Strategic Toolkit and the Innovation Frontier

From Economic Theory to Strategy to 2030

Prof. Dr. Joerg Osterrieder

Digital Finance – Intensive Course: Day 7C

Day 7A named the economic failures. Day 7B applied them to platform power. Day 7C asks what playbooks reliably build (and disrupt) that power, and where the frontier moves next.

Same Playbook, Opposite Outcomes

Analyst A: "Ant Group started as a payments app in 2004. Ten years later, 1bn+ users, the world's largest fintech."

Analyst A: "Both used data, both went mobile-first, both expanded products. The economics looks identical. Why are the outcomes so different?"

Analyst B: "Nubank started as a credit card in Brazil in 2014. Ten years later, 100m+ users, the largest digital bank in Latin America."

Prof: "Same economics. Different strategy. Today we name the strategy choices that explain both."

Strategy frameworks isolate the design choices that explain why two firms with similar economics end up in very different places.

The Puzzle

Same era, different markets, both winners:

- Ant Group, 2004 to 2020: payments to credit to asset management. 1bn+ users, USD 37bn IPO blocked by regulators.
- Nubank, 2014 to 2024: credit card to digital bank to investments. 100m+ users, public listing, profitable from 2023.

Both apply the same 7B forces: switching costs, network effects, data flywheel. Both built moats. Both are still standing.

Yet:

- Different revenue mix (Ant: data and credit; Nubank: interchange and cards)
- Different regulator outcome (Ant restructured; Nubank licensed)
- Different durability source

The question:

Same economic forces.
Why different strategy?

Strategy frameworks are the missing layer between 7A theory and 7B platform power.

We need 3 strategy frameworks, 2 diffusion models, and 2 innovation lenses to read both cases.

Two firms with the same 7B forces can land in opposite places because their strategy choices (who to serve, what to monetise, how to sequence) differ. 7C names those choices.

Day 7C: Four Sections

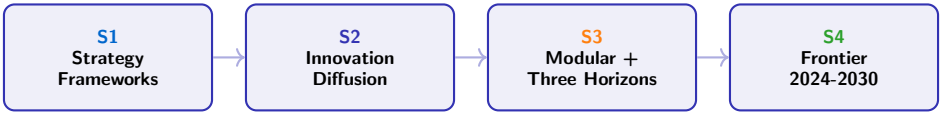

Seven new frameworks today. Each built from a concrete example before the formal definition. Each applied to Ant Group, with Nubank as comparator.

Frameworks are layers, not rivals. Porter for industry, BMC for firm, Helmer for durability. Rogers and Bass for adoption. Henderson-Clark and Three Horizons for innovation. Each answers a different question.

Bridge from 7B

7B closed with: platform power is measurable.

Lerner index, switching cost taxonomy, data flywheel speed, four barrier types, mechanism failure modes.

7C asks two follow-up questions:

1. If power is measurable, what playbooks *reliably build* it?
2. What playbooks *predict its disruption*?

Theory (7A) tells you the failure. Applied (7B) tells you the moat. Strategy (7C) tells you the move.

Theory, applied, strategy: the natural completion of the day. 7C is the synthesis layer that turns what we know into what we do.

"Strategy is the design of mechanisms, not the choice of moves."

Three frameworks today, each answers a different question:

Porter (1979): where is profit possible in this industry?

BMC (Osterwalder, 2010): does our firm cohere across 9 boxes?

Helmer (2016) – 7 Powers: is our moat actually durable?

Each built from scratch. No prior strategy course assumed.

Strategy is mechanism design with a budget constraint. The same forces 7A and 7B taught are now design parameters: switching costs, network effects, IO barriers are levers, not just observations.

Porter's 5 Forces: Why Some Industries Are More Profitable

Motivating puzzle:
The EU payments market is hyper-competitive (Stripe, Adyen, Mollie, Klarna). The Chinese mobile payments market is a duopoly (Alipay, WeChat Pay). Same product. Same era. Why?

Porter's claim (1979):
Industry structure determines profitability before any firm makes a single move. Five forces act on every industry simultaneously.

Strategy first asks: which industry are we in? Strategy second asks: which forces can we change?

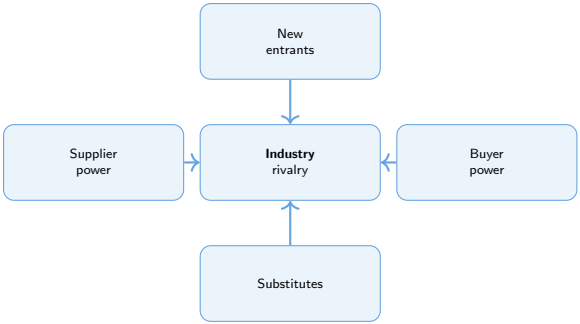

Industry structure determines profitability before any firm makes a move. Strategy is the choice of which forces to attack, which to accept, and which to redesign.

Porter Applied: Ant Group, China Payments, 2018

Five forces, scored for Ant in 2018:

- **Rivalry:** LOW. Only WeChat Pay competes at scale. Effectively a duopoly.
- **New entrants:** LOW. Regulatory licence required; network effects already tipped.
- **Supplier power (cloud, banks):** MEDIUM. Alibaba Cloud captive; partner banks fragmented.
- **Buyer power (merchants, users):** LOW. Merchants must accept Alipay; users single-home.
- **Substitutes:** LOW in 2018, RISING. DeFi protocols and CBDC (e-CNY) emerging.

Verdict: four forces favourable, one watch-list. Highly profitable industry position.

Digital adaptation:

Porter's substitutes arrow no longer points at competing products. It points at competing *protocols*: DeFi replaces lenders, CBDC replaces stablecoins, AI agents replace advisors.

Substitution is now protocol-level, not product-level.

Porter still works in digital finance; the substitutes arrow now points to protocols and infrastructure, not consumer products. Read Porter through the 7B platform lens for full diagnostic power.

Business Model Canvas: 9 Boxes on a Napkin

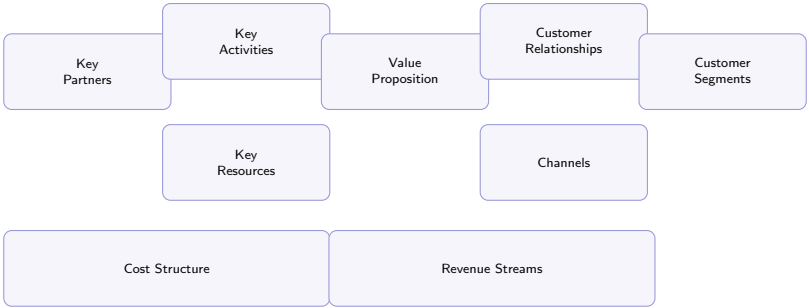

Osterwalder (2010): a business model is a coherence test. The 9 boxes must hang together. Change one box and the others must adjust, or the model breaks.

A business model is a coherence test. The 9 boxes must hang together. A regulator can break a business model by attacking any one of them; understand which.

BMC: Ant 2020 vs Ant 2026 (Restructuring Effects)

Ant 2020 (pre-IPO halt):

- **Customer segments:** 1bn+ users, 100m+ SME merchants
- **Value prop:** payments + credit + wealth + insurance in one app
- **Revenue:** commission, credit interest, fund management, insurance
- **Key resources:** Sesame Credit data, Alipay graph
- **Key partners:** Alibaba, Yu'e Bao funds, regional banks

Ant 2026 (post-restructuring):

- **Customer segments:** unchanged
- **Value prop:** payments + regulated wealth (credit walled off into subsidiary)
- **Revenue:** commission UP weighting; credit interest DOWN (subsidiary capital cost)
- **Key resources:** Sesame Credit data shared with PBOC bureau
- **Key partners:** +PBOC credit bureau; same banks; voting control transferred (Jack Ma trust)

Boxes changed: revenue model, key resources, key partners, value proposition. Customer-facing boxes mostly stable.

A regulator can break a business model by attacking any of 9 boxes. The Ant restructuring targeted three boxes simultaneously: resources (data sharing), partners (bureau), revenue (capital requirements on credit).

