AML AT SCALE

The Structural Limits of Financial Crime Compliance in a Global Fintech

A Case Study Analysis of Revolut

An Evidence-Based Assessment of the AML Capacity Gap

Andréas Hobbelin
Independent AML Practitioner

Working paper — June 2026
Based exclusively on publicly available information

Author Note, Scope and Disclosures

Andréas Hobbelin is an independent AML practitioner with over fifteen years of experience in anti-money laundering, financial crime investigation and compliance. He served as a Senior AML Officer at Swedbank Group during the bank's internal investigation into its Baltic operations between 2016 and 2020. Aspects of his work in that period are referenced in Honungsfällan (2020). He has subsequently held senior AML roles at European fintech firms, has acted as an AML expert witness in financial-crime proceedings, and works on AML technology and advisory.

Scope. This paper makes no allegation of money laundering, terrorist financing, sanctions violation or criminal wrongdoing by Revolut, by any Revolut entity, or by any individual associated with Revolut. It is a structural and analytical assessment, based exclusively on publicly available information, of whether Revolut illustrates structural limits in the ability of an AML/CFT control framework to scale with the size, complexity and risk profile of a global fintech. Every substantive claim carries an evidence label, explained in the section that follows. Analytical opinions are the author's own and are identified as such.

A note on what this paper is. It is not, at its core, a critique of Revolut. It is an inquiry into whether the AML supervisory architecture — designed in an era of nationally-bounded institutions — is yet equipped for the scale, velocity and cross-border asymmetry that the largest fintechs now present. Revolut is the case study not because it is the worst actor, but because it is the clearest available public example: the largest, fastest-growing, most product-diverse and most geographically extended European fintech, with the most substantial public supervisory record. The concept it illustrates applies across the sector.

Disclosures. The author holds no financial interest in Revolut, holds no position (long or short) in any Revolut security or instrument, and has received no funding, commission or consideration from any party in connection with this paper. The author holds a personal retail Revolut account as an ordinary customer. The author worked in forensic services at BDO Norway between 2013 and 2016; this preceded any period in which BDO LLP (United Kingdom) acted as Revolut's auditor, and the author never worked on any Revolut audit or engagement. In January 2024, the author was approached by Revolut's recruitment team regarding a Senior Compliance Manager (MLRO) Norway position; the author attended interviews in February 2024 and did not receive an offer. None of the foregoing has had any bearing on the analysis or conclusions, which rest entirely on the cited public evidence.

How to Read This Paper

Every substantive claim in this paper carries one of four labels. The labels separate what is established fact from what is the author's interpretation — the discipline that distinguishes analysis from allegation. A reader who keeps the labels in mind can see, at every point, exactly how much weight the underlying evidence bears.

[F] — Fact. A claim directly evidenced by a cited public source.

[EBI] — Evidence-Based Inference. A claim that follows logically from cited facts. The inference is the paper's; every step is traceable to evidence.

[AO] — Analytical Opinion. A claim drawing on professional judgment to characterise the significance of evidence. A reader of equivalent experience could reasonably reach a different view.

[OQ] — Open Question. A question the public evidence raises but cannot answer. It is identified for the party who can answer it — and is not an assertion that the answer is adverse.

Specialist terms are briefly explained on first use, so that the analysis is legible to a general professional reader without prior AML expertise. Sources are listed in full at the end of the paper.

Contents

Executive Summary

Revolut is the most consequential European challenger bank of the last decade. Between 2016 and early 2026 it grew from roughly 300,000 customers to more than 70 million, from £2.4 million in revenue to £4.5 billion, and from a single UK electronic-money licence to a portfolio spanning a Lithuanian banking licence supervised by the European Central Bank, a UK banking licence in full operation since March 2026, a Mexican banking licence, and authorisations or applications across India, Colombia, the United Arab Emirates and the United States. The Group operates in 40 markets, is licensed as a bank in over 30 of them, was valued at roughly $75 billion in a late-2025 secondary share sale, and has stated an ambition of 100 million customers by mid-2027. By any measure of scale, product range, geographic reach or speed of expansion, it occupies a position in the European financial system without close precedent. **[F]**

This paper asks one structural question: whether the anti-money-laundering ("AML") controls, the supervisory architecture and the governance machinery available to identify, assess, manage, mitigate and report money-laundering and terrorist-financing ("ML/TF") risk can scale at the rate that the customer base, transaction volume, product complexity and geographic reach are scaling. **[AO]**

To frame the question, the paper introduces the **AML Capacity Gap**: the structural divergence between the size, complexity and risk of an institution's customer and transaction population, and its investigative, supervisory and governance capacity to discharge its obligations in respect of that population. The central idea is that a firm can satisfy every formal requirement of modern AML regulation — documented risk assessments, calibrated monitoring systems, model-governance frameworks, a reporting officer, board oversight — while in operational reality being unable to investigate, conclude on and report a meaningful proportion of the suspicion that its own systems generate or should generate. Where that divergence becomes wide enough, AML ceases to be a comprehensive control system and becomes, in substance, a filtering and prioritisation exercise. The question this paper raises is whether Revolut illustrates the conditions under which that happens. **[AO]**

The paper is explicit about what it is not. It does not conclude that a material AML Capacity Gap exists at Revolut. The public record cannot establish that it does — and, equally, cannot establish that it does not, because the operational metrics that would settle the question are not disclosed, and because the failures that would evidence a Gap (suspicion that is never detected, or never reported) are by their nature invisible. The paper's conclusion is calibrated precisely to that position: given the convergence of an exceptional requirement and consistent signals, proportionate capacity should be **demonstrated rather than assumed**; and, given the asymmetry of access to operational data, the practical ability to resolve the question rests primarily with the

firm, which holds the data, and its supervisors, who hold the powers to inspect it. The case study establishes the concept; it does not convict the subject. **[AO]**

The unicorn and the risk surface. One observation frames the whole paper. The same scale, growth velocity and product breadth that produced an exceptional valuation also produced an exceptional ML/TF risk surface. The public record contains detailed, audited, quantified proof of the former — customers, revenue, balances, transaction volumes — and conspicuous silence on the operational proof of the latter. A reader is invited to hold those two facts together: a firm celebrated for the scale of what it has built has disclosed almost nothing about whether the machinery for controlling the financial-crime risk that scale generates has kept pace. This paper does not claim that machinery has failed. It observes that the public record does not show it has succeeded, and argues that, at this scale, that is a question which should be answered rather than presumed. **[AO]**

A warning, stated as risk and not as prediction. The conditions this paper documents — exceptional scale, rapid growth, product and geographic complexity, an undisclosed operational capacity, a supervisor's finding that suspicious transactions were "not always properly identified," and a control model increasingly dependent on automation at the suspicion-detection stage — are, as Part III sets out, precisely the conditions under which large-scale ML/TF failures have historically occurred at other institutions and gone undetected until external exposure. This paper does not predict such an outcome at Revolut, and has no evidence that one is occurring. It identifies why, on the public record, no advance assurance exists that one would be caught from outside the firm. A reader who absorbs this paper before any future event should come away knowing where the structural vulnerability would lie, what would manifest it, and why it would be invisible until late. That is the paper's purpose: not to allege, but to ensure that the warning, if it is ever needed, was on the record. **[AO]**

Six findings support this position, each developed in the body. **First**, Revolut's disclosure is asymmetric: it quantifies fraud prevention in detail (£600 million prevented in 2024, specialist headcount, AI productivity claims) but discloses no AML operational metric — no suspicious-activity-report volumes, no alert volumes, no measure of what its systems miss — even though both fall within its own definition of financial crime, and even though the proceeds of fraud are, as a matter of law, criminal property whose handling can trigger the money-laundering reporting obligation. **[F] Second**, three independent supervisory and audit signals converged in 2025: the Bank of Lithuania's €3.5 million fine — which the authority stated was the largest it has ever issued — finding that the firm was "not always properly identifying suspicious monetary operations or transactions"; a French supervisory inspection that drove customer-due-diligence remediation across the European Economic Area and globally; and internal-audit findings on financial crime confirmed

by the Board under a remediation plan. **[F] Third**, the only external data on the firm's financial-crime exposure — UK fraud and scam data — places Revolut consistently at the elevated end of its peer group, including as the UK firm with the highest volume of fraud complaints to the Financial Ombudsman Service in two consecutive years. **[F] Fourth**, no single supervisor sees Revolut's consolidated, group-wide ML/TF picture; each sees a fragment, and the whole exists only inside the firm. **[AO] Fifth**, Revolut's pursuit of a US national bank charter brings the most far-reaching enforcement architecture in global finance into direct relationship with the Group — while the variable that drives that exposure, its US-dollar transaction flow, is undisclosed. **[F regarding the disclosure gap; AO regarding the implication] Sixth**, the decision that suspicion exists is a regulated act that, under current law, a human must make and be accountable for; AI can prepare, triage and accelerate that work but cannot perform the judgment, placing a hard limit on how much of the AML decision chain automation can absorb. **[F regarding the legal framework; AO regarding the limit]**

