges three pillars:

Pillar 2 (info): credit scoring reveals types.

Pillar 3 (platforms): platforms internalise cross-group externalities.

Coordination layer: governance votes pick between multiple stable equilibria.

Mechanism design IS equilibrium selection.

Tipping: When Markets Become Winner-Take-Most


Markets “tip” when:

- 1 Network effects are strong (NE slope is steep)
- 2 Multi-homing is costly (users stick to one platform)
- 3 Switching costs are high (historical data locked in)

The S-curve:

- Below critical mass: platform is unattractive, slow growth
- At tipping point: NE creates self-reinforcing growth
- Above tipping point: dominant platform captures most of market

Finance examples: Visa vs Amex in Europe; Alipay vs WeChat Pay in China; Uniswap vs SushiSwap in DeFi.



Winner-take-most does not mean winner-take-all. Alipay and WeChat Pay coexist because they entered the market together and each bootstrapped a different user base (Alibaba vs Tencent).

Case: Stripe vs PayPal – Different Sides, Different Subsidies

PayPal (1999 origin):

- Target: consumers making peer-to-peer payments
- Bootstrap: eBay buyers needed to pay sellers
- Subsidised side: *buyers* (free P2P, sign-up bonus)
- Revenue: sellers pay 2.9% + \$0.30 (same as Stripe)

Stripe (2010 origin):

- Target: developers building payment-enabled products
- Bootstrap: Y Combinator startup ecosystem
- Subsidised side: *developers* (7-line API, free sandbox)
- Revenue: merchants pay 2.9% + \$0.30

Same pricing, different logic:

Both charge 2.9%+\$0.30.
But PayPal subsidised the consumer side;
Stripe subsidised the builder side.

Result:

PayPal won consumer payments in 2000s.
Stripe won developer payments in 2010s.

Who wins the 2020s depends on who controls embedded finance.

Platform competition is competition over which side to subsidise. The winner is the platform that correctly identifies where the scarcest – and therefore most externality-generating – side is.

Pillar 4: Creative Destruction (Schumpeter, 1942)

Joseph Schumpeter's insight (1942): Capitalism does not improve industries gradually. It *destroys* them and rebuilds them.

Creative destruction: An innovation wave replaces an existing product category entirely.

- The car did not improve the horse-carriage industry – it ended it
- The ATM did not improve teller efficiency – it made most teller roles obsolete
- DeFi does not improve Barclays – it removes Barclays from specific functions

Key implication: The threat to incumbents is not better competition within the category. It is *elimination* of the category.

Why it matters:


Incumbents that improve their existing product (sustaining innovation) may still lose to a company that makes the product irrelevant.

Optimising a horse-carriage does not protect against Ford.

This is Christensen's innovation dilemma applied at the category level.

Schumpeter's insight: the correct unit of competitive analysis is the innovation wave, not the firm. By the time an incumbent responds, the wave has already reshaped the market.

Five Waves of Financial Innovation



Attacked: search TC

Attacked: verification TC

Attacked: underwriting TC

Each wave uses a new technology substrate to attack a different TC category. The waves do not replace each other – they stack.

We are in the middle of Wave 5. The question for strategy: which TC category remains high enough to build a business on? Hint: underwriting information asymmetry is still expensive.

Sustaining vs Disruptive Innovation (Christensen, 1997)


Clayton Christensen's distinction:

Sustaining innovation: Makes existing products better for existing customers. Incumbents are *good* at this.
Example: JPMorgan improving its mobile banking app.

Disruptive innovation: Starts as worse performance on *mainstream metrics* but offers new features (cheaper, simpler, more accessible) to a segment incumbents do not serve.

- Low-end: targets *overserved* customers
- New-market: targets *non-consumers*

Disruption crosses over when performance improves enough to serve mainstream customers, but by then the incumbent cannot respond.



Most fintech companies are disruptive, not sustaining: they start by serving customers that banks do not want (unbanked, micro-SME, low-volume traders), then move upmarket.

The Innovator's Dilemma: Why JPMorgan Cannot Build Revolut

Christensen's question: Why do successful companies consistently miss disruptive innovation even when they see it coming?

The answer: rational short-termism

- 1 Best customers are high-margin, demanding customers
- 2 Those customers want sustaining improvements, not disruptions
- 3 Resources flow to where returns are highest: the mainstream
- 4 The disruptive segment (unbanked, micro-SME) is unprofitable by current cost structures
- 5 By the time the disruptor improves, it is too late to respond

Why JPMorgan cannot build Revolut: A \$0/month account with no overdraft fees is incompatible with JPMorgan's cost structure, branch model, and compliance overhead.

The dilemma:

Every good manager at JPMorgan is doing the right thing for JPMorgan's current business model.

But collectively, those rational choices add up to missing the wave.

Christensen called this "the best-managed companies fail" paradox.

The innovator's dilemma predicts that incumbents will respond to disruption only when their core business is already threatened – by which point the disruptor has built a significant lead.

