



BIS Bulletin

No 134

Supervisory screening with large language models:
finding divergences

Fernando Perez-Cruz, Kumar Rishabh and Ayush Uchil

26 August 2026

BIS Bulletins are written by staff members of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks. The authors are grateful to Mathias Drehmann, Johannes Ehrentraud, Jon Frost, Gaston Gelos, Pablo Hernández de Cos, Jermy Prenio, Daniel Rees, Bertrand Rime, Vatsala Shreeti, Frank Smets, Nikola Tarashev, Kevin Tracol, Christian Upper and Leanne Zhang for helpful comments, to Fanni Leppanen and Jiaqi Zhang for research assistance, to Nicola Faessler and Danielle Ritzema for administrative support and to Alison Arnot for editing.

The editors of the BIS Bulletin series are Gaston Gelos and Frank Smets.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2026. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN: 2708-0420 (online)

ISBN: 978-92-9259-979-9 (online)

Supervisory screening with large language models: finding divergences

Key takeaways

- *We develop a procedure based on large language models to help supervisors review prospectuses against capital rules. This procedure ranks potential divergences for their consideration. We apply it to Additional Tier 1 (AT1) prospectuses from European global systemically important banks.*
- *We validate the procedure on anonymised documents from Credit Suisse and Yes Bank, recovering rule-prospectus differences that later mattered in disputes over AT1 loss absorption.*
- *This procedure could support other supervisory comparisons, including reviews of national rules against the international standards they implement.*

When a bank issues an instrument intended to count as regulatory capital, supervisors assess whether the proposed terms meet the applicable capital rules. Those proposed terms are set out in the instrument's prospectus. They govern how and when it absorbs losses and establish its seniority relative to other capital instruments. Supervisory review therefore involves two texts: the regulation, which states the conditions an instrument must satisfy; and the prospectus, which sets out the contractual terms on which it is issued. Differences between them are common. Most prospectuses add detail or reflect the discretion allowed by the regulation. A few, however, may change how the rule operates.

Distinguishing permissible variation from wording that may contravene the rule requires supervisory expertise. This Bulletin describes a large language model (LLM)-based procedure that assists supervisory review by directing supervisors to the passages that are most likely to matter. The LLM compares the applicable regulation with the corresponding contractual terms in a prospectus. It identifies provisions that address the same requirement, flags wording that may alter its meaning and ranks candidate divergences by how consistently they recur across separate model runs. Supervisory and legal experts can then assess whether a flagged difference is permitted, material or actionable.

We apply the procedure to Additional Tier 1 (AT1) instruments issued by 10 global systemically important banks (G-SIBs), comparing each prospectus with the applicable capital rules. The procedure returns 79 recurrent candidate divergences for expert review. We also test the procedure on AT1 instruments of Credit Suisse in Switzerland and Yes Bank in India. Here, divergences between prospectuses and applicable rules later became consequential for investors and featured in legal challenges.

Screening prospectuses against capital rules

The procedure compares the full regulatory package governing an instrument with its prospectus. For each comparison, the analyst supplies the regulatory documents, the prospectus and a structured

prompt.¹ The LLM then identifies corresponding provisions in the regulatory package and prospectus, quotes the relevant passages and flags wording in the prospectus that may alter the requirement’s scope, conditions or consequences. It assigns each candidate divergence to one of five types:


- **Narrowing.** The prospectus adds a condition or limitation not imposed by the rule.
- **Expansion.** It extends the requirement beyond the circumstances covered by the rule.
- **Omission.** It omits a condition, trigger, procedural step or other element of the rule.
- **Ambiguity.** The prospectus’s wording leaves it unclear how it relates to the rule.
- **Contradiction.** The two texts prescribe inconsistent conditions, sequencing or consequences.

These labels describe how the texts differ; they do not establish that the prospectus breaches the rule. The Online Annex illustrates each type with examples from the analysis.