The proportionality of AML capacity to fintech scale is, in the author's assessment, the central supervisory question of European fintech in the late 2020s. The answer to it — for Revolut, for its peers, and for the supervisory architecture being built around them — will define what proportionate AML at scale means in practice. **[AO]**

PART I — THE QUESTION

1. The AML Capacity Gap: Concept and Thesis

1.1 The concept

Modern AML regulation is built on the risk-based approach: the globally-mandated method, set out in Financial Action Task Force ("FATF") Recommendation 1 and transposed across the UK and EU frameworks, by which a regulated firm assesses the ML/TF risk it faces and calibrates its controls in proportion to that risk. The risk-based approach is, by design, a system of prioritisation. It does not require a firm to examine every transaction with equal intensity; it requires the firm to focus resources where risk is highest and to apply lighter measures where risk is lower. A firm that monitored everything identically would be misapplying the approach, not complying with it. **[F]**

The AML Capacity Gap is a structural condition that can arise within a correctly-specified risk-based system. It is the divergence between two things: on one side, the size, complexity and risk profile of the customer and transaction population an institution serves; on the other, the institution's actual investigative, supervisory and governance capacity to identify, assess, manage, mitigate and report the ML/TF risk that population generates. Where the first outgrows the second, the institution may continue to satisfy every formal requirement — it has a risk assessment, monitoring systems, a reporting officer, board committees — while becoming operationally unable to investigate, conclude on and report a meaningful proportion of the suspicion its systems generate or should generate. The formal architecture remains intact; the operational reality hollows out beneath it. **[AO]**

The Gap is therefore not a synonym for non-compliance, and still less for wrongdoing. It is a question about the relationship between requirement and capacity — a question that the formal indicators of compliance are not designed to answer, because a firm can possess all of them and still be overwhelmed. **[AO]**

1.2 What this paper claims, and what it does not

This paper claims that Revolut exhibits every structural condition under which the AML Capacity Gap arises, and that the convergence of those conditions makes the question of proportionate capacity pressing rather than theoretical. It claims that, on the public record, that question cannot be resolved in either direction, and that the responsibility for resolving it rests with the parties that hold the operational data. **[AO]**

It does not claim that money laundering has occurred at Revolut. It does not claim that Revolut has breached its AML obligations. It does not claim that Revolut under-reports suspicion. It does not claim that any supervisory finding establishes a systemic failure. It does not rely on any non-public information, and it does not imply the existence of any. Each of these limits is maintained throughout, and the evidence labels make the boundary visible at every step. **[AO]**

1.3 The thesis, stated precisely

The thesis is this: where an institution's requirement profile is as exceptional as Revolut's, and where the publicly-observable signals are as consistent as those this paper documents, proportionate AML capacity should be demonstrated rather than assumed — and the public record does not contain that demonstration. The argument is not that the Gap is proven. It is that the conditions which would make a Gap likely are present and documented; that the evidence which would confirm or dispel it is held only by the firm and its supervisors; and that the appropriate response to that situation is neither to presume failure nor to presume adequacy, but to require demonstration. **[AO]**

This framing is deliberate, and its discipline is the source of its strength. A stronger claim — that comprehensive AML compliance is impossible at Revolut's scale — is unavailable on the public record, because proving it would require the firm's actual operational parameters, which are not disclosed; it is also legally and analytically fragile, because a single counter-disclosure could refute it. The calibrated claim is more durable precisely because it does not depend on data the author does not have. It rests instead on what is visible: the scale, the signals, and the silence. **[AO]**

2. The Legal Architecture

The analysis depends on three features of AML law, stated here for the general reader because the rest of the paper builds on them. **[F]**

2.1 The suspicion threshold is low

Under the UK Proceeds of Crime Act 2002 ("POCA"), a person in the regulated sector commits an offence if they fail to disclose where they know or suspect, or have **reasonable grounds** to know or suspect, that another person is engaged in money laundering (section 330). The proceeds of crime — including the proceeds of fraud — are "criminal property" (section 340), and dealing with criminal property is itself money laundering. The disclosure takes the form of a Suspicious Activity Report ("SAR") submitted to the national Financial Intelligence Unit (in the UK, the National

Crime Agency). The leading authority, R v Da Silva, holds that suspicion requires only "more than a fanciful possibility" — not proof, not even belief. The equivalent EU framework, transposed in Lithuania where Revolut's principal bank is licensed, applies a substantively identical standard, with reports filed to the national Financial Intelligence Unit as Suspicious Transaction Reports ("STRs"). The threshold that triggers the obligation is, by design, low. **[F]**

2.2 The test is objective

Liability does not depend only on whether the firm actually suspected. It arises where a reasonable person in the regulated sector, with the same information, **would have** had grounds to suspect — whether or not anyone in the firm reached that conclusion. The benchmark is the reasonable AML professional, not the firm's software and not its commercial convenience. A monitoring system calibrated so that it does not surface what a reasonable professional would have found does not discharge the obligation; it relocates the failure. **[F]**

2.3 The decision is human and personally accountable

The reporting obligation is owed by the firm and discharged through a named, personally-accountable officer — in the UK, the Money Laundering Reporting Officer, an individual approved by the regulator under the Senior Managers and Certification Regime, with equivalent personal-accountability regimes across the EU. Criminal and regulatory liability attach to persons. A model cannot hold the role, cannot stand behind the judgment, and cannot, under current UK and EU supervisory expectations, make the suspicion decision on the firm's behalf. This is not a transitional technological limitation; it is a structural feature of how AML liability is constructed. Its consequence for the role of artificial intelligence is developed in Section 8. **[F regarding the framework; AO regarding the implication]**

PART II — THE EVIDENCE BASE

3. Growth Trajectory, 2016–2026

This section establishes Revolut's growth as quantitative fact before any analysis is performed, because the Capacity Gap thesis rests on a comparison between the growth of the requirement and the growth of the capacity to meet it. The figures below are drawn from Revolut's audited annual reports and Companies House filings, with the 2025 figures confirmed against the Annual Report 2025 and the firm's March 2026 results announcement. **[F]**

Between 2016 and early 2026, the customer base grew roughly 230-fold; annual revenue grew roughly 1,875-fold, from £2.4 million to £4.5 billion; pre-tax profit moved from sustained losses to £1.7 billion in 2025 (up 57% year on year); total customer balances reached £50.2 billion; and the workforce grew from a few hundred to more than 16,000. The firm reports analysing more than 10 billion transactions in 2025 alone. **[F]**

Year-end	Customers	Revenue	Pre-tax profit	Principal milestone
2016	~0.3m	£2.4m	(loss)	UK e-money authorisation
2018	3.5m	£58.2m	(loss)	Lithuanian banking licence approved; sanctions self-report
2020	14.5m	£222m	(loss)	EU base transitioned to the Lithuanian bank
2021	16m+	£636m	£26m	First full-year profit; UK banking licence applied for
2022	26.2m	£923m	(£25m)	First Lithuanian AML fines; FRC audit-quality finding
2023	38m	£1.8bn	£438m	UK flagged-account review reported; BDO qualified 2021 accounts
2024	52.5m	£3.1bn	£1.1bn	UK banking licence (restricted); French inspection; ECB supervision begins
2025	68.3m	£4.5bn	£1.7bn	€3.5m Lithuanian AML fine; >10bn transactions analysed; >1/3 of staff in financial-crime prevention
Q1 2026	>70m	—	—	UK full authorisation (Mar); US charter applied (Mar); Italian fine (Apr)

Sources: Revolut Annual Reports 2023–2025; Companies House filings; regulator decisions. Transaction-volume figures cited in the text are annual figures as reported by Revolut.