Three Disruption Types in Digital Finance

Low-End Disruption

Target: overserved customers paying too much

Example: Robinhood vs Merrill Lynch – free brokerage for retail traders willing to accept simplified service

New-Market Disruption

Target: non-consumers not served at all

Example: M-Pesa for unbanked Kenyans – first financial product for 40% of adult population

Ecosystem Disruption

Target: value chain, not a single firm

Example: Stripe Atlas – company formation + banking + payments as a one-API bundle

Most fintech unicorns combine types: Nubank is low-end (no fees) AND new-market (70M previously underbanked Brazilians). The most powerful disruptions attack multiple incumbent weaknesses simultaneously.

Ecosystem Innovation: The Platform-of-Platforms Model

The next level of disruption: Not just a better product, but a platform that enables *others* to build products.

Stripe's ecosystem:

- **Core:** payment processing API
- **Stripe Atlas:** company formation in 180 countries
- **Stripe Connect:** marketplace payment routing
- **Stripe Treasury:** embedded banking (fintech-in-a-box)
- **Stripe Radar:** ML fraud detection

The TC-reduction per layer: Each product eliminates one TC that blocked startups from offering financial services.

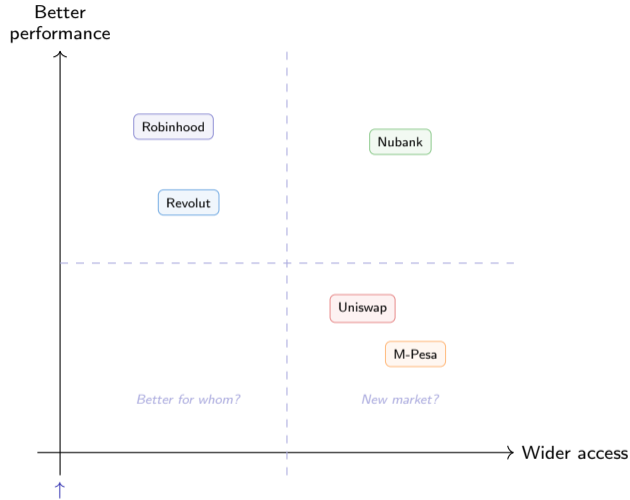
Plaid's model:

Plaid does not compete with banks. It makes fintechs possible by removing the TC of connecting to bank account data.

Stripe + Plaid together reduced the cost of starting a fintech from \$2M (compliance + dev + banking) to \$20K (API calls only).

Ecosystem innovators are TC-removers for other innovators. They compound disruption: by lowering the cost of innovation, they accelerate the pace of subsequent waves.

Class Exercise: Classify These Fintechs (5 minutes)



Discuss in pairs (4 min): Where would you place Stripe, Ant Group, and Aave on this matrix, and why?

There is no single right answer: the correct placement depends on which dimension of performance and which customer segment you take as the reference point.

Schumpeter's waves do not accelerate forever. Three forces slow each wave:

1. Regulatory capture

- Successful fintechs get banking licences
- Banking licences come with banking regulation
- Revolut in 2025 faces the same Basel III that slows JPMorgan

2. Incumbent NE moats

- Visa has 3.9B cardholders; no new entrant can replicate this instantly
- Data moats: Ant Group's 16 years of transaction data cannot be bought

3. Institutional inertia

- Employers pay salaries to bank accounts; employers are slow to switch
- Government transfers (welfare, pensions) locked into legacy systems

Timing is everything in disruption: enter too early (no regulation) and grow fast; enter too late (post-regulation) and face the same constraints as the incumbent you wanted to replace.


The paradox:

The more successful a fintech disruptor is, the more it becomes the thing it disrupted.

Robinhood now faces the same regulatory burden it disrupted. Alipay became so systemically important that it was regulated like a bank.

Disruption always contains the seeds of its own plateauing.

Synthesis: The Four Economic Villains of Digital Finance



⇒ M-Pesa, Wise, Aave



⇒ Zhima Credit, Plaid



⇒ Stripe, Uniswap, Visa



⇒ Revolut, Nubank

The rule: Every successful fintech solves at least one villain. The most powerful platforms solve two or more.

Map any fintech to its villains and you understand its strategy, its moat, and its regulatory exposure. Ant Group addresses all four: TC (payments), asymmetry (Zhima), platform (Alipay), incumbent (disrupted state banks).

The Central Thesis: Digital Finance Is Economics in Software

Traditional view: A payment app is a technology product. Success depends on UI, speed, and reliability.

Economic view: A payment app is an economic mechanism. Success depends on which TC it eliminates, which information asymmetry it resolves, and which network externality it activates.

Why this matters: The economic view predicts:

- Which features are defensible (vs copyable)
- Which pricing decisions are sustainable
- Which markets the platform can enter
- Where the regulatory pressure will come from

The economic lens:

Every product feature
is an answer to a market failure.

Every pricing decision
is a two-sided market choice.

Every business model
is a TC-reduction or
asymmetry-resolution play.

Every incumbent threat
is Schumpeter or Christensen
in action.