Because LLM outputs are stochastic, identical inputs can yield different responses. To address this, we repeat the procedure five times for each regulatory package–prospectus comparison. Each run uses the same texts and prompt and has no access to the outputs of the other runs. We consolidate equivalent findings and rank them by how often they recur across the five runs. The resulting screening file records the paired passages, divergence type, recurrence and a short explanation of the difference. Supervisory experts can then decide whether the differences are permitted, material or actionable. Graph 1 summarises the process from document comparison to a ranked queue for expert review.

LLM screening ranks candidate divergences for expert review

Graph 1



Source: Authors’ illustration.

Screening AT1 prospectuses at scale

We first apply the procedure to AT1 instruments issued by 10 European G-SIBs.² AT1 instruments are perpetual securities that count towards banks’ regulatory capital under the Basel III regulatory framework. In normal times, they trade much like subordinated debt, paying discretionary, non-cumulative coupons. Under stress, AT1s support the issuing bank in two distinct ways. First, cancelling coupons conserves capital. And second, once a contractual trigger is breached or the authorities determine that the bank has reached the point of non-viability, the instruments are written down or converted into common equity,

¹ Online Annex A provides the structured prompt used throughout the analysis.

² Groupe Banque Populaire Caisse d’Epargne, the remaining European G-SIB, had no AT1 instrument outstanding at the consolidated level at the time of our review and is therefore excluded.

thereby imposing losses on investors. Coelho et al (2026) study how these features should be designed and the trade-offs among them.


AT1 instruments are characterised by four sets of terms:

- **Mechanical triggers** set the capital threshold at which loss absorption begins.
- **Point of non-viability** (PONV) provisions specify the circumstances in which an authority can require loss absorption.
- **Coupon** terms govern when payments can be cancelled without default.
- **Hierarchy** determines whether common equity or AT1 bears losses first.

The procedure helps supervisors to answer the contractual implementation question: once the applicable rules define these features, does the prospectus preserve them, or does it add contractual conditions that may constrain loss absorption or change how losses are allocated?

The 10 G-SIBs comprise six European Union (EU) banks, three United Kingdom (UK) banks and one Swiss bank, spanning three regulatory regimes. We compare each bank’s AT1 prospectus against the applicable capital rules: the Capital Requirements Regulation together with the Bank Recovery and Resolution Directive in the EU, the Prudential Regulation Authority Rulebook in the UK, and the Swiss Capital Adequacy Ordinance in Switzerland. Across the five independent runs,³ we consolidate equivalent findings and retain 79 candidate divergences that appear in at least four runs,⁴ which are depicted in Graph 2.

Candidate divergences across G-SIB AT1 instruments Graph 2



PONV = point of non-viability. Panels A and B show the counts for the 79 candidate divergences, grouped by AT1 feature (panel A) and by jurisdiction (panel B), each split by divergence type. Bar heights are counts.

Source: Authors’ calculations.

The flagged divergences cluster in the provisions governing loss absorption and vary in type across jurisdictions. Grouped by AT1 feature (Graph 2.A), mechanical triggers and PONV account for roughly two

³ We use GPT-5.5 Pro for the results reported here. The prompt can be used with any other state-of-the-art LLM, although results will vary according to the LLM’s capabilities.

⁴ The procedure returned 576 individual divergences. After equivalent findings within and across runs were merged, these yielded 193 distinct candidate divergences: 56 appeared only once, 31 in two runs, 27 in three, 34 in four, and 45 in all runs. The supervisor receives an ordered list of potential issues; the graphs below are based on the 34 + 45 recurrent discrepancies.

thirds of the total, concentrating the queue on provisions that determine when loss absorption begins, either mechanically or through supervisory judgment. All five types of divergence are present, but their mix differs across AT1 contract terms. Grouped by jurisdiction (Graph 2.B), most common are narrowings in Switzerland, expansions in the UK and omissions in the EU. Since the number of banks differs across jurisdictions, these counts describe the composition of flagged candidates rather than their prevalence across regulatory regimes. Whether a flagged difference is legally meaningful or warrants supervisory action remains a matter for expert review.