The structural point is not the commercial achievement but its consequence. A control framework built to monitor one billion transactions a year is not the same kind of system as one built to monitor ten billion; at the larger scale, monitoring

necessarily becomes machine-dependent, because the alert volumes cannot be reviewed by people. Revolut crossed this threshold continuously, not once, and at each step the framework had to be redesigned rather than merely enlarged. Whether each redesign occurred in time, with adequate testing and proportionate human capacity behind it, is the question the remainder of the paper examines. The growth is the firmest fact in the paper; it is established almost entirely from the firm's own disclosures, and it is not seriously contestable. **[F/AO]**

4. Entity and Supervisory Architecture

Revolut is not a single company in a single country. The Group consolidates roughly fifty entities; the principal regulated ones include Revolut Ltd (the UK electronic-money institution, supervised by the Financial Conduct Authority); Revolut Bank UK Ltd (the UK bank, supervised prudentially by the Prudential Regulation Authority and for conduct by the FCA, fully authorised in March 2026); Revolut Bank UAB (the Lithuanian bank through which EEA operations are conducted, supervised for AML by the Bank of Lithuania); Revolut Holdings Europe UAB (under direct European Central Bank prudential supervision since 1 January 2024); a Mexican bank; and a growing US presence currently operating through a partner bank. **[F]**

This architecture has a consequence that Section 10 develops: no single supervisor receives, or has the mandate to compile, the consolidated group-wide ML/TF picture. The FCA sees UK conduct; the Bank of Lithuania sees the EEA bank's AML; the ECB sees prudential health; host-state supervisors see local branches; US authorities see US activity. The cross-border cooperation mechanisms that exist — supervisory colleges, FIU-to-FIU channels, memoranda of understanding — are designed for supervisors **sharing** their separate views, not for assembling a single consolidated view. That view exists in exactly one place: inside Revolut. This is not an accusation; it is a structural feature of any multi-entity, cross-border group, whose consequences are simply most pronounced at this scale. **[F regarding the architecture; AO regarding the consequence]**

5. The Disclosed Control Framework — and the Disclosure Asymmetry

5.1 What the firm discloses

Revolut's disclosed control framework has matured substantially. By 2025 it describes real-time transaction monitoring, daily sanctions screening, risk-based customer due diligence, machine-learning and computer-vision tools, a central AI platform with access to external large-language-model providers, an explicit model-risk-management category, and — the headline disclosure — more than one-third of its global workforce, implying over 5,300 people, working in financial-crime prevention. Risk and compliance functions grew more than 42% in 2025. The firm credits AI with an "almost 10x" increase in the number of cases reviewed daily. These are real and substantial investments, and the paper credits them. **[F]**

5.2 The asymmetry at the centre

The most analytically significant feature of Revolut's disclosure requires nothing more than reading its own annual report. The firm quantifies **fraud** with commercial precision: £600 million of fraud prevented in 2024, tens of millions of fraud-warning messages to users, 200-plus fraud specialists, the "almost 10x" case-review figure. Across the same reports, it discloses **no AML operational metric whatsoever**: not the number of SARs or STRs it files; not the volume of alerts its monitoring generates; not the proportion of alerts it closes automatically without human review; not whether or how it tests for the suspicious activity its systems miss; and not how many of its financial-crime staff are AML investigators as opposed to fraud, sanctions or onboarding personnel. **[F]**

Both fraud and money laundering fall within the firm's own definition of financial crime, and both are run by the same organisation; yet one is quantified in detail and the other described only as an architecture. The contrast matters because of the legal link established in Section 2: the proceeds of fraud are criminal property, and handling criminal property is money laundering. Fraud that completes through Revolut accounts — the fraud the £600 million did not prevent — generates, by operation of law, a population of activity in which money-laundering suspicion is highly likely to arise. In many consumer-fintech environments, fraud-related proceeds are likely to form a major component of the AML reporting exposure — and that is precisely the component for which Revolut publishes the upstream prevention numbers and none of the downstream reporting ones. **[F regarding the pattern; AO regarding its significance]**

The paper does not infer under-reporting from this asymmetry, and Section 9 maintains that limit strictly. The observation is narrower and survivable: the disclosure asymmetry removes the firm's most direct means of demonstrating that no Capacity Gap exists. A firm confident that its AML identification and reporting were proportionate to its fraud exposure could disclose summary metrics to show it. The silence is not evidence of a Gap; it is the removal of the evidence that would dispel one. **[AO]**

5.3 The audit dimension

Revolut's audit history is part of the disclosure-quality picture. The UK Financial Reporting Council ("FRC") found that the firm's 2020 audit, performed by BDO, was "inadequate" and that "the risk of an undetected material misstatement was unacceptably high," citing revenue recognition and the testing of payment processes. BDO subsequently issued a qualified opinion on the 2021 accounts, stating it had been unable to obtain sufficient assurance over the completeness and occurrence of £477 million of revenue — approximately 75% of the £636 million reported — because the firm's IT systems were not designed in a way that allowed the relevant controls to be tested; the accounts were filed roughly five months late. BDO did not qualify the firm's going-concern status and independently confirmed customer cash balances. In August 2025, Revolut announced that EY would replace BDO from the 2026 financial year. **[F]**

The relevance to AML is indirect and is stated as such: a firm whose external auditor could not, for a period, fully test the IT systems underlying revenue recognition was, in the same period, relying on IT systems to monitor transactions for financial crime. The paper does not assert that the audit findings demonstrate any AML control deficiency; the two systems are distinct. It observes that the documented IT-control difficulties and the AML monitoring both depend on the same underlying data and systems infrastructure, and that the audit history is therefore part of the context in which the disclosure asymmetry should be read. **[AO]**

6. The Supervisory and Audit Record

6.1 The record

Across the period, Revolut entities have been the subject of continuous and escalating supervisory engagement. Stated as fact, without inference of wrongdoing: a self-reported sanctions-screening failure (2018); the FRC's "inadequate" audit finding (2022) and BDO's qualified opinion on the 2021 accounts (2023); four Bank of Lithuania enforcement actions between 2022 and 2025 totalling roughly €3.82

million; the ECB's assumption of direct prudential supervision of the EU holding company (January 2024); a UK banking-licence "mobilisation" period of roughly twenty months — unusually long — before full authorisation in March 2026; and an Italian consumer-protection fine of €11.5 million (April 2026, under appeal). **[F]**

6.2 The 2025 convergence

Three of these signals concentrated in 2025, and all three bear on financial-crime controls. **[F]**

Signal (2025)	Source	Substance
€3.5m fine	Bank of Lithuania	Stated by the authority to be the largest fine it has ever issued; finding of "violations and shortcomings in the monitoring of business relationships and operations (transactions)" that led to the firm "not always properly identifying suspicious monetary operations or transactions." No confirmed money laundering found; resolved by administrative agreement and a corrective plan.
EEA-wide & global remediation	French supervisor (ACPR)	An inspection of French operations drove customer-due-diligence enhancements applied across the EEA and globally — disclosed in the firm's own annual report.
Internal-audit findings	Board Audit Committee	Internal audit identified financial-crime findings, now under a Board-monitored Control Enhancement Plan.