Technology determines what is possible. Economics determines what is valuable, what is defensible, and what will be disrupted. Build technology. Think economics.

Pillar 1: Transaction Costs

Key result: TC \rightarrow intermediary

Diagnostic: which TC does it eliminate?

Cases: M-Pesa (search/verify), Wise (FX)

Pillar 2: Info Asymmetry

Key result: adverse selection unravels markets

Signals (costly to fake) + screening menus

Cases: Zhima Credit, FICO score

Pillar 3: Two-Sided Markets

Key equation: $p_A^* = c_A - \Delta v_{B \leftarrow A} / \epsilon_A$

Subsidise high-externality side

Cases: Stripe vs PayPal, Visa

Pillar 4: Innovation Theory

Schumpeter: waves replace categories

Christensen: disrupt from below

Cases: Robinhood, M-Pesa, Stripe Atlas

These four frameworks are complementary. In Day 7B, we apply them simultaneously to Ant Group's rise and the platform competition dynamics that followed.

Synthesis Case: Ant Group Through All Four Lenses

TC Lens

Alipay eliminated verification TC via Taobao escrow

Asymmetry Lens

Zhima Credit reversed info gap via transaction data

Platform Lens

Two-sided market: merchants pay, consumers free

Innovation Lens

New-market disruption of state banks (unbanked 700M)

Ant Group: the largest financial institution never built on a banking licence. Four economic mechanisms, one super-app.

In Day 7B we ask: how did this become a problem for regulators, and what economic tools did they use to respond? The same frameworks that explain the rise also explain the fall.

Discussion: Apply the Four Lenses (5 minutes)

Choose one of these companies:

- Uniswap (decentralised exchange)
- Nubank (Brazilian neobank)
- Wise (international transfers)
- Stripe (developer payments)
- CBDC issued by European Central Bank

For your company, identify:

- 1 Which TC does it reduce?
- 2 Which information asymmetry does it address?
- 3 Is it a two-sided platform? If so, who is subsidised?
- 4 Is it disruptive (Christensen) or sustaining? What type?

Hint for Uniswap:

TC: eliminates market maker search cost and counterparty risk

Asymmetry: smart contract enforces terms without credit screening

Platform: two-sided (liquidity providers and traders)

Disruption: new-market (for non-bank traders globally)

This framework is portable: it works for traditional banks, fintechs, DeFi protocols, CBDCs, and insurance companies. One toolkit, every business model.

Rochet-Tirole pricing rule:

$$p_A^* = c_A - \frac{\Delta v_{B \leftarrow A}}{\epsilon_A}$$

Subsidise side A when A's joining creates high value for B and A is price-elastic.

Metcalfe's Law:

$$V \sim n(n-1) \approx n^2$$

Direct NE: value grows as square of users.

AMM preview (Day 7B):

$$x \cdot y = k$$

Constant product market maker. Price = y/x .

In Day 7B we add: Lerner index $L = (P - MC)/P$ for market power, the full AMM derivation, and Klemperer's switching cost taxonomy.

Today (7A): the toolkit

- ✓ Transaction cost economics (Coase)
- ✓ Information asymmetry (Akerlof)
- ✓ Two-sided market pricing (Rochet-Tirole)
- ✓ Innovation theory (Schumpeter, Christensen)

Next (7B): the application

- How do platforms compete? (switching costs, winner-take-all)
- How do AMMs work economically? ($x \cdot y = k$ derived)
- How is market power measured? (Lerner index)
- Ant Group's full story: rise, power, and regulatory response
- What should regulators do? (open discussion)

The thread:

The Ant Group case appears in Day 7A three times (slides 15, 17, 39).

In Day 7B it becomes the central case study – the full economic biography of the world's most interesting fintech.

Day 7A gives you the vocabulary. Day 7B tests whether you can read a real platform war using that vocabulary.

Investor: “I want to invest in the next Stripe. How do I know which startup has a real moat?”

Investor: “That sounds like an economics course, not a pitch deck.”

Analyst: “Simple. Which transaction cost does it eliminate? Which information asymmetry does it resolve? Which side of a two-sided market does it subsidise?”

Analyst: “Every durable moat in digital finance is an economics paper published before 2003. The pitch deck just describes how the paper is being implemented.”

The most powerful analytical tools in digital finance are not Excel models or Python scripts. They are the economic frameworks you now have. Use them first.

You can now:

- 1 Name the TC that any fintech eliminates and predict its pricing model
- 2 Identify adverse selection risk in a credit product and explain the screening mechanism used
- 3 Apply the Rochet-Tirole rule to determine which side of a platform should be subsidised
- 4 Classify any fintech innovation as sustaining, low-end disruptive, new-market, or ecosystem
- 5 Apply all four lenses simultaneously to Ant Group

The master thesis of Day 7:

“Digital finance is economics in software. Every design choice creates an incentive structure. Every incentive structure has a name.”

Tomorrow in Day 7B:
Platform Wars.

Economics is the theory. Software is the implementation. The economic lens predicts which implementations will survive and which will be disrupted – before the market tells you.