


CBDC Business Models: Public Money, Private Rails

The cash is public, the rails are private — the business model lives in between

Digital Finance


The Mint



"The state mints the money. Somebody else mints the moat."

vs.

The Toll Road



Why Does a Public-Money Project Need a Private-Sector Business Model at All?


The Two-Tier Necessity

Central banks issue the money (the legal-tender claim), but they have no retail storefront, no onboarding desk, no support line, and no incentive to build one. A public-good claim still has to be delivered to wallets, phones, and merchant tills by somebody — and that carrying is a service. In business-model language, *friction* = any cost, delay, or hassle between a customer and the service they want.

The answer almost every CBDC programme reaches is structural rather than ideological: a two-tier model where the issuer mints, and licensed intermediaries distribute, custody, and service. Those intermediaries — payment service providers (PSPs), anything from a licensed fintech wallet to a retail bank acting as distributor — own the customer relationship. The intermediation function is unbundled from the money function.

- **Two tiers:** tier one = central bank issues the money; tier two = licensed intermediaries (PSPs, commercial banks) distribute it. This mini-lecture uses “two-tier” in this sense throughout.
- The cash itself is a public good (*public good* here = the unit of account is free at the point of use and non-discriminatory — anyone eligible gets the same claim).
- The rails carrying it accrue costs that somebody must absorb.
- Whoever runs the wallet sees every transaction and earns every fee — in BM language, a *moat* = a competitive advantage that rivals cannot easily copy.
- The friction the BM solves is not money creation — it is reach.

Core tension: the issuer is the central bank, but the commercial BM lives at the PSP layer — so who captures value and who absorbs cost is decided entirely by the license and fee architecture, not by monetary policy.




How Does mBridge Turn Cross-Border Settlement Into a Two-Sided Platform?

The mBridge Case

The pain point: cross-border payments have always been intermediated by chains of correspondent banks (a *correspondent bank* = a middleman bank holding an account in another country on your bank's behalf; a single wire can go through three to five of these before landing) — each link adds a fee, a delay, and a reconciliation risk. mBridge is a multi-central-bank DLT-based (distributed-ledger) wholesale settlement prototype that collapses that chain by letting participating central banks settle directly on a shared ledger.

A *two-sided platform* = a service that sells access to one user group (e.g. commercial banks) by first attracting another (e.g. central banks). Each side wants the other to show up first — the “chicken-and-egg” problem.

- **Multi-sided platform:** central banks on one side, commercial banks on the other, corridor pairs underneath (a *corridor pair* = a single pairing like US-dollar-to-HK-dollar — each pair is a separate route with its own volume, fee, regulation).
- **Network effects:** each new central bank that joins makes the platform more useful to every commercial bank that needs a corridor in or out of that jurisdiction.
- **Chicken-and-egg:** the platform launched with a small founding bloc (HKMA, PBoC, Bank of Thailand, Central Bank of the UAE — four central banks plus BIS Innovation Hub — committed settlement liquidity across a handful of pre-selected corridor pairs at launch, so the first commercial banks who joined saw working corridors immediately), rather than trying to onboard everyone at once.
- The result: a settlement venue that competes with the correspondent-bank chain by removing intermediary layers entirely.



Platform economics applied to public infrastructure: the more central banks on one side, the more corridors available — and the more useful the settlement venue becomes.

How Does e-CNY Start Narrow as a Wallet and End Broad as a Platform?

Clayton Christensen (Harvard Business School) argued disruptors start narrow and cheap, earn trust, then expand upward.

Christensen's Disruption Cycle Applied

Unbundling = pulling one service out of a historical bundle and offering it alone;
rebundling = stacking adjacent services onto that foothold once trust is established.

Phase One — Unbundling: the e-CNY pilot launched as a single proposition — a digital cash wallet for a handful of pilot cities — Shenzhen, Suzhou, Chengdu, plus Winter-Olympics venues — during the early rollout. No lending, no investments, no merchant tools. Just a way to hold and spend a digital claim on the central bank.