BMC Comparison: Ant 2024 vs Nubank 2024

Ant 2024 (China):

- Segments: super-app cross-sell to existing 1bn users
- Value prop: integrated payments + wealth + insurance in one app
- Revenue: commission heavy; credit walled off
- Key resources: transaction data, Sesame Credit (now shared)
- Regulatory posture: license-after-scale, restructured

Nubank 2024 (Brazil + LatAm):

- Segments: underbanked first, then upsell to banked middle class
- Value prop: no-fee credit card + simple digital banking
- Revenue: interchange + card interest + cross-sell investments
- Key resources: brand trust, low-cost servicing, low CAC via referral
- Regulatory posture: license-first, full bank charter

Same industry. Different canvas. Different durability source. Different regulator outcome.

Same industry, different canvas, different outcome. Nubank chose license-first and underbanked-first; Ant chose scale-first and license-after. Strategy choices reshape every other box.

Helmer's 7 Powers: What Makes a Moat Durable?

Motivating question (Buffett):
Many businesses have a moat. Few moats stay wide.
What makes a moat *durable*?

Helmer (2016) named 7 sources:

1. Scale Economies
2. Network Economies
3. Counter-Positioning
4. Switching Costs
5. Branding
6. Cornered Resource
7. Process Power

Each is a specific economic source of advantage. "Moat" alone is too vague.

1. Scale: avg cost falls with volume

2. Network: value rises with users

3. Counter-position: incumbent cannot copy

4. Switching: cost to leave

5. Branding: trust premium

6. Cornered res.: unique input

7. Process: tacit operational know-how

"Moat" is vague. Helmer's 7 Powers names the specific economic source of durable advantage. A company can hold multiple Powers; the rare ones survive a Power transfer when regulators or technology shift.

Helmer Scorecard: Ant 2020 vs Ant 2026

Power transfer under restructuring:

- **Scale:** HIGH in both eras (cost per transaction near zero)
- **Network:** HIGH (2020), MEDIUM (2026, after interoperability mandates)
- **Counter-positioning:** HIGH (2020), LOW (2026, banks now offered super-app features)
- **Switching:** HIGH in both (user habit + merchant QR codes)
- **Branding:** HIGH in both
- **Cornered resource:** HIGH (2020, Sesame data), MEDIUM (2026, shared with PBOC)
- **Process power:** MEDIUM in both

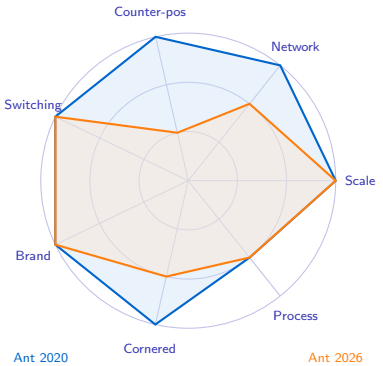

Restructuring is a Power transfer: regulators redistribute Powers from one party to another. Ant kept scale, switching, brand; lost half of network, counter-positioning, and cornered resource.

Section 1 Summary: Three Frameworks, Three Questions

When to reach for which:

Porter Where to play. *Which industry? Which forces?*

BMC How to play. *Do our 9 boxes cohere?*

Helmer How durable. *Is our moat actually deep?*

Frameworks are layers, not rivals. Use all three on any digital finance firm.

Ant scorecard recap:

- Porter (2018): four forces favourable, substitutes rising
- BMC (2020 to 2026): three boxes changed by regulator (revenue, resources, partners)
- Helmer: lost half of three Powers (network, counter-pos, cornered)

Nubank chose differently on every layer; outcomes differ accordingly.

Frameworks are not rivals; they are layers. Use Porter for industry, BMC for firm, Helmer for durability. Each answers a question the others ignore.

Frameworks tell you WHAT to do.

Diffusion theory tells you WHO you do it FOR, and WHEN.

Two tools today:

Rogers (1962): adoption proceeds in 5 named segments.

Bass (1969): adoption rate is governed by 2 parameters.

Together they describe the entire life cycle of a new financial product, from launch to saturation, in two simple equations.

Adoption is a process, not an event. Two-thirds of fintech launch budgets burn on the wrong segment because the founder mistook innovators for early majority.