The April 2025 fine is the single most significant item in the record, and its language is what matters. The Bank of Lithuania did not find a paperwork lapse; following a scheduled inspection, it found that the firm was not always properly **identifying** suspicious transactions — a finding about the operation of the detection systems themselves, at the step on which all subsequent reporting depends. The supervisor also confirmed that no money laundering was established, and the firm has remediated. Whether the remediation has resolved the underlying cause, rather than the specific instances identified, is a question only re-inspection can answer. **[F regarding the finding; OQ regarding remediation]**

6.3 Balance, and the sector-wide pattern

The record also contains clear positives, most notably the PRA's grant of full banking authorisation in March 2026 after a thorough and unusually extended process — a supervisory judgment that the firm had met the conditions for authorisation. And the dynamic the paper describes is sector-wide, not Revolut-specific. The FCA fined Starling Bank £28,959,426 in September 2024 for financial-crime control failings, finding that its measures "did not keep pace" with growth from 43,000 to 3.6 million customers; Starling itself accepted that its controls "failed to keep pace with the growth of the business." Germany's BaFin fined N26 €9.2 million in May 2024 for systematically late suspicious-activity reporting, following an earlier €4.25 million

fine. In each case the regulator's theme was identical: compliance capacity failing to scale with rapid growth. The point is not that Revolut is uniquely deficient; it is that the structural pressure producing this sector-wide pattern is most acute at the largest and fastest-growing institution, which is why Revolut is the clearest case through which to examine it. **[F/AO]**

PART III — THE ANALYSIS

7. The Arithmetic of AML at Scale

7.1 Purpose and the disclaimer that governs the section

Because Revolut discloses no AML operational figures, this section illustrates the **order of magnitude** of the problem its scale creates. The purpose is not to estimate Revolut's actual numbers — which cannot be done from the public record — but to show why scale alone makes proportionate capacity something that must be demonstrated rather than presumed. Every figure in this section is an illustrative band built from assumption ranges commonly discussed within the AML industry; none is a claim about Revolut's actual operations. The bands are deliberately wide, and that width is itself the finding: the public record does not permit a narrower estimate, and only the firm and its supervisors hold the data that would. **[AO]**

7.2 The model

The arithmetic of transaction monitoring runs in a chain. A monitored transaction population generates alerts at some rate. A proportion of alerts are closed automatically or at first-line review as obvious false positives. The remainder are promoted to cases for investigation. A proportion of cases result in a SAR or STR. Each investigation consumes investigator time, so the case volume implies a required investigator headcount given an assumption about how many cases one investigator can handle in a year. The parameters below are the ranges the industry commonly discusses; their applicability to Revolut is itself uncertain, because it depends on the firm's undisclosed monitoring design. **[AO]**

Parameter	Lower	Mid	Upper
Alert rate (per monitored transaction)	0.10%	0.50%	2.00%
Auto-closure rate (of alerts)	60%	85%	95%
Alert-to-case promotion rate	5%	15%	30%
Case-to-SAR conversion rate	5%	15%	30%
Investigator productivity (cases/FTE/year)	200	500	1,500

Parameter ranges reflect assumptions commonly discussed within the AML industry; they are not independently validated and are not Revolut's figures.

7.3 What the bands imply

Applying these ranges to the disclosed scale (treating the "over 10 billion transactions analysed" figure as an upper bound on the monitored population, itself a generous assumption examined in 7.4) produces the following illustrative bands. The figures are rounded to make clear that no precision is claimed. **[AO]**

Illustrative annual output	Lower	Mid	Upper
Alerts generated	~10m	~50m	~200m
Alerts auto-closed	~6m	~42m	~190m
Investigation cases	~200,000	~1m	~3m
Suspicious reports filed	~10,000	~170,000	~900,000
Investigators implied (mid productivity)	~400	~2,000	~6,000

Illustrative bands only — not Revolut's figures and not a precise model. The width of the bands is the finding.

7.4 "Transactions analysed" is not "transactions monitored"

A critical caution governs the input figure. Revolut reports over 10 billion transactions "analysed," not "monitored against ML/TF typologies." The two are not the same: "analysed" plausibly includes fraud scoring, authorisation checks, and other real-time processing that is not AML transaction monitoring in the regulatory sense. The true AML-monitored population could be smaller. The model does not depend on the exact figure; it depends on the structural point that, at any figure in this range, the implied case and investigator volumes are large enough that proportionate human capacity cannot be assumed without evidence. **[AO]**

7.5 The capacity comparison

The disclosed financial-crime workforce of more than 5,300 is not all investigators; it includes onboarding, fraud, sanctions and support functions. Industry-typical allocation suggests dedicated AML investigators might be a minority of that total — plausibly in the low thousands at most. Set against the bands above, that capacity would comfortably cover the lower case and fall short of the mid and upper cases. The assumption that best reflects the genuine complexity of cross-border investigation — where a single case can span several jurisdictions, languages and counterparties, and cannot be resolved at high throughput — pushes the requirement toward the upper end and the realistic productivity toward the lower end, widening rather than closing any gap. **[EBI regarding the allocation; AO regarding the comparison]**

The honest conclusion is not that capacity is insufficient. It is that the public record does not establish that capacity is **sufficient**, and that insufficiency is consistent with

several plausible configurations. The arithmetic does not prove a Gap; it demonstrates why, at this scale, the absence of disclosed operational metrics leaves the question genuinely open — and why the burden of narrowing the bands lies with those who hold the data. **[AO]**

8. What AI Can and Cannot Be Relied On to Do

8.1 The capability is real

Revolut's strategic answer to the scale problem is, in large part, artificial intelligence, and it credits AI with an "almost 10x" increase in cases reviewed daily. The capability is genuine and the paper does not dispute it: AI meaningfully improves alert triage, pattern detection, network analysis, the summarisation of customer histories, and the drafting of investigation narratives and reports. Supervisors across the FCA, ECB, EBA and FATF endorse its responsible use. The Capacity Gap analysis does not locate its constraint at the capability layer. **[F/AO]**

8.2 The precise limit: evidencing the calibration

The constraint is best stated not as "AI cannot make the decision" — true, but easy to caricature — but in its sharper form. AI may lawfully assist detection, triage, summarisation and drafting; but the regulated institution must be able to **evidence** that its automated prioritisation, its automated closure of alerts, and its automated escalation all remain calibrated to the legal threshold of suspicion — the low, objective standard set out in Section 2. The burden is not on a critic to show the calibration is wrong; it is on the firm to demonstrate it is right, because the threshold the calibration must track is set by law and the resulting decision is owed by an accountable person. AI moves the work; it does not move the standard, and it does not move the accountability. **[F regarding the legal architecture; AO regarding the implication]**

This framing converts the "10x cases reviewed" claim from a reassurance into a definable question. The figure is consistent with two very different realities. In one, AI has prepared roughly ten times more cases for human judgment — which requires the human judgment capacity to have grown correspondingly, or the additional throughput is illusory. In the other, AI is disposing of more cases without a human judgment step — which is lawful only where the closure logic is demonstrably calibrated below the suspicion threshold and the supervisor can inspect that calibration. These have materially different regulatory consequences, and the public record does not distinguish them. **[F regarding the disclosure; OQ regarding the interpretation]**

8.3 The false-negative invisibility problem

This connects to the most consequential invisibility in the field. The errors that matter most in AML are false negatives — suspicious activity the system never flags, closes automatically, or reviews too shallowly to identify. A false positive is visible: it enters a queue, consumes review time, and is closed, leaving a record. A false negative produces no artefact at all, because by definition it is the case the system never surfaced. There is no alert to audit, no case to review, no record that the activity was ever considered. **[AO]**

The consequence is profound for both supervision and self-assessment: an institution can measure its false-positive rate precisely (the queue is the dataset) but cannot measure its false-negative rate without a deliberate testing methodology — seeding known-suspicious patterns, re-reviewing closed populations against enhanced criteria, or comparing detection against external intelligence on confirmed laundering. Whether Revolut conducts such testing, with what methodology and what results, is not disclosed, and is among the most diagnostic of the unanswered questions, because the false-negative rate is the single most direct measure of whether the monitoring system is finding the suspicion the law requires it to find. The Bank of Lithuania's finding — that the firm was not always **identifying** suspicious transactions — is, in this light, an external observation of precisely the failure mode that is internally invisible absent deliberate testing. **[OQ regarding the testing; AO regarding the significance]**

8.4 Four mechanisms by which suspicion fails to reach human judgment

Four general mechanisms, documented across the AML literature and in supervisory findings sector-wide, can cause genuine suspicion to fail to reach the human judgment the law requires. The paper does not assert that any is occurring at Revolut; it asserts that all four are structural risks at any institution running AI-augmented monitoring at scale, that the firm's disclosures do not allow any to be excluded, and that the Bank of Lithuania finding is consistent with at least the first two. These four are the specific failure modes a reader should watch for, and against which any future event should be measured. **[AO]**