Phase Two — Trust Earned: pilot users spent the wallet balance for transit, retail, and small payments. Familiarity built gradually as more merchants accepted it and more cities were added to the rollout. The PBoC's design choice (free-at-the-point-of-use, with no transaction fee passed to either merchant or citizen) forced Alipay and WeChat Pay into a defensive posture — they had to keep their rails open to e-CNY transfers or lose transit and government-subsidy flow.

Phase Three — Rebundling: merchant SDKs (software development kits — a bundle of code a merchant integrates to accept the wallet), programmable (the money carries its own rule-code — expiration dates, merchant categories, auto-disbursement conditions — enforced at the rail layer, not by a human at a bank desk) red-envelope subsidies (China's digital lucky-money — e-CNY made the subsidy itself programmable, so a city could deliver stimulus that had to be spent within a week at local merchants), payroll integration, vending integration, and government disbursement rails. The wallet that began as digital cash becomes a programmable distribution platform for any state-issued or merchant-issued payment instruction.

The irony: a public-goods CBDC ends up rebundling more than a commercial wallet

ever could, because programmability is itself a state instrument.

Christensen's disruption arc applies even to public-good projects: e-CNY entered as digital cash, earned habit, then rebundled into a programmable distribution platform.




Where in the Issuance-to-Circulation Stack Does Value Actually Accrue?

The Three-Tier Value Chain (Drex Lens)

Value chain = the ordered sequence of activities a service passes through (issue → distribute → service → settle). Evans and Wurster (BCG) argued that when information is cheap, each link can be split off to a specialist — the chain *deconstructs*. In a CBDC, the split is vertical: top = issuer (no BM), middle = PSP (the only commercial tier), bottom = user (the base, not a separate BM layer).

- **Issuer — Central Bank:** mints, supervises. Zero revenue. Pure cost centre.
- **Intermediary — PSPs:** custody, KYC, wallet, VAS. The commercial BM lives here.
- **User — Households, merchants:** hold, spend. Savings accrue when fees stay low.
- **Data revenue:** PSPs see every transaction; analytics, targeting, fraud-score APIs are the invisible fourth stream.
- **Vacuum vs displacement:** whether the middle tier fills a gap or displaces an incumbent decides if it can earn at all.

The critical insight: wider user base = larger *economic rent* (excess profit above cost) the PSPs can earn per user. Fixed onboarding cost ÷ large base = low cost per user; fees × users = revenue that compounds with reach. Drex — the Banco Central do Brasil's CBDC — was designed exactly around how much middle-tier rent should go to incumbent banks versus new entrants.



Pyramid widens downward; rent compounds upward.

Evans and Wurster's deconstruction logic on a vertical stack: only the middle tier carries a BM, and only it captures the rent that the wider base generates.

Is the Sand Dollar's Licence Regime a Permanent Moat or a Passing Window?

The Arbitrage Tension

Regulatory arbitrage = a firm earns profit specifically because it faces a lighter rulebook than its competitors, not because it is better at the underlying business. The advantage lasts only as long as the rulebook gap does.


The Sand Dollar — the Central Bank of the Bahamas' retail CBDC — arrived with a specific regulatory *carve-out* (a narrow exception to the normal licensing rule): PSPs authorised under a new tier could distribute the digital claim without holding a full commercial-banking licence. The intent was inclusion — but the side effect was a real arbitrage window.

- **Lighter prudential burden:** the new PSP-tier requires less capital, fewer filings, less supervisory scrutiny than a full bank — lower fixed cost to operate.
- **Faster onboarding rules** for end-users than the conventional bank-account standard.
- **First-mover positioning** for entrants who can scale before the perimeter widens to include them.