Rogers' 5 Adoption Segments

Five segments, fixed proportions:

- Innovators: 2.5% (try anything)
- Early adopters: 13.5% (opinion leaders)
- Early majority: 34% (pragmatists)
- Late majority: 34% (sceptics)
- Laggards: 16% (resistance)

Crypto launches always look the same: 2.5% pump, 13.5% notice, then 84% never adopt unless the product crosses to early majority.

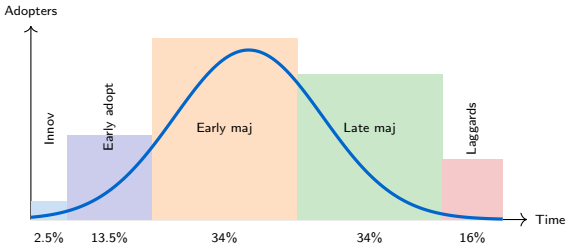

Adoption follows a predictable distribution. Ignore the segment your product is currently in and the launch budget burns on the wrong audience.

Rogers Applied: Four Fintech Adoption Stories

Where each company sits today:

- **M-Pesa (Kenya, since 2007):** late majority. Started with unbanked early adopters; crossed early majority around 2010 via agent network; near saturation today.
- **Robinhood (US, since 2014):** stuck at early majority. Regulatory backlash (PFOF, GameStop) plus brand damage stalled mass-market crossing.
- **Nubank (Brazil, since 2014):** late majority. Crossed early majority via referral plus card waiting list; now broad bank-account demographics.
- **Revolut (UK, since 2015):** early majority. Strong with millennial knowledge workers; still works to cross into older mass market.

Per-segment product:

Innovators want: novelty, bragging rights, raw feature access

Early adopters want: smooth UX, professional application

Early majority wants: proven scale, low switching risk

Each segment needs a different value prop. Product-market fit is per-segment, not global.

Each segment requires a different value proposition. Product-market fit is per-segment, not global. Most fintech failures are not at idea stage but at segment crossing.

The Bass Model: Two Parameters, Whole Curve

Bass (1969): model the rate of adoption with just two parameters.

$$\frac{dN}{dt} = \left(p + q\,\frac{N}{m}\right)(m - N)$$

Read this verbally:

- p = innovation push (ads, awareness)
- q = imitation pull (word-of-mouth)
- m = total addressable market
- $N(t)$ = adopters at time t

Rate of adoption = (innovation push) + (word-of-mouth pull), capped by remaining market.

Why this matters:

Just two parameters p and q fit a whole adoption curve.

Once p and q are estimated from 2-3 years of data, you can project the remaining curve.

The closed-form solution lives in notebook 7g; on this slide we keep the rate form because it makes the mechanics legible.

Bass: two parameters fit the entire curve. Innovation push (p) is what you spend on ads. Imitation pull (q) is what users do to other users. The ratio $p : q$ is the strategic input.

Bass Fitted to Nubank, 2014-2024

Nubank customer counts (millions):

- 2014: 0.05
- 2017: 3
- 2019: 15
- 2021: 48
- 2023: 85
- 2024: 105

Fitted Bass parameters (LatAm market $m \approx$ 350m):

- $p \approx 0.004$ (low innovation push, mostly viral)
- $q \approx 0.55$ (very strong word-of-mouth)
- $p : q$ ratio about 1:140

2027 forecast band: 145-170m.

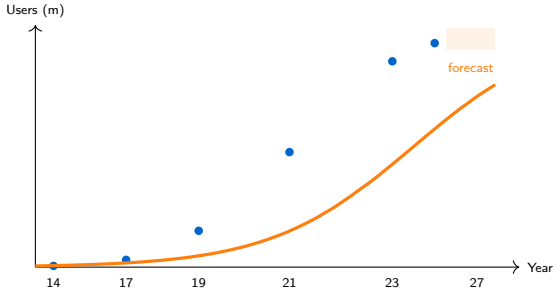

Bass forecasting is the most concrete strategy tool you have for a growth-stage fintech. The Nubank curve is dominated by q; viral referral is the engine, not paid acquisition.

Bass Comparison: Alipay vs Nubank

Alipay 2004-2018 (rough fit):

- $p \approx 0.03$ (high innovation push via Taobao integration)
- $q \approx 0.40$ (strong word-of-mouth, but p dominates early years)
- $p : q$ about 1:13

Nubank 2014-2024:

- $p \approx 0.004$, $q \approx 0.55$
- $p : q$ about 1:140

Two adoption curves, two very different growth engines. Alipay rode an e-commerce platform; Nubank rode a referral network.

Strategic reading:

High p, low q: invest in distribution, partnerships, advertising.

Low p, high q: invest in referral, NPS, community.

Both Ant and Nubank "won" but the optimal marketing playbook is opposite.

Diagnose the $p : q$ ratio before you set marketing budget.

Different markets have different $p : q$ ratios; the ratio is the strategic input. Same product can need opposite marketing playbooks across geographies.

Two lenses on how innovation actually happens inside firms:

Henderson-Clark (1990) – Architectural innovation:
a 2x2 that distinguishes 4 innovation types
by what changes: components, architecture, or both.

McKinsey Three Horizons (2000) – Temporal innovation:
H1 extend core, H2 emerging, H3 viable options.

Together: how innovation moves through an organisation,
and why incumbents miss the architectural shifts.

Innovation is organisational, not just technological. The hardest innovation to defend against is architectural; the hardest to fund is H3.

Henderson-Clark: Four Types of Innovation

The Henderson-Clark 2x2 (1990):

Two axes: what changes about *components*, and what changes about the *architecture* that connects them.

Why incumbents miss architectural innovation:

- They own all the components
- They have deep R&D in each component
- They miss that competitors are re-wiring the SAME components into a NEW architecture
- By the time they notice, the new product is mainstream

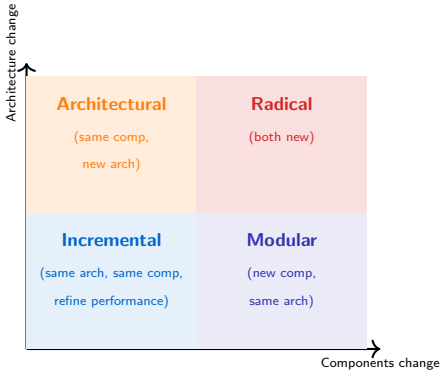

Architectural innovation re-wires the SAME components; incumbents stay focused on component-level R&D and lose the architecture. This is the deepest version of Christensen.

Henderson-Clark Applied: JPMorgan vs Revolut

Same components, different architecture:

- **KYC engine:** JPMorgan has one, Revolut has one
- **Payment rails:** both
- **Card issuance:** both
- **Custody:** both
- **FX engine:** both
- **Lending engine:** both

Revolut's architectural move:

- Mobile-first single app (vs JPM's siloed apps)
- Real-time FX inside payments (vs JPM's settle-and-quote)
- Universal account (vs JPM's product-by-product)

The architectural innovation test:

Did the new entrant invent new components, or re-wire the same components into a new architecture?

Revolut: same components, new architecture. JPM cannot copy the architecture without breaking its own product silos.

Architecture is the deeper moat.

The hardest innovation to defend against is architectural; the new entrant uses the same building blocks but rearranges them. Incumbents lose because reorganisation is more painful than R&D.

McKinsey Three Horizons

Why do banks fail at fintech investments?
Because their investment portfolio is over-weighted to H1 and starved on H2 and H3.

Three Horizons (Baghai-Coley-White, 2000):

- **H1:** extend the existing core
- **H2:** build emerging businesses
- **H3:** create viable options for tomorrow

Ideal allocation: approximately 70 / 20 / 10.
Real allocation at most banks: closer to 95 / 4 / 1.

Three Horizons translates the innovator's dilemma into budget terms.

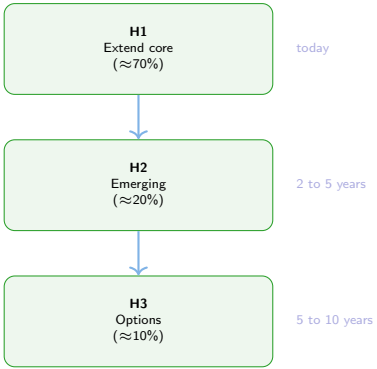

Three Horizons reveals incumbent under-investment in H2 and H3; this is Christensen's dilemma stated in budget form. Banks defend H1 perfectly and starve the future.