The first is **calibration to capacity rather than risk**: thresholds set so that alert volume matches available investigator headcount, rather than so that it captures the activity a reasonable professional would treat as suspicious. Where capacity becomes the calibration driver, the system has implicitly accepted that some suspicious activity will not be surfaced. The second is **typology coverage**: monitoring detects what its scenarios and models are designed to detect, so novel or adversarially-adapted laundering methods, and schemes that do not resemble the institution's existing scenarios or training data, are detected late or not at all — a risk that grows with product and geographic breadth. The third is **auto-closure above**

the threshold: automated closure that disposes of alerts a reasonable professional might have treated as suspicious, where the closure logic is not independently validated against the suspicion threshold, creating a blind spot whose size is invisible to the institution. The fourth is **investigator depth**: where capacity is below the level the case volume requires, the institution faces a choice between queue backlog and reduced review depth, and reduced depth produces false negatives at the judgment step — suspicion that a more thorough review would have identified, missed under time pressure. This fourth mechanism connects directly to the capacity arithmetic of Section 7. **[AO]**

8.5 External-model dependency

Revolut's AI platform includes access to external large-language-model providers alongside proprietary models. Where customer due diligence, screening or case preparation depends on external models, a component of the suspicion-identification pipeline runs on infrastructure the firm does not control and may not be able to fully explain to a supervisor; and the legal requirement to explain why a suspicion finding was or was not reached — readily satisfied for a decision made, far harder for a non-decision produced by a model — intersects with data-protection and banking-secrecy constraints in an area no settled supervisory framework yet governs. This is noted as an emerging structural consideration, not a current deficiency. **[AO]**

9. Suspicious-Activity Reporting in Context

9.1 The red line

This section requires the greatest discipline, because the temptation it must resist is the inference of under-reporting — an inference the public record cannot support and which, if drawn, would convert the paper from analysis into allegation. The line is absolute and is stated at the outset: the paper does not assert, imply, or invite the inference that Revolut under-reports suspicious activity. Revolut's SAR/STR volumes are not disclosed; without them, no assessment of reporting adequacy is possible, and none is attempted. **[AO]**

9.2 The legal relationship, and the proportionality question

What can be stated is the legal relationship established in Section 2. Fraud that completes through Revolut accounts generates criminal property; dealing with criminal property is money laundering; suspicion of money laundering triggers the reporting obligation. The fraud-prevention function and the AML reporting function

are therefore sequential, not parallel: detected-and-completed fraud is an input to the reporting obligation. **[F]**

The external data establishes that Revolut sits at the elevated end of UK fraud-exposure metrics. Drawing on Payment Systems Regulator firm-level data, Financial Ombudsman Service complaint data, and Action Fraud reports: Revolut has appeared among the firms with elevated authorised-push-payment fraud received per unit of activity; and it was the UK firm with the highest volume of fraud complaints to the Financial Ombudsman Service in two consecutive years. These are fraud-exposure metrics, not AML-reporting metrics, and the paper treats them strictly as such. **[F]**

The proportionality question follows, and it is legitimate: given elevated fraud exposure, and given the legal chain that converts fraud into an AML reporting obligation, is the firm's fraud-to-AML pipeline converting that exposure into proportionate reporting output? The question is generated directly by the firm's own disclosed fraud data combined with the settled operation of the law. It is also unanswerable from the public record, because the reporting output is not disclosed. It is an open question, identified for the parties who can answer it. **[OQ]**

9.3 The symmetric invisibility, and the over-reporting counter-incentive

The proportionality question is subject to a symmetry the paper states plainly. Just as the public record cannot establish that the pipeline is deficient, it equally cannot establish that it is adequate. The firm may convert its elevated fraud exposure into fully proportionate, high-quality reporting; the public record is consistent with that. It may not; the public record is equally consistent with that. The invisibility protects the firm from an adverse inference and denies it the ability to demonstrate adequacy through public data. Only a supervisor with access can resolve it. **[AO]**

A sophisticated objection sharpens this further, and the paper incorporates it directly. Large regulated firms face a strong incentive to **over-report** defensively — to file SARs of marginal value to demonstrate diligence and avoid the far greater cost of being seen to under-report. Defensive over-reporting is real and well-documented across the sector. But it does not weaken the thesis; it sharpens it. The Capacity Gap is not a claim about report **volume**; it is a claim about whether **substantive** suspicion is investigated, concluded on and reported. A high volume of low-value defensive reports can coexist with — and even mask — undetected high-value suspicion: the volume tells a supervisor nothing about whether the cases that mattered were caught. This is precisely why the paper benchmarks against the **identification** step (the subject of the Bank of Lithuania finding) rather than against raw report counts, and why it treats no particular reporting volume, high or low, as evidence either way. **[AO]**

10. Supervisory Fragmentation and the US Dimension

10.1 The fragmentation, precisely stated

The proposition that a cross-border fintech can become "too big to supervise" is rhetorically powerful and analytically imprecise, and the paper does not assert it. Dismantled, it reduces to a single defensible claim. It is not that the relevant supervisors are individually incapable — the record shows active, capable engagement by the Bank of Lithuania, the FCA, the PRA, the ECB and the ACPR. It is that the supervisory architecture is structurally fragmented, such that the consolidated, group-wide ML/TF picture is held only by the firm and is seen by no single external authority. The ACPR inspection illustrates both sides: a host-state conduct supervisor identified an issue that drove global remediation — so the architecture is not inert — yet the issue with EEA-wide and global implications was found by a host-state supervisor in a French inspection, rather than by any authority with a consolidated remit, because none exists. **[F regarding the engagement; AO regarding the fragmentation]**

10.2 The US dimension

The most consequential dimension is the United States, and it connects to a specific disclosure gap. US authorities — OFAC, the New York Department of Financial Services, the OCC, FinCEN and the Department of Justice — have repeatedly asserted jurisdiction over non-US banks on the basis that US-dollar transactions clear through the US financial system, bringing the activity within US reach regardless of where the institution or its customers are located. The historical enforcement record is consistent: BNP Paribas (US$8.9 billion, 2014), Standard Chartered, HSBC (US$1.92 billion, 2012) and Danske Bank (a US$2 billion US resolution in 2022 concerning its Estonian branch's dollar flows) were all reached because the activity was dollar-denominated and cleared through the US, not because of where the customers were. In each case the enforcement reach followed the dollar flow. **[F]**

The variable that determines this exposure — the volume of US-dollar transactions flowing through the firm — is not disclosed in Revolut's group accounts. Meanwhile, Revolut already routes dollar activity through a US partner bank (Lead Bank, Kansas City) and, on 5 March 2026, applied to the Office of the Comptroller of the Currency and the Federal Deposit Insurance Corporation for a US national bank charter, to be named Revolut Bank US, N.A. A national charter would deepen its relationship with the US dollar-clearing system and bring it into direct supervision by the most extraterritorially capable AML and sanctions enforcement architecture in global finance. The paper does not assert that Revolut has any US-related exposure; it observes that the variable which would determine such exposure is undisclosed, that the historical record establishes dollar flow as the operative trigger, and that the

firm's trajectory runs toward deeper, not shallower, engagement with that architecture. **[F regarding the record and the disclosure gap; AO regarding the implication]**

11. Historical Analogues: The Documented Shape of Scale Outrunning Control

11.1 Why this section exists

The preceding sections establish a set of conditions. This section asks what those conditions have looked like elsewhere, because the most disciplined way to assess a risk that has not yet materialised is to examine cases where the same conditions did materialise. The analysis that follows is pattern recognition, not prediction. It does not assert that Revolut will follow the path of the institutions described; it identifies the documented shape of how large-scale ML/TF failures have occurred at comparable institutions, so that a reader can judge for themselves how far the conditions resemble those cases — and so that, if a future event ever occurs, the pattern was identified in advance rather than only in hindsight. **[AO]**