Testing the procedure on known AT1 divergences

A further test of the procedure's supervisory value is whether it can identify rule-prospectus differences that later proved consequential. Credit Suisse and Yes Bank provide two such cases. In both, the wording in AT1 contracts written years earlier became consequential only when authorities had to determine whether and how losses would be imposed. Using documents available before those events, we apply the procedure specified above and test whether it identifies relevant differences and classifies them correctly.

The historical validation raises an additional concern: the LLM may already have encountered these cases in its training data or because it searched the internet for them. We therefore remove names, dates and other direct identifiers from each regulatory package and prospectus before applying the procedure. Supervisors applying the procedure *ex ante* would not need to do this step, and we did not do it in the previous exercise. Anonymisation alone cannot rule out prior model knowledge, so Online Annex C tests whether the findings depend on the supplied text. On the original anonymised documents, the procedure identifies the relevant divergence in both cases. When we separately modify the relevant wording in the prospectus and the regulatory text to bring the two into agreement, the model no longer returns the original finding. This illustrates that the finding responds to the wording supplied to the model rather than simply reproducing the known case.

Credit Suisse dispute: a narrower public support condition for PONV

The Credit Suisse case concerns how the prospectus defined public support under the PONV trigger. Article 29 of the Swiss Capital Adequacy Ordinance requires AT1 terms to permit full write-off or conversion at PONV, at the latest before the bank receives public support or when the Swiss Financial Market Supervisory Authority (FINMA) orders the measure to avert insolvency (Swiss Federal Council (2012)). Relative to Article 29, the Credit Suisse prospectus added a capital adequacy condition: the support had to improve, or be about to improve, the bank's capital adequacy. The procedure flags this condition and classifies it as a narrowing in all five runs (Graph 3.A).

That narrowing became consequential when Credit Suisse received extraordinary liquidity assistance in March 2023. On 19 March, FINMA ordered the full write-down of Credit Suisse's AT1 instruments, with a nominal value of about CHF 16.5 billion, as part of the measures enabling its takeover by UBS (FINMA (2023)). FINMA treated the assistance that Credit Suisse received as satisfying the contractual viability event; bondholders challenged that interpretation. The bondholders' legal challenge is still ongoing. Regardless of its eventual outcome, the case shows how a rule-prospectus divergence can remain consequential after issuance and become a focus of legal dispute.

Yes Bank: a regulatory change in loss-absorption sequencing

The Yes Bank case provides a different validation test because the applicable capital rule changed between bond issuances. Yes Bank issued three AT1 bond tranches in 2013, 2016 and 2017, under two versions of the Reserve Bank of India (RBI) capital framework. The 2013 bond prospectus reflected the rules then in force and required common equity to absorb losses before the bond could be written down. In 2015, the RBI amended the framework so that common equity no longer had to absorb losses before AT1 could be written down. The 2016 and 2017 bonds reflected the amended framework. By 2020, Yes Bank's outstanding AT1 bonds therefore differed on whether common equity had to absorb losses first.


For validation, we compare all three prospectuses with the RBI framework in force in 2020 (RBI (2015)). The procedure flags the 2013 wording in all five runs and classifies it as a contradiction: the RBI framework did not require common equity to absorb losses first, while the prospectus did. The 2013 prospectus complied with the rules in force when the bond was issued. However, because the bond remained outstanding after the framework changed, it retained a sequencing condition that was no longer present in the RBI framework in force in 2020. The procedure does not return the same finding for the 2016 and 2017 bonds, whose terms reflected the amended framework.

The difference mattered when the RBI intervened in March 2020. The RBI placed Yes Bank under moratorium and approved a restructuring plan in which the State Bank of India and other investors recapitalised the bank. The 2016 and 2017 AT1 bonds, with a combined value of INR 84.15 billion (approximately USD 1.1 billion at the time), were written off. The 2013 bond was not.

AT1 wording differences become consequential under stress

Graph 3

A. Credit Suisse: a qualