

ResearchTech

Snappi in Context

Piraeus, Revolut, and the Contest for Greece's Trusted Credit Relationship

A comparative analysis of Snappi Bank S.A. against its majority shareholder, Piraeus, and the dominant digital challenger in the Greek market, Revolut — June 2026.

Table of Contents

Nº	Section
1.	The thesis: Snappi is a credit business wearing an app, and the relationship that decides it is still in play.
2.	Snappi is built as a bank, not an app — and that distinction is the strategy.
3.	The intake valve is leaking exactly where trust is formed — and the most important proof is not yet public.
4.	Piraeus is a strong digital incumbent and a capital-rationed parent — and both facts constrain Snappi.
5.	The parent's limited integration is managed separation by design, not neglect.
6.	Revolut has won everyday spending in Greece without becoming a bank Greeks trust with their lives.
7.	Revolut wins the normal day and loses the bad one — and in Greece it is still not a local bank.
8.	Side by side: each is strong on a different axis, and none has closed the primary relationship.
9.	What the comparison reveals.
10.	The regulatory clock is the one constraint Snappi cannot manage by working harder.
11.	Recommendations.
12.	What matters less — and what only Snappi's own numbers can settle.

This report compares Snappi with the two institutions it was set against: **Piraeus**, the systemic Greek bank that owns the majority of it, and **Revolut**, the digital challenger that has already won a large share of everyday spending in Greece. It is a comparative analysis; it reaches conclusions and offers recommendations, but it is not a sales proposal. Where the evidence is firm, it is stated plainly. Where it is soft, contested, or reported by management rather than independently verified, it is flagged in the line. Figures internal to Snappi's own operations are identified as such and are not asserted here as fact — the institution holds the real numbers, and several of the conclusions below can only be settled by them.

1. The thesis: Snappi is a credit business wearing an app, and the relationship that decides it is still in play.

Beneath the interface, Snappi is a credit business. Under the 0.2% European cap on debit interchange, daily banking does not pay for a bank: covering a launch-year cost base on interchange alone would require on the order of **900,000 active spenders at roughly €500 a month — close to 9% of the country**, and several times the current base. The economics only close if Snappi lends, and lends well, to a Greek population the bureau cannot see. That single fact reorders everything. The app is not a product to be polished; it is the **intake valve for the salary and cash-flow data on which the lending engine depends**.

The contest that matters, then, is not user count, app rating, or "AI." It is the **trusted, local, salary-anchored, credit-bearing primary relationship** — and the comparison that follows shows that one version of it is still contestable. Piraeus already holds deep incumbent primacy, and Revolut has already won everyday spending; what neither has claimed is the **digital-native, locally licensed, credit-led primary account, above all among younger and mobile-first Greeks**. Snappi is the only one of the three attempting to combine a Greek licence, local rails, human resolution, and local credit in a single proposition — and the window to do so is measured in quarters, not years.

Three issues decide whether it can — and all three protect the same asset:

1. **The front door.** App and onboarding reliability is upstream of the primary relationship, and therefore upstream of the underwriting data that is the only durable moat.
2. **Credit governance before the clock.** The flagship automated loan collides with a dated, specific change in European consumer-credit law on **20 November 2026**.
3. **A narrow local moat, and a disciplined reading of the parent.** The defensible edge is local trust, rails, and credit — not generic digitality — and the parent is a bounded asset, not a distribution shortcut.

The chain underneath all three: **app reliability → the primary relationship → salary and cash-flow data → thin-file underwriting → the only durable moat.**

2. Snappi is built as a bank, not an app — and that distinction is the strategy.

What it is. Snappi is **Greece's first ECB-licensed neobank**, holding a full banking licence rather than the lighter electronic-money permission most challengers run on. The licence is part of the product: it carries deposit protection, the right to take deposits and to lend, and the institutional trust that lighter models cannot replicate. It is a joint venture, **Piraeus 55% / Natech 45%**, with Natech also serving as the core-banking vendor — a related-party arrangement worth noting once, plainly. Public launch was September 2025. Reported traction, drawn from parent-group disclosures and **stated by management rather than independently audited**, is on the order of **100,000 onboarded customers, 170,000 app users, and 70-plus Pay Later merchants**, a material share of them acquired through state-benefit programmes (Fuel-Pass, Youth-Pass). The stack is a zero-fee current account with a Greek IBAN, IRIS instant payments, a debit card, savings, two live lending products, and 24/7 human support in Greek and English.