11.2 The recurring pattern

The major European ML scandals of the last decade share a structure. Danske Bank's Estonian branch processed an estimated €200 billion of largely non-resident flows over roughly nine years before external exposure; the activity was cross-border, the local entity's controls did not scale with the volume it was handling, and the home-state supervisor did not hold the consolidated picture that would have revealed the scale. Swedbank's Baltic operations followed a comparable shape: rapid growth in a cross-border, non-resident customer base, monitoring that did not keep pace, and a failure that became visible only through external investigation and journalism rather than through the institution's own detection. ABLV in Latvia, and earlier cases such as FBME, repeated the essential features: scale and cross-border complexity outrunning the capacity — and the supervisory visibility — to control it, with the failure invisible from outside until a late and sudden exposure. **[F regarding the cases; AO regarding the common structure]**

The common features are precise, and they map onto the four mechanisms of Section 8.4: rapid growth in a complex, cross-border customer and transaction base; monitoring and investigative capacity that did not scale with that growth; fragmented supervision in which no single authority held the consolidated picture; and — decisively — a failure that generated no visible external signal until exposure came from outside the institution. In each case, the absence of a visible problem was

not evidence that the controls were working; it was a property of the false-negative invisibility described in Section 8.3. The institutions did not look, from outside, as though they were failing. That is precisely the point. **[AO]**

11.3 The warning, stated as risk and not prediction

The conditions this paper documents at Revolut — exceptional scale and growth velocity, product and geographic complexity, a cross-border structure in which no single supervisor holds the consolidated picture, an undisclosed operational capacity, a supervisor's finding that suspicious transactions were not always identified, and a control model increasingly dependent on automation at the detection stage — are the same conditions that were present, and documented after the fact, in the cases above. **[AO]**

This is the paper's warning, and its limits are as important as its substance. The paper does not predict an ML/TF failure at Revolut. It has no evidence that one is occurring, and the firm may have operational capacity, fully proportionate to its risk, that the public record simply does not show. What the paper asserts is narrower and, for that reason, durable: that the conditions under which undetected large-scale failures have historically occurred are present and documented here; that the false-negative invisibility which characterised those cases means no advance external assurance exists that a failure, if it were occurring, would be visible before late exposure; and that the appropriate response is therefore demonstration of proportionate capacity, not presumption of it. A reader who absorbs this section before any future event will know where the structural vulnerability would lie, which of the four mechanisms would manifest it, and why it would not be visible from outside until late. If such an event never occurs, this section will have been a documented and falsifiable caution that the firm and its supervisors were well-placed to answer. If one ever does, the pattern will have been identified in advance, on the evidence, without accusation. Either way, the warning is on the record. **[AO]**

PART IV — SYNTHESIS AND CONCLUSIONS

12. Assessment: The AML Capacity Gap

The assessment has three sides: the requirement, the capacity, and the relationship between them.

The **requirement** is the firmest finding in the paper, and is built almost entirely from Revolut's own disclosures and official external data: an exceptional and still-accelerating scale (Section 3); an unusually broad product set, each line carrying its own typology surface (Sections 4–5); a wide, demanding and multi-jurisdictional footprint that multiplies the capacity required (Sections 4, 10); and a fraud-exposure profile at the elevated end of every available external measure, generating an elevated reporting obligation by operation of law (Section 9). This requirement profile is exceptional and is not seriously contestable. **[F/AO]**

The **capacity** is the side the public record cannot establish. The control architecture is disclosed in qualitative terms (Section 5); the operational metrics that would show whether it functions at proportionate scale — investigator headcount, alert and report volumes, auto-closure calibration, false-negative testing, the operational meaning of the AI productivity claim — are not (Sections 5, 7, 8, 9). The arithmetic of Section 7 establishes that, under several plausible configurations, the disclosed workforce is below the implied requirement; it does not establish that capacity is insufficient, only that the public record does not establish it is sufficient. **[AO]**

Between the two sit **five convergent signals**, each individually bounded, none alone decisive, pointing consistently in one direction: the disclosure asymmetry (Section 5); the convergence of three supervisory and audit findings in 2025 (Section 6); the Bank of Lithuania finding on the identification of suspicion specifically (Sections 6, 9); the elevated fraud exposure that generates an elevated reporting obligation (Section 9); and the AI limit at the human-judgment step (Section 8). **[F/AO/OQ as labelled in the body]**

These signals converge on a question, not a conclusion, because the disclosure gap is **symmetric**: its absence prevents anyone outside the firm from proving a Gap exists, and equally prevents the firm from demonstrating one does not. The failures that would evidence a Gap are structurally invisible (Section 8.3). The question is genuinely open in both directions, and the openness is itself produced by the lack of disclosure. **[AO]**

But the requirement and the capacity are not symmetrically placed. The requirement is established and exceptional; the capacity is undisclosed; five signals

converge; and the historical record (Section 11) shows that these precise conditions have, elsewhere, accompanied undetected failure. When an exceptional, well-evidenced requirement meets an undisclosed capacity against that background, the appropriate conclusion is not that the Gap is proven, and not that it is disproven, but that proportionate capacity should be demonstrated rather than assumed — and that, given the asymmetry of access to operational data, the practical ability to resolve the question rests primarily with the firm and its supervisors. The public record does not contain the demonstration. **[AO]**

The conclusion is deliberately bounded. It does not assert that a material Gap exists; it does not infer control failure from the supervisory findings, each of which is bounded and partly remediated; and it does not infer under-reporting. It asserts only that the convergence makes the question pressing rather than theoretical, and that the evidence which would settle it is held by those best placed to provide it. The case study establishes the concept; it does not convict the subject. **[AO]**

13. Counterarguments and Falsifiability

A conclusion that cannot withstand its strongest counterarguments is not analysis. This section states the strongest objections at full strength and concedes where the concession is genuine. **[AO]**

"You are describing the risk-based approach itself." The risk-based approach is, by design, a system of filtering and prioritisation; every institution filters, prioritises and accepts residual risk. To call this a defect is to mislabel the regulator-endorsed standard. **Conceded in full.** What survives is the distinction the paper holds throughout: filtering is not the issue; the issue is whether the filtering is calibrated to **risk** (legitimate) or to **capacity** — set so that alert volume matches headcount rather than the level of suspicion the law requires. That a system filters tells us nothing about which driver governs it, and that driver is undisclosed. The objection defeats a claim the paper does not make. **[AO]**

"Absence of evidence is not evidence." Most institutions do not disclose AML metrics; treating Revolut's non-disclosure as significant holds it to a standard its peers escape. **Conceded.** Non-disclosure is the sector norm and is not, by itself, evidence of a Gap, and the paper never claims it is. What survives is narrower: the **asymmetry** between detailed fraud disclosure and absent AML disclosure within the same reports is distinctive even where each half is unremarkable; and the conclusion rests not on non-disclosure as evidence of a Gap but on the requirement-plus-signals combination, with non-disclosure explaining only why the question cannot be

settled externally. Non-disclosure is identified as the thing the firm could remedy, not as proof of anything. **[AO]**

"The supervisory findings show the system working." The Bank of Lithuania identified issues and fined the firm; the ACPR drove remediation; the PRA granted authorisation after a thorough process. This is supervision functioning as intended, not evidence of a Gap. **Conceded, and important.** The paper does not cite the findings as proof of a Gap; it cites their convergence as one of five signals, and the Lithuanian finding for what the supervisor found — that suspicion was not always identified — regardless of subsequent remediation. Remediation of identified instances does not establish that the underlying cause was resolved. Supervision functioning and the question remaining open are not mutually exclusive; both are true. **[AO]**

"Large firms over-report defensively, so the real risk is too many reports, not too few." The objection a sophisticated AML reader raises fastest, and conceded as a real and well-documented phenomenon. But it sharpens rather than weakens the thesis: the Capacity Gap is about whether substantive suspicion is investigated and reported, not about report volume, and defensive over-reporting can coexist with and even mask a Gap. This is why the paper benchmarks against the identification step, not raw report counts (Section 9.3). **[AO]**

"The dynamics are sector-wide, so singling out Revolut is unfair." Conceded, and it is the paper's own position. The dynamics are sector-wide; Revolut is the case study because it is the clearest available public example, not because it is uniquely deficient. A concept is best introduced through its clearest case, and the paper is explicit that the concept applies across the sector and to the supervisory architecture being built around it. **[AO]**