Two strategies emerge. The *compounder* uses the light-rule window to build heavy compliance infrastructure early — when rules tighten, it is already built and the entrant now has a moat. The *rider* simply enjoys lower costs for as long as the window lasts — when rules tighten, the rider loses its edge overnight. Same window, opposite outcomes.

The pattern is identical to e-money licences in other markets (for example: EU e-money institution licences under PSD2, Singapore MAS payment-services licence) — same structural carve-out, same compounder-vs-rider outcome fork.

Sand Dollar's PSP-tier carve-out is a real arbitrage window. Entrants who invest in compliance early convert it into a moat; those who ride it lose it when it closes.



Why Does Drex Find Different Adoption Conditions Than Helvetia in Its Domestic Market?

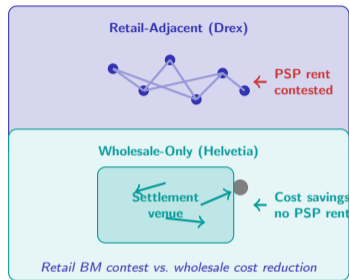
The Context Comparison

Retail CBDC = citizens and merchants hold digital central bank money directly.

Wholesale CBDC = only banks hold it, for interbank settlement. The question is not “CBDC yes/no?” — it is “retail or wholesale?”; the BM is fundamentally different in each case.

Drex — Brazil’s retail CBDC pilot — was designed for a large market with deep fintech penetration and an already-dominant instant-payments rail (Pix, the BCB’s domestic rail used by most Brazilian adults). Helvetia — a joint SNB, BIS Innovation Hub, and SIX Digital Exchange experiment — was designed at the wholesale layer for a small open economy (Switzerland) with a deeply banked population and no retail gap to fill.

- Drex inherits a retail context where the pyramid base is already digitally fluent — the BM contest is whether incumbent banks or new entrants capture the middle tier.
- Helvetia operates in a wholesale context where the question is whether interbank settlement can be made cheaper, not whether end-users will adopt anything new at all.
- Same technology family, fundamentally different BMs: Drex has a retail value-capture question; Helvetia has a settlement cost-reduction question.
- In Drex, PSPs compete for rents; in Helvetia, banks save settlement fees but no new revenue stream is created. BM opportunity exists only in Drex’s retail design.
- The lesson: where the CBDC sits in the retail-vs-wholesale spectrum determines whose BM matters and how the rent is split.




Where on the retail-vs-wholesale spectrum the CBDC sits determines whose BM matters: PSP-tier rent contest in the retail case, settlement-cost contest in the wholesale case.

What Five Tests Predict Whether a CBDC Programme Creates Lasting Public-Plus-Private Value?

The Five-Test Synthesis (CBDC Lens)

- 1 **Friction test:** does the CBDC actually remove a real retail or wholesale friction — reach, cost, settlement delay — or is it a public-relations exercise dressed as infrastructure?
- 2 **Platform test:** are network effects present at the PSP layer or at the inter-jurisdictional layer (mBridge-style), and does each new participant make the next one more valuable?
- 3 **Rebundling test:** can the wallet add adjacent programmable services — subsidy, payroll, conditional disbursement — once trust is established?
- 4 **Infrastructure test:** is the CBDC filling a vacuum (no existing instant-payments rail) or displacing an incumbent rail that already works?
- 5 **Arbitrage test:** is the licence carve-out for new PSPs widening or narrowing? Will entrants convert it into a compliance moat or lose it when the perimeter widens?



A CBDC programme that passes at least three of these tests has a defensible split between public-good issuance and private-sector distribution. Programmes that pass only the friction test become politically popular but commercially hollow.

A CBDC that passes three or more tests has a defensible split between public-good issuance and private-sector distribution. Programmes that pass only one become politically popular but commercially hollow.

The Pitch

PUBLIC MONEY
FOR EVERYONE



"The mint is public; the moat is private. The BM lives in the gap."

vs.

The Future



Wallet by
Subscription

*The cash is free.
The wallet is not.*