Full Snappi dossier
online

<https://app.research.tech/share/d/lu0a8Fhg5KJeRpBs9lDWvIFRsNzjp5T>

The economic spine, stated as arithmetic. The interchange cap makes the model's logic unambiguous: payments cannot fund the bank, so credit is not an upside feature to be added later — it is the business. Every other fact in this profile follows from that.

The genuine differentiator is human support. Across review surfaces, 24/7 live human support — named agents, no bot wall — is the single most-praised element, and it is uncommon in a field that has competed on automation. It is a real asset. Its strategic weight, however, depends entirely on the front door working; the evidence on the next page shows it is currently doing a different job.

The lending engine is already live, and its terms matter. Because the arithmetic makes lending existential, the products are worth stating precisely:

- **Cash Now**, marketed as "**2-1-0**" — a two-minute application and a sub-ten-second automated decision, presented publicly as a no-human-touch instant-loan journey, live since 23 December 2025. €150–€1,000; 0% nominal interest; a one-time operational fee of €15 (€150–€500) or €30 (€501–€1,000), 70% refunded if repaid within 30 days; a late fee of €3 per month capped at €9.
- **Pay Later** (live since November 2025) — four interest-free instalments across the merchant base, the first payable immediately. One observable detail, developed on page 10: the published terms leave the cumulative late-fee cap **blank**, where Cash Now's is set.

The thin-file gap is real. Teiresias, the Greek bureau, records formal credit and defaults — not rent, utilities, or telecom. The population Snappi is built to serve — younger, gig, seasonal, first-job earners — is therefore genuinely invisible to the incumbent model: creditworthy but unscored. Underwriting them on captured cash-flow data is sound in principle. Whether it works in practice is a question only Snappi's own data can answer.

3. The intake valve is leaking exactly where trust is formed — and the most important proof is not yet public.

The experience behind the headline. The store headlines — **iOS 4.2 (≈3,200 ratings), Google Play 4.0 (≈5,750 reviews)** — sit above a materially harsher written-review pattern on the channels where customers write rather than merely rate; one third-party aggregation places the written average near **2.4 stars** (*directionally well corroborated, the exact figure less so*). The complaints do not scatter. They cluster at the precise moments that decide whether a customer will trust the institution with a salary: **eGov/KYC identity loops** (worn Greek IDs, OTP timing, verification restarts), **IRIS/QR and online-payment execution**, and **login and stability**. A late-May 2026 redesign polished the surface; the clusters persist.

The support paradox. Snappi's most-praised feature co-occurs with its highest complaint volumes. Human resolution at the moment of fear — a frozen payment, a failed transfer — is genuinely constitutive of a trusted bank, and it is the durable advantage Revolut's automation cannot easily match (page 7). But today it is also a **rescue layer**: it scales with defects rather than with customers, repairing by hand the trust the product breaks at the front door. Both are true at once — the seam is real, and it is currently over-deployed on avoidable failures it should not have to absorb. (*No internal cost figure is offered for this: at Snappi's scale it would be a modelled guess about numbers the institution holds and the analyst does not.*)

Registered is not monetizable. A structural detail most outside observers miss: customers who do not complete full (LIVEID) verification face a first-credit-from-an-EU-account requirement and a **€15,000 annual credit cap**. A restricted account is a different financial animal from a verified one, so headline customer counts overstate the economically active — and creditable — base.

A moat that survives a Greek IBAN. Distribution of state programmes is a narrow but durable advantage: a passported competitor cannot route Greek state funds, so the subsidy-acquisition channel persists even after Revolut localises (see page 6). Its value, though, turns on a question the headcount cannot answer: whether Youth-Pass and Fuel-Pass cohorts convert into salary-deposit primacy at rates comparable to organically acquired users, or simply add non-monetizable accounts of the kind the previous point describes.

The economics of being wrong are thin and refund-sensitive. The public terms show how little room the instant-credit model leaves for mispricing. A €15 fee on a €150–€500 loan is roughly 10% of principal gross — but the 70% early-repayment refund cuts the realised fee on a good borrower to about €4.50, near 3% of principal, with the full fee retained mainly from slower or worse payers. The margin then left to absorb defaults, funding, and operating cost is narrow and refund-sensitive — and instant, frictionless access to €500 is exactly the product that draws adverse selection, where risk underpriced at origination tends to surface two cohorts later rather than at the point of sale. The same thinness runs through the liability side: a book funded by a rate-sensitive, zero-fee deposit base paying 0.853% AER is a funding-quality question as much as an origination one. Whether the thin-file book actually sits inside that margin is not something public terms can settle; the cohort loss curves, reserved to the final page, decide it.

The one set of numbers that converts the thesis into a result. The first Cash Now cohorts have now matured, which means approval, take-up, delinquency, vintage and loss data already **exist internally** — alongside salary-deposit share, active-usage rates, and the unpublished FY2025 financials. Until seasoned, through-cycle loss curves are visible, the lending moat is a **credible thesis, not a proven result**. A small contrast makes the point: the savings rate is published to four decimals, while no active-usage figure appears anywhere. *The most consequential questions about this bank sit on the far side of its own firewall — they return as the close of this report.*

4. Piraeus is a strong digital incumbent and a capital-rationed parent — and both facts constrain Snappi.

Every claim in this profile is drawn from Piraeus's own disclosures and is presented as intelligence, not as an audit of the parent.

Not a legacy laggard. The Piraeus consumer app rates **4.7 on iOS (≈188,000 ratings) and 4.6 on Google Play (≈171,000 reviews)** on the Greek storefronts, against **3.0–3.2 million digital clients**, with the overwhelming majority of transactions digital. *(For the record: the "4.8 from ≈3,200 ratings" figure that has circulated is the U.S. storefront — a small regional slice. The Greek-market figures above are the relevant ones.)* This forecloses a comfortable story: Snappi cannot win simply by being digital-first, because the parent — and the wider field, with Eurobank and NBG also rating in the high 4s — is already digitally credible. Piraeus's own youth products, the **"1 Fund" account (ages 18–30, 3% on balances to €5,000)** and a mini e-loan, target precisely Snappi's demographic.



Full Piraeus dossier
online

<https://app.research.tech/share/d/GxODbmtDN-AreohxH8KEInS-QSjMDovnB>

Capital is the master variable, and it is rationed. Piraeus runs an ambitious plan on the **thinnest core-capital buffer of the four Greek systemic banks — CET1 ≈12.6%**, the least headroom above its own capital targets (against ≈18.8% at NBG, ≈15.6% at Eurobank, ≈15.0–15.4% at Alpha). That turns every initiative into a zero-sum draw on the same scarce capital; management has acknowledged that the first quarter of 2026 consumed close to half the year's loan-growth budget while the dividend stepped up. Documented support to Snappi is real but measured — on the order of **€23 million across 2025** — holding the stake at 55%. The implication for Snappi is concrete and not adversarial: **the group cannot fund everything at once, so Snappi's claim on capital is strongest when it is the most capital-light, fastest-to-traction item on the roadmap — and when its lending economics are proven before anyone must choose between Snappi's growth and the parent's distribution.**

The single load-bearing assumption. The capital trajectory leans on roughly **+75 basis points of relief from a "Danish Compromise" insurance-consolidation treatment that has no confirmation from the supervisor, no public timeline, and no disclosed contingency**, and must resolve by end-2027. It is the kind of unverified, load-bearing input that decides whether an otherwise-sound plan holds — and, by extension, how much patience the parent can extend to a subsidiary.

5. The parent's limited integration is managed separation by design, not neglect.

The strategy, in the parent's own words. Piraeus's 2026–2030 Capital Markets Day deck (March 2026) sets out a "two-pillar strategy: attack neobanks via Snappi, defend the young and tech-savvy portion of our own clientele via the Piraeus app." (Source: *Piraeus Financial Holdings, 2026–2030 Capital Markets Day presentation, March 2026.*) Read plainly, the limited integration after eighteen months is **managed separation by design**. The corollary follows from the capital page: migrating three-million-plus digital clients into a zero-fee subsidiary is value-neutral-to-dilutive at the group level. The realistic posture toward the parent is therefore to seek what it can rationally and lawfully give — milestone-based capital, reusable risk and AI tooling, and acquisition funnels that do not cannibalise the core — rather than a customer hand-over that will not come.

A structural wound, not a one-off charge. Law 5264/2025 forces the conversion of Swiss-franc loans to euros at preferential rates, with banks bearing the cost. The shape matters more than the headline **€57 million provision**: the converted book is locked near 2.30–2.90% against a mortgage yield around 4.81% — a permanent reduction in net interest income, not a single charge. (*Larger required-provision estimates circulate; they rest on applying a broad loan-market share to the sector's exposure, so they are a warning light, not an underwriting conclusion — the borrower-level detail is not public.*)

The shape of the risk is timing. The costs — net-interest-income compression and the franc conversion — are front-loaded; the benefits — insurance synergies, Snappi, the AI programme, efficiency — are back-loaded. Under a thin buffer, that sequence is the ordinary way bank plans miss even when each pillar is individually sound.

Two contrasts worth carrying. NBG acquired comparable insurance economics capital-light, through a minority stake plus a ten-year distribution agreement; Piraeus took the 100%, capital-absorbing route into Ethniki — a genuine question for the peer with the least capital to spare. And headline fee growth of roughly 32% is largely inorganic; stripping the first-time insurance consolidation leaves **organic growth closer to 19–20%**.

The hard limit to design around. Piraeus holds a decade of "Yellow" loyalty data and a **€200 million, three-year AI programme** (with Accenture and Anthropic, **no quantified return yet disclosed**). But Greek bank-secrecy law and data-protection rules make raw transfer of that customer data to Snappi legally fraught — and, in any cleanly reusable form, likely unlawful, since consent does not reliably cure the constraint. The safe planning assumption is to structure for reusable tooling, feature architectures, and governed model outputs — not raw data.

6. Revolut has won everyday spending in Greece without becoming a bank Greeks trust with their lives.

The scale is real. Two million-plus Greek users (≈18% of the population), 244 million transactions in 2025 (up roughly 40%), and — the figure that matters more than the headcount — **68% of their Greek card transactions are already domestic.** That last number is the tell: this is everyday spending leaving Greek-bank cards, not holiday money. The app is excellent, **4.8 on both stores.** Nothing here argues the product is weak. The argument is that the product is not the same thing as a bank people route their lives through.

Broad, but shallow. Despite the reach, the depth is thin. Revenue per customer runs around **\$101 — below the ≈\$155 neobank average, and roughly six times below the ≈£435 (about \$565) that established incumbents earn per customer.** Only about a **quarter of Revolut's under-35 users treat it as their primary account** (growing, but from a low base), and its **loan-to-deposit ratio is ≈6%.** A large, low-yielding deposit base it has not turned into well-underwritten assets is a cost, not a moat. The picture is a high-volume spending-and-FX utility, not a primary bank.

The ceiling appears structural, not executional. This is the finding that travels beyond Revolut. No challenger — Monzo and Starling included, after a decade — has durably pushed much past roughly a quarter-to-a-third of its base into primary use, and at least one peer's primary share has recently edged down, not up. Switching inertia is deep; product polish moves it only at the margin. The salary relationship is won on trust and a genuine reason to route money, not on features — which means it is not closed simply because a competitor has the better app.

The lending engine is not where the headline implies. Revolut's ≈£2.2bn loan book sits in its **Lithuanian** entity; its newer ring-fenced banks launch on regulatory-minimum balance sheets and must build underwriting, collections, and credit risk from zero. Lending is a local institutional discipline, not a product tab that travels with the brand. The consolidated numbers make Revolut look like a running lender across its markets; in most of them, including the one that matters here, it is not. *(A note on the celebrated "33% efficiency ratio": it is a bespoke definition — administrative cost excluding sales and marketing, over gross profit — flattered by peak rates and an electronic-money cost base, not the standard cost-to-income figure peers report.)*



Full Revolut dossier
online

<https://app.research.tech/share/d/Aqn0OqefwbRcEo8B3GV3Tnvi8UZkvBpt>

7. Revolut wins the normal day and loses the bad one — and in Greece it is still not a local bank.

Status in Greece, as of early June 2026. For all the everyday-spending share, Revolut is **not yet a local bank**. It still has **no live Greek IBAN** (customers hold a Lithuanian IBAN); deposits sit under the **Lithuanian** guarantee scheme, **not** the Greek one; and it offers **no Greek consumer credit** — no loans, credit cards, or overdrafts. The local virtual branch and Greek-IBAN migration were reportedly approved in principle by the Bank of Greece and targeted for 31 March 2026, but the launch is not yet live and is now expected later in the year (*press-reported; go-live unconfirmed*). The practical consequence is that Greeks pair Revolut for daily spend with a Greek-licensed bank for salary, tax, and state payments. The window before that changes is measured in quarters — open, but not a margin of safety: an institution with Revolut's engineering and capital can close a local-IBAN-and-rails gap far faster than a domestic incumbent could, which is precisely why the window is narrow.

A regulator has now priced the trust gap. Revolut carries the best store ratings in the market alongside the weakest regulated-trust record of its peer set. In April 2026 the Italian competition authority fined Revolut **€11.5 million in total — of which €5 million was specifically for aggressive and deceptive account-suspension practices** (the remainder for misleading investment promotion and opaque IBAN-migration information); Revolut is appealing. The fine penalised an **architecture**, not a staffing shortfall: accounts frozen "without adequate notice, reasoning, or assistance." In the UK ombudsman data, Revolut shows the **worst non-APP fraud-and-scam uphold rate among its neobank peers — about 37%, against an industry norm near 32%**.

The distinction that matters is prevention versus resolution. Revolut's fraud *prevention* is genuinely strong — large volumes stopped, high intervention success. Its fraud *resolution* is the weakest in its set. Improvement in one does not offset failure in the other, because trust is destroyed not when an algorithm catches a fraud, but when a frightened customer cannot reach a human to unfreeze their money. (*The anchor is the uphold rate, not absolute complaint volume, which is largely a scale artifact.*) Revolut's new in-app assistant extends the same automation-first posture the regulator faulted in its dispute handling — a data point on the direction of travel.

Why this is the open ground. The seam is not Revolut's app; it is the moment something goes wrong. A locally licensed, locally insured, human-answered bank is the exact alternative to "my money is frozen and a machine is handling it." The contest for user count is settled. The contest for the **primary, salary-anchored, locally-insured, human-resolved relationship — and the lending that rides on it — is not.**

8. Side by side: each is strong on a different axis, and none has closed the primary relationship.

App ratings are platform-separated; the written/resolution column is kept distinct from the headline, because the gap between them is the point.

Dimension	Snappi	Piraeus	Revolut
Licence / type	ECB-licensed Greek bank	Greek systemic bank	Lithuanian bank (EU passport)
Deposit guarantee	Greek (TEKE)	Greek (TEKE)	Lithuanian scheme , not Greek
Greek IBAN / IRIS	Yes / Yes	Yes / Yes	No GR IBAN (LT only) / No IRIS
Consumer credit in Greece	Yes — Cash Now, Pay Later	Yes — full suite	None
App rating — iOS	4.2	4.7	4.8
App rating — Google Play	4.0	4.6	4.8
Written / resolution reality	≈2.4 written; support is a rescue layer	Established incumbent at scale	High stars; worst peer FOS uphold (≈37%); €11.5M AGCM fine
Primary-relationship depth	Unproven (no disclosed salary-deposit share)	Deep (incumbent)	Shallow (≈25% of under-35s primary; ≈6% loan-to-deposit)
Support / resolution model	24/7 human	Branch + digital	Automation-first
Lending regulatory exposure	Live automated credit → CCD2 Nov 2026	Established	Minimal (no Greek credit)
Posture toward Snappi	—	55% parent and "two-pillar" competitor	Scale competitor; no Greek credit yet

The read beneath the table. Each institution is strong on a different axis: Piraeus on local trust and scale, Revolut on polish and reach, and Snappi — alone among the three — on holding a Greek licence, deposit protection, local rails, integrated local credit, **and** human resolution in one place. Piraeus proves that incumbents are already digitally credible, so "digital-first" is not by itself a differentiator. Revolut proves that scale and a superb app have not produced a trusted *local primary* relationship. The space none of the three has closed is the intersection: the **trusted, credit-anchored, primary relationship**. Snappi's proposition is the only one aimed squarely at it — which is promising, and unproven.

9. What the comparison reveals.

The profiles converge on three conclusions. They are stated as interpretation, and they set up what should and should not command management attention.

1. The front door is part of the credit engine, not a layer of polish on top of it. For a branchless lender, app reliability is upstream of everything that matters financially. A failed onboarding, a late one-time password, an IRIS payment that times out — read backwards, each is not a one-star review but a customer who never becomes primary, and therefore never produces the salary and cash-flow data the underwriting edge is built on. The comparison sharpens this: Piraeus's digital credibility means there is no UX grace period, and the data-starvation cost of first-mile friction is paid directly into the lending thesis. First-mile reliability is credit infrastructure.

2. Trust under stress is the open competitive seam — and it is the one Snappi is built to win. Revolut has won the normal day and, on the regulator's own evidence, is weakest on the bad one; the automation that built its efficiency is what fails a frightened customer and what an authority has now fined. Snappi's 24/7 human model is the precise answer to that failure — **but only once the product stops manufacturing the emergencies.** The seam is real and a giant cannot quickly copy it; it converts to advantage only after page 3's front door is fixed.

3. Piraeus is a strategic asset, not a shortcut. The parent supplies credibility, capital, and potentially reusable tooling. It does not imply free customer migration, unlimited support, or movable data. The two-pillar strategy and the capital floor say the same thing from two directions: broad distribution is not coming, and patience is rationed. The productive relationship is one of specific, milestone-tied compacts — not a wait for hand-over.

A sober comparable underlines the stakes. Starling is the closest precedent for a licensed local challenger thriving beside a global giant — proof the model can work — but its path to profit rode a large share of government-guaranteed, pandemic-era lending income that is not available here. The structural lesson holds; the timeline does not. Expect a longer, more capital-disciplined road to the same place.

10. The regulatory clock is the one constraint Snappi cannot manage by working harder.

This is the part of the picture hardest to generate from inside the company, because it connects the live product to a dated external rule — and it lands first on the flagship.

CCD2 is the near clock — 20 November 2026. The revised Consumer Credit Directive applies in Greece on that date and reaches **micro-loans and BNPL explicitly**, so both Cash Now and Pay Later are in scope. It grants the consumer a **right to human intervention** and a clear explanation of an automated assessment; it **forbids a decision based solely on automated processing**; and it **bars special-category and social-media data** from creditworthiness, requiring the assessment to rest on relevant financial information rather than credit history alone. A two-minute application resolved by a sub-ten-second decision with **zero human interaction** — the "2-1-0" flow as it operates today — is the precise shape the directive is written to constrain. The work this implies is concrete, not a slogan about AI: a human-in-the-loop override, a clear explanation of the decision, and an auditable reason-for-decline, operational before the date.

The AI Act is the far clock — 2 December 2027. Credit-scoring systems are classified high-risk, with documentation, oversight, and explainability obligations applying on that date (moved later by the May 2026 Digital Omnibus). The trap a risk team will recognise: an engine kept under continuous development can trip the "substantial modification" threshold and forfeit the transitional grace. CCD2 arrives first and bites the flagship first; the AI Act is the second wave.

The data line, drawn precisely. The behavioural signals associated with instant, app-native underwriting — how a handset is held, the age of an email account, device and session patterns — are genuinely useful for **fraud and identity**, and remain available there. They are **legally fragile the moment they inform the credit decision itself**. The lawful-for-credit set is narrower and, fortunately, stronger: consented open-banking cash-flow, verified income and employment, and bureau data — precisely the data the primary relationship produces. The durable edge is not "AI"; it is lawful, explainable cash-flow underwriting that can defend itself to a supervisor.

Two observable signals, and one draft to watch. Snappi's open engineering roles are model **builders** (predictive analytics, quantitative risk, collections); there are no visible model **governors** — validation, model-risk, or AI-compliance roles. Separately, Pay Later's published terms leave the cumulative late-fee cap **blank** where Cash Now's is set — a product-control and governance signal of exactly the kind the builders-but-no-governors gap predicts: easy to fix, but the sort of detail a supervisor reads as a tell, and too important to leave open in a live lending product. A Greek draft law capping total consumer-loan repayment at 30–50% over principal is, as of now, **draft, not enacted**; on a flat €15 fee per €150 it would not bite Cash Now, but it is worth tracking for interest-bearing or longer-dated products.

11. Recommendations.

Three priorities follow from the comparison, mapping to the three issues named at the outset. They are the strongest defensible reading of where management attention, capital, and roadmap should go.

1. Treat front-door reliability as credit infrastructure, not user experience. The highest-return engineering work is not new features; it is closing the eGov/KYC, OTP, IRIS/QR, login, and stability failures that prevent customers from becoming primary. The test of success is not a higher app rating but a higher salary-deposit conversion. A useful discipline is to separate *good attrition* — customers a bank should not want — from *trust-destructive friction*, the creditworthy customers lost at the door, and to treat every support contact caused by a broken flow as a defect that should have been designed out.

2. Build credit governance before automated lending scales, not after the law forces it. The clearest near-term imperative is an explainable, auditable, human-reviewable credit-decision layer over the existing models — a human-intervention override, a regulator-grade reason-for-decline, and a clean separation between fraud-useful and credit-lawful data — operational ahead of November 2026. Ownership of model governance should be named, and the matured loan cohorts treated as the central proof pack.

3. Play the moat narrowly: compete on trust, and resource through the parent deliberately. The right contest with Revolut is not feature parity but the one it structurally cannot win quickly — local-bank trust, human resolution at the moment of fear, Greek rails, and local credit; feature-matching should be limited to genuine table stakes. The parent should be used deliberately rather than waited on: the realistic asks are milestone-tied capital, reusable risk and AI tooling and governed model outputs, and non-cannibalising acquisition funnels — not broad customer migration or raw-data access. Together these decide whether the narrow local moat is held before the window closes.

The internal metrics that would show management whether each priority is working are the ones the public record cannot supply — set out on the final page.

12. What matters less — and what only Snappi's own numbers can settle.

The weakest uses of management attention. Selectivity is the point, and several issues that attract attention are not the ones that decide the outcome:

- Chasing broad feature parity with Revolut — genericity is the threat to a local moat, not the cure for it.
- Treating app ratings as a vanity metric — when the app is the bank, they are the leading indicator of lending-data acquisition cost.
- Optimising user count without primary behaviour — a registered user who deposits no salary is a number, not data.
- Waiting for Piraeus to migrate its customers — its own strategy and capital position say it will not.
- Calling AI the moat without governance — the supervisor will not weigh how good the model is, only whether it can be explained.
- Assuming the parent's data can move — raw transfer is legally fraught and not a safe planning assumption.
- Reading support praise as proof the journey is healthy — today it is the opposite signal.

What only Snappi's numbers can answer. Public evidence can frame the open ground and the constraints. It cannot settle the thesis. These figures can, and they sit inside the institution:

- **Salary-deposit / primary-account share** — the master metric; the line between a payments app and the credit business the economics require.
- **Fully verified, monetizable user base** — how much of the headcount can actually be lent to.
- **Subsidy-acquired versus organic cohort quality** — whether Youth-Pass and Fuel-Pass users reach salary-deposit primacy at organic rates, or remain non-monetizable headcount.
- **Onboarding completion by failure reason** — identity, OTP, payment, stability, decomposed, because each points to a different fix.
- **Cash Now and Pay Later cohort performance** — approval, take-up, repeat use, delinquency, collections, cure and roll-rate behaviour, vintage and loss curves, and risk-adjusted contribution *after* the full cost of risk, funding, and operating cost, by product.
- **Deposit stickiness and salary-inflow as funding** — whether the zero-fee, rate-sensitive deposit base is stable enough to fund a growing lending book.
- **Avoidable support contacts per activated customer, by root cause** — whether support is becoming the differentiator it can be or remaining the rescue layer it is.
- **Where model governance sits** — Snappi, Natech, or an independent line — against the November 2026 and December 2027 clocks.
- **FY2025 post-launch economics** — the audited base the rest of the picture rests on.

Public evidence can identify the open ground, the competitive seam, and the dated constraints. Snappi's own numbers determine whether the thesis is becoming a bankable result — and whether the bank claims the trusted, local, credit-bearing primary relationship before the window closes.