"AI may close the gap faster than the analysis assumes." Conceded as a genuine contingency. The human-judgment limit is a feature of current law and supervisory expectation, not a permanent fact; if the frameworks evolve to permit accountable AI decision-making at the suspicion stage, the analysis would require revision. The paper assesses the constraint as it currently exists and frames it as contingent; the present-tense conclusion stands on the present-tense frameworks. **[AO]**

The falsifiability set. The thesis is testable, and would be substantially refuted — the question resolved in the firm's favour — by disclosure or supervisory confirmation of: SAR/STR volumes proportionate to the firm's fraud exposure and peers; independently validated false-negative testing showing low miss rates; investigator capacity matched to the case volumes the scale generates, at adequate review depth; monitoring demonstrably calibrated to risk rather than capacity, with auto-closure validated below the suspicion threshold; and confirmation, by a supervisor with consolidated group-wide access, that operational capacity is

proportionate. The paper invites the firm and its supervisors to produce exactly this. A thesis with specified falsifiers is analysis, not accusation. **[AO]**

14. Public-Market Scrutiny as a Forcing Function

A public listing, which several disclosed signals suggest the firm is preparing for — the EY auditor transition from the 2026 financial year, disclosed internal-controls readiness work, and the US charter application — would operate, for the matters this paper analyses, as a forcing function. Pre-listing due diligence, listing-prospectus risk disclosure subject to securities-law liability, ongoing public-company reporting, and (for a US listing) internal-control attestation would compel the firm to assemble, examine and partially disclose precisely the evidence the public record currently lacks. A US listing or charter path would converge the most demanding disclosure regime with the most far-reaching enforcement architecture. The paper does not predict the outcome; it observes that the listing process is the structural event most likely to compel the assembly and partial disclosure of the evidence that would resolve the question this paper raises — and that the question is therefore likely to be engaged, through that process, whether or not supervisors engage it first. **[F regarding the signals; OQ regarding the outcome]**

15. Recommendations and Open Questions

The recommendations are framed neutrally: as opportunities to demonstrate, not as assumptions of failure. **[AO]**

The **firm** holds the data that would settle the question. It could demonstrate proportionate capacity — to its supervisors as a matter of course, and to the extent it judges appropriate to the public — by producing the evidence in the falsifiability set. It is better placed than any other party to do so, because it holds the evidence. **[AO]**

The **supervisors** — the Bank of Lithuania and the ECB in respect of the EEA bank, the UK regulators, and the EU's new central AML authority (AMLA) as it begins direct supervision of selected institutions from 2028 — hold the powers to inspect investigator capacity, alert and report volumes, the calibration of automated closure against the suspicion threshold, and false-negative testing, and are the only parties able to assess the consolidated group-wide picture. A specific and pressing challenge is whether the supervisory community will develop, in time, the technical capability to inspect AI-driven AML systems to the depth their own expectations require — a capability AMLA is, in part, intended to build. **[AO/OQ]**

The **policy architecture** faces a question the case study sharpens: whether the largest obliged institutions should be expected to disclose summary AML operational metrics — such as alert-to-report conversion ranges, model-driven auto-closure rates, and investigator headcount per unit of scale — sufficient to allow proportionality to be assessed without compromising the confidentiality of specific suspicions or tipping off criminals. Such disclosure would let the largest firms demonstrate proportionate capacity publicly, narrow the symmetric information gap this paper repeatedly encounters, and bring AML disclosure closer to the standard those same firms already apply to fraud. The costs are real and would have to be weighed; the paper raises the question and does not prescribe the answer. **[AO/OQ]**

The consolidated open questions, each identified for the party positioned to answer it: whether the consolidated group risk assessment addresses the heterogeneous footprint with internal consistency; what the firm's SAR/STR volumes are and how they relate to its fraud exposure; what proportion of "transactions analysed" is subject to AML monitoring; what the alert, auto-closure and conversion rates are, and on what basis auto-closure is validated against the suspicion threshold; what proportion of the financial-crime workforce are AML investigators; whether the "almost 10x" AI gain represents capacity in suspicion judgment or throughput in preparation; whether the firm conducts false-negative testing, and with what results; whether operational capacity is proportionate to requirement; what the model-level cause of the Bank of Lithuania finding was, and whether remediation addressed the cause rather than the instances; and whether the supervisory architecture, including AMLA, will arrive in time with the technical capability and the consolidated visibility the question requires. **[OQ]**

Closing

The AML Capacity Gap is a defined, observable phenomenon: the divergence between the size, complexity and risk of an institution's customer and transaction population and its capacity to discharge the resulting obligations. Revolut exhibits every condition under which the Gap arises — an exceptional, well-evidenced requirement; an undisclosed operational capacity; a convergence of supervisory and audit signals; a symmetric information gap; the human-judgment limit that artificial intelligence cannot relieve; and a cross-border structure in which no single supervisor holds the consolidated picture. The paper does not conclude that a material Gap exists. It concludes that the question is pressing rather than theoretical, that proportionate capacity should be demonstrated rather than assumed, and that the practical ability to resolve it rests primarily with the firm and its supervisors. **[AO]**

This is, finally, a warning rather than a verdict — and the distinction is the point. The conditions documented here are the conditions under which large-scale ML/TF failures have historically occurred and remained invisible until late external exposure. The paper does not predict that outcome at Revolut and offers no evidence that it is occurring. It establishes, on the public record and in advance of any event, where the structural vulnerability would lie, what would manifest it, and why it would not be visible from outside until late. A reader who had this paper before any future event would have seen the warning; a firm and a set of supervisors who have it now have a specified, falsifiable set of questions they are uniquely placed to answer. If the question is answered and proportionate capacity is demonstrated, the paper will have prompted a useful demonstration and been refuted in the firm's favour. If it is not, the warning will have been on the record, stated as risk and bounded by evidence, without ever crossing into accusation. **[AO]**

The proportionality of AML capacity to fintech scale is the central supervisory question of European fintech in the late 2020s. The answer to it — for Revolut, for its peers, and for the supervisory architecture being built around them — will define what proportionate AML at scale means in practice. **[AO]**

BACK MATTER

Glossary

ACPR — Autorité de contrôle prudentiel et de résolution, the French prudential and conduct supervisor.

AML / CFT — Anti-Money Laundering / Combating the Financing of Terrorism.

AML Capacity Gap — As defined in this paper: the structural divergence between the size, complexity and risk of an institution's customer and transaction population and its investigative, supervisory and governance capacity to identify, assess, manage, mitigate and report ML/TF risk under the risk-based approach.

AMLA — The EU Anti-Money Laundering Authority, which begins direct supervision of selected high-risk obliged entities from 2028.

Auto-closure — The automated disposition of a monitoring alert without human review; supervisorily acceptable only where validated as operating below the suspicion threshold.

BaFin — Germany's Federal Financial Supervisory Authority.

Bank of Lithuania — The central bank and AML supervisor of Revolut Bank UAB.

ECB SSM — The European Central Bank Single Supervisory Mechanism, which assumed direct prudential supervision of Revolut Holdings Europe UAB from 1 January 2024.

EMI — Electronic Money Institution, the status under which Revolut Ltd has operated since 2015.

False negative — Suspicious activity a monitoring system fails to detect; structurally invisible absent deliberate testing.

FATF — The Financial Action Task Force, the global AML standard-setter; Recommendation 1 establishes the risk-based approach.

FCA / PRA — The UK Financial Conduct Authority (conduct and AML) and Prudential Regulation Authority (prudential).

FIU — Financial Intelligence Unit (the UK's National Crime Agency; Lithuania's FCIS; the US FinCEN).

FRC — The UK Financial Reporting Council, which found Revolut's 2020 audit inadequate.

MLRO — Money Laundering Reporting Officer, the personally-accountable individual responsible for the suspicion-reporting decision.

OCC — The US Office of the Comptroller of the Currency, to which Revolut applied for a national bank charter in March 2026.

POCA — The UK Proceeds of Crime Act 2002; section 330 (failure to disclose), section 340 (criminal property).