Three Horizons Applied: JPM, Revolut, DeFi

Portfolio mix today:

- **JPMorgan:** mostly H1 (Onyx, JPM Coin, Payments). Tokenisation pilots are H1 in disguise (extending custody). Limited H2; almost no H3 outside Quorum.
- **Revolut:** natively H2. Wealth, crypto, robo-advisory, business accounts. H3 ambition through trading platform and superapp positioning. Almost no H1 legacy to defend.
- **DeFi protocols:** pure H3 by construction. No legacy to defend; the whole business is the option.

Open vs closed innovation: H3 plays tend to be *open* (Aave grants programmes, Visa Innovation Center). H1 plays tend to be *closed* (JPMorgan internal R&D).

H2 is the gap:

Most fintech disruption happens in the H2 gap that incumbents leave open.

Open innovation (Aave grants, sandbox programmes) is the H3 norm; closed innovation is the H1 norm.

The boundary choice (open vs closed) is within each horizon, not above them.

The H2 gap is where most fintech disruption happens; open vs closed is the boundary choice within each horizon. Banks lose because they leave H2 to startups and DeFi.

Christensen's law: profit moves to the bottleneck, which moves with the architecture.

The DeFi stack is modular by design:

- L1 (Ethereum, Solana): infrastructure
- L2 (Optimism, Arbitrum): scaling
- DEX (Uniswap, Curve): venue
- Aggregator (1inch, Paraswap): routing

Each layer:

- Subject to its own Bass dynamics (different p and q)
- Governed by different Helmer Powers
- Sits at a different McKinsey Horizon

Multi-S-curve view:

Treat each layer as its own S-curve. Adoption timing differs by layer.

Profit migration: as a layer commoditises, profit moves up or down the stack to the next bottleneck.

Modular stacks are not one product; they are a portfolio of Bass curves, each at a different horizon.

A modular stack is a multi-S-curve; treat it as such, not as one product. The right strategy for L1 differs from the right strategy for an aggregator at the top of the stack.

The 2030 forecast follows from 2026 facts. Not prediction. Trajectory analysis: which economic forces are currently building, and where does their logic lead before a regulator or rival intervenes?

The 2030 forecast follows from 2026 facts. Four anchors today: AI, tokenisation, CBDC, regulatory frontier. Each carries the seven Day 7 lenses forward.

AI in Finance 2024-2026: Where We Actually Are

Already in production, not vision:

- **LLM credit underwriting:** Upstart (USD 7bn+ loans originated using ML), Klarna 2024 disclosures showing GenAI in customer service plus risk scoring.
- **Robo-advisory at scale:** Wealthfront and Betterment manage tens of billions; Vanguard Personal Advisor hybrid model dominates.
- **Agentic finance pilots:** JPMorgan IndexGPT, Bloomberg AI plus Anthropic vertical model, Stripe Atlas plus LLM onboarding.
- **Generative documents:** regulatory filings, KYC summaries, RFP responses.

The open question is regulatory rate, not technology rate.

Where 2026 differs from 2023:

Open-source models (Llama 3, Mistral, DeepSeek) collapse the price of inference.

Vertical fine-tunes replace generic LLMs for finance use.

Agent frameworks (OpenAI Assistants, Anthropic MCP) commoditise tool use.

Model layer is racing toward commodity.

AI in finance is already in production; the open question is regulatory rate, not technology rate. Most LLM use is internal automation, not external product.

AI in Finance: Where Does the Value Accrue?

Stack of value capture:

- **Model layer (OpenAI, Anthropic):** high R&D cost, low margins per call; commoditising.
- **Vertical model layer (Bloomberg + Anthropic, JPM Athena):** domain data is the moat.
- **Distribution + data layer (banks, exchanges, neobanks):** customer relationship and licence portfolio.
- **Orchestration layer (startups, copilots):** workflow integration.

Strategic question: which layer captures durable margin?