Risk-based approach — The FATF-mandated framework under which controls are calibrated in proportion to assessed risk.

SAR / STR — Suspicious Activity Report (UK) / Suspicious Transaction Report (EU), filed to the national FIU on suspicion of money laundering.

Suspicion threshold — The low, objective standard (R v Da Silva: "more than a fanciful possibility") that triggers the reporting obligation.

Sources and References

This paper is based exclusively on publicly available information, verified as at the date of writing (mid-2026). Principal claims are referenced below to named, dated sources.

Revolut scale and 2025 results

1. Customer base (68.3 million retail customers at year-end 2025, up 30% from 52.5 million; over 70 million by early 2026; 16 million added in 2025); revenue £4.5 billion; pre-tax profit £1.7 billion (up 57%); business customers 767,000; total balances £50.2 billion; lending £2.2 billion — Revolut Group Holdings Ltd, Annual Report 2025 (assets.revolut.com/pdf/annualreport2025.pdf) and Revolut press release, "Revolut reports record profit of $2.3bn for 2025" (24 March 2026); corroborated by Reuters and the Guardian (March 2026).
2. "Over 10 billion transactions analysed" in 2025; "over a third of the global workforce" in Financial Crime Prevention; Risk and Compliance functions grew "more than 42%"; AI enabling an "almost 10x" increase in cases reviewed daily; 200+ fraud specialists; fraud prevented (£475m in 2023; £600m in 2024) — Revolut Annual Report 2025 and prior-year annual reports.
3. Operating in 40 markets, licensed as a bank in over 30; ambition of 100 million customers by mid-2027; ~$75 billion valuation (November 2025 secondary share sale) — Revolut Annual Report 2025; Revolut corporate statements; Reuters; The Paypers (March 2026).

Regulatory and supervisory facts

4. UK full banking authorisation: the PRA lifted mobilisation restrictions on Revolut Bank UK Ltd on 11 March 2026, following a mobilisation period of roughly 20 months (restricted authorisation July 2024; application 2021); approximately 13 million UK customers — Revolut, "Revolut Launches UK Bank" (11 March 2026); reported by Tech.eu, The Block, MoneyWeek, Retail Banker International.
5. Bank of Lithuania €3.5 million fine, Revolut Bank UAB, 8 April 2025, following a scheduled inspection; stated by the Bank of Lithuania to be the largest fine it has ever issued; finding of "violations and shortcomings in the monitoring of business relationships and operations (transactions)" that "led to (Revolut) not always properly identifying suspicious monetary operations or transactions"; no confirmed money laundering; resolved by administrative agreement and corrective plan — Bank of Lithuania statement (8 April 2025); Reuters.
6. Four Bank of Lithuania enforcement actions 2022–2025 totalling roughly €3.82 million (two AML-related, two relating to the Capital Requirements Regulation / late reporting) — Bank of Lithuania records; sector analysis.
7. ACPR inspection driving EEA-wide and global customer-due-diligence remediation; internal-audit financial-crime findings under a Board-monitored Control Enhancement Plan — Revolut Annual Report 2025 (the firm's own disclosure).
8. ECB assumption of direct prudential supervision of Revolut Holdings Europe UAB from 1 January 2024 — Revolut Annual Report 2025; ECB/SSM.
9. US national bank charter application filed 5 March 2026 with the OCC and FDIC, to be named Revolut Bank US, N.A.; current US services via partner bank Lead Bank (Kansas City) — Revolut, "Revolut Files U.S. Bank Charter Application" (5 March 2026); American Banker, Banking Dive, PYMNTS, FinTech Futures.
10. Italian Antitrust Authority (AGCM) fine of €11.5 million (2 April 2026), under appeal — AGCM decision; press reporting.
11. Audit history: FRC finding that Revolut's 2020 audit (BDO) was "inadequate" with "the risk of an undetected material misstatement … unacceptably high" (FRC, July 2022); BDO qualified opinion on the 2021 accounts, unable to verify £477 million of revenue (~75% of £636 million) due to IT systems design, accounts filed ~5 months late; BDO fee £4.5 million for 2021; EY to replace BDO from FY2026 — Financial Times, Irish Times, CityAM, PYMNTS (2022–2023); Revolut announcement (August 2025).

External financial-crime data

12. UK fraud-exposure data placing Revolut at the elevated end of its peer group: Payment Systems Regulator firm-level APP fraud data (2022–2024); Financial Ombudsman Service data showing Revolut as the highest UK firm for fraud complaints in two consecutive years; Action Fraud report counts — PSR; FOS (via Which? and FOI disclosures); BBC.

Legal framework and comparators

13. Proceeds of Crime Act 2002 (UK), sections 330 and 340; R v Da Silva [2006] EWCA Crim 1654; the EU AML framework as transposed in Lithuania; UK Senior Managers and Certification Regime; FATF Recommendations 1 and 20; National Crime Agency SARs Annual Report 2023/24; Lithuanian FCIS statistics and 2024 National Risk Assessment.
14. Comparators and historical analogues: Starling Bank, FCA Final Notice, £28,959,426 (27 September 2024); N26, BaFin, €9.2 million (9 May 2024) and €4.25 million (2021); BNP Paribas (2014), Standard Chartered, HSBC (2012), Danske Bank (2022 US resolution), ABLV and FBME — FCA; BaFin; US Department of Justice and regulatory records; contemporaneous reporting.

Appendix A — Quantitative Model Parameters

This appendix sets out the parameters underlying the illustrative bands in Section 7. Every figure is an assumption range commonly discussed within the AML industry, not a claim about Revolut's actual operations, and not independently validated. The bands are illustrative of order of magnitude only.

Parameter	Lower	Mid	Upper
Alert rate (per monitored transaction)	0.10%	0.50%	2.00%
Auto-closure rate (of alerts)	60%	85%	95%
Alert-to-case promotion rate	5%	15%	30%
Case-to-SAR conversion rate	5%	15%	30%
Investigator productivity (cases/FTE/year)	200	500	1,500
Monitored population (upper bound)	10bn transactions analysed (2025)		

The "transactions analysed" figure is treated as an upper bound on the monitored population; the true AML-monitored population may be smaller (Section 7.4). The width of the resulting bands reflects genuine uncertainty arising from the disclosure gap, and is the section's central point.

Appendix B — Supervisory and Enforcement Matrix

This appendix records public supervisory and enforcement matters involving Revolut entities. Inclusion does not imply that any matter constitutes evidence of money laundering or criminal conduct; each is recorded as the relevant authority or source has characterised it.

Date	Authority	Matter	Outcome
2018	FCA (UK)	Sanctions-screening period without effective filtering; self-reported	No enforcement action
2022	FRC (UK)	2020 audit (BDO) "inadequate"	Published finding (revenue / IT)
2022	Bank of Lithuania	AML internal-control deficiencies	€50,000
2022	Bank of Lithuania	Late statutory financial statements	€70,000
2023	BDO (auditor)	Qualified opinion on 2021 accounts (~75% of revenue)	Qualified opinion; ~5 months late
Jan 2024	Bank of Lithuania	CRR large-exposure breach; self-reported	€200,000
Jan 2024	ECB SSM	Direct prudential supervision assumed	SI supervision (AML stays with BoL)
Jul 2024	PRA (UK)	Banking licence with mobilisation	~20-month mobilisation
2025	ACPR (France)	Inspection of French operations	EEA-wide / global CDD remediation
8 Apr 2025	Bank of Lithuania	"Not always properly identifying suspicious … transactions"	€3.5m; stated largest ever by the authority; no confirmed ML
2025	Internal Audit / Board	Financial-crime findings	Control Enhancement Plan
11 Mar 2026	PRA (UK)	Mobilisation restrictions lifted	Full authorisation (positive milestone)
5 Mar 2026	OCC / FDIC (US)	National bank charter application	Pending
2 Apr 2026	AGCM (Italy)	Consumer-protection matters	€11.5m; under appeal

Cumulative Bank of Lithuania penalties 2022–2025: approximately €3.82 million. The pattern is one of continuous and escalating engagement, with three financial-crime signals concentrated in 2025, balanced by the positive milestone of full UK authorisation in 2026.