Helmer applied to AI in finance:

Cornered resource? At the model layer: thin (open-source catching up). At the distribution layer: strong (regulated deposit base).

Counter-positioning? Open-source models counter-position closed models on price.

Most fintech AI value will accrue at the orchestration + distribution layer, not the model layer.

Most fintech AI value will accrue at the orchestration and distribution layer, not the model layer. Models are commoditising; distribution licences and proprietary data are not.

Tokenization 1: Stablecoin Reserves as Monetary Mechanism

Major stablecoins, 2026 reserve mix (approximate):

- **USDC (Circle):** 80%+ short T-bills, 15% cash, residual repo
- **USDT (Tether):** mixed T-bills, secured loans, gold, BTC reserves
- **PYUSD (PayPal, Paxos):** short T-bills and cash
- **Tokenized money market funds (BUIDL, OnChain USDY):** short Treasuries

The monetary lever:

- Reserve mix determines monetary aggregate impact
- T-bill heavy = inert to commercial bank reserves
- Cash heavy = drains M2 directly
- Run risk depends on liquidity of reserve mix

Scale, 2026:

USDC: about $50bn
USDT: about $120bn
PYUSD: about $1bn
Tokenised MMFs: BlackRock BUIDL about $1.7bn; Ondo and Franklin Templeton scaling.

Stablecoins are private monetary instruments whose reserve choice is a de facto monetary policy.

Stablecoins are monetary policy in private hands; reserve composition is the monetary lever. A USDC reserve shift toward cash would reshape commercial bank liquidity.

Tokenization 2: CBDC vs Commercial Money

The issuer question: a CBDC is a central-bank liability; commercial bank money is a network of bank IOUs.

Trade-offs:

- Settlement risk: lower with CBDC (single issuer)
- Disintermediation risk: higher for commercial banks (deposits leak to central bank)
- Privacy: CBDC creates a complete payment record at the central bank
- Programmability: both can be coded, but rule-setting differs

Where each region stands 2026:

- China e-CNY: largest rollout, USD 250bn+ cumulative volume
- ECB digital euro: preparation phase 2025-2026, design choices on hold-limits
- UK BoE: design phase, decision deferred
- US Fed: paused; political risk dominant

Why progress is uneven:

Each jurisdiction weighs disintermediation risk against payment system efficiency.

China: state already owns big banks; disintermediation risk is internal.

EU: 2000+ banks; disintermediation risk is systemic.

Do not extrapolate from one region to another.

CBDC progress is uneven; do not extrapolate from one region to another. The same technology has very different strategic implications in different financial systems.

Tokenization 3: Monetary Policy Transmission

The new transmission channel:
When stablecoins and tokenized funds hold T-bills, central bank repo operations affect their yields, which affects their adoption, which affects deposit migration.

Examples in 2026:

- **BlackRock BUIDL** (about $1.7bn): tokenized money-market fund holding short Treasuries
- **Ondo USDY:** retail-facing tokenized Treasury exposure
- **Franklin Templeton OnChain US Government Money Fund:** on-chain shares
- **JPMorgan Onyx Digital Assets:** bank-issued tokenized collateral

The Fed has a stablecoin policy whether it wants one or not.

Frontier convergence:

Monetary mechanism (7B Section 3) and tokenisation (S4 here) are the same problem from two angles:

Who issues, who holds, what backs it, and what flows through the rate corridor?

The frontier converges; do not study these in separate silos.

The frontier converges; monetary mechanism and tokenization are the same problem from two angles. The Fed already prices stablecoin reserves indirectly through repo and T-bill yields.

The Regulatory Frontier: Consolidation, 2024-2026

Active regimes today:

- **EU MiCA:** stablecoin reserves, custody, market abuse; fully phased in 2024-2026
- **EU DMA:** platform interoperability and gatekeeper obligations
- **Draft EU PSD3:** open finance extension, instant payments mandate
- **US GENIUS Act (Stablecoin):** federal stablecoin framework
- **UK FCA crypto perimeter:** authorisation regime for crypto firms
- **MAS, FINMA:** stablecoin and tokenisation guidelines

The pattern: convergence on four planks.

- Capital
- Custody
- Disclosure
- Interoperability

Implication for strategy:

Regulatory arbitrage windows are closing.

A 2024 strategy that relies on offshore domicile or unlicensed custody has a 2026-2027 expiry date.

2025-2026 is the consolidation phase; the rules of 2030 are being written now.

2025-2026 is the regulatory consolidation phase; the rules of 2030 are being written now. Strategies that assume the regulatory ambiguity of 2018 misread the moment.

Forecasting Exercise + Section 4 Summary

Class exercise (5 minutes):
Pick one: Ant 2030, Nubank 2030, JPMorgan 2030, Stripe 2030.

Project that company onto the 7C toolkit:

1. Porter: which force shifts most?
2. BMC: which box changes?
3. Bass: where is the curve in 2030?
4. Three Horizons: what mix?

No single right answer. A good forecast names the dependencies, not the outcome.

What forecasting is NOT:

It is not prediction.
It is structured uncertainty management.

Good forecasters list:
what must hold for this story to be true,
what would falsify it,
and which signals come first.

Apply the 7C toolkit the same way to any 2030 question.

Forecasting is not prediction; it is structured uncertainty management. The 7C toolkit lets you project any company onto a set of named dependencies that you can monitor.

Disruptors of Disruptors

Christensen recursion: every 2026 winner is on the 2030 disruption list.

- **Stripe** (today: dominant payments infra) is being disrupted by **embedded AI agents** that route payments without an explicit dashboard
- **Stablecoins** (today: private monetary instruments) are being disrupted by **CBDCs** (central bank rails)
- **Nubank** (today: digital-bank disruptor) is being disrupted by **AI underwriting** that lets any e-commerce platform issue cards
- **JPMorgan Onyx** (today: tokenised collateral) is being disrupted by **public-chain settlement** once regulation allows

The recursion test:

For each current winner, ask:

Which 7A pillar did they ride?

Which 7B force protects them?

Which 7C framework names the playbook their challenger will use?

Every winner of 2026 is on the disruption list for 2030.

Every winner of 2026 is on the disruption list for 2030. Christensen's recursion never ends; only the names of the winners and disruptors change.

The Day 7 Grand Toolkit

Day 7A (theory):

- Transaction costs (Coase)
- Information asymmetry (Akerlof)
- Two-sided markets (Rochet-Tirole)
- Innovation theory (Schumpeter, Christensen)
- Behavioral finance, game theory, welfare

Day 7B (applied):

- Switching costs, network effects, data flywheel
- Lerner index and IO barriers
- Mechanism design failure modes

Day 7C (strategy):

- Porter 5 Forces
- Business Model Canvas
- Helmer 7 Powers
- Rogers diffusion segments
- Bass model (rate form)
- Henderson-Clark architectural
- McKinsey Three Horizons
- Frontier 2024-2030 anchors

Memorise three. Reach for the right one instinctively for everything else.

Day 7 gives you a toolkit. Memorise three; reach for the right one instinctively. The lenses are layered, not competing; together they describe any digital finance situation.

Same Playbook, Opposite Outcomes – Resolved

Analyst A: "So Ant and Nubank really did use the same economic forces."

Analyst A: "So strategy IS the variable that explains the divergence."

Analyst B: "Different Porter scoring, different BMC, different Helmer mix, different Bass $p : q$, different Horizon allocation."

Prof: "Same economics. Different strategy. Two winners. Two playbooks. Now you can read both."

Same economics, different strategy. Two winners, two playbooks. The strategy frameworks isolate what differs and let you read each case on its own terms.

Day 7 Synthesis: Engineering with Money

"Digital finance is engineering with money."

Economics tells you what is possible. (7A)
Strategy tells you what is profitable. (7B and 7C)
Innovation tells you what is coming. (7C)

Three layers. One day. Sixteen lenses on the wall behind every digital finance decision you will analyse in your career.

Apply them to every pitch deck, every regulatory filing, every news headline. The lenses do not go stale; the data they read does.

Digital finance is engineering with money. Economics tells you what is possible. Strategy tells you what is profitable. Innovation tells you what is coming. Apply the toolkit. It does not go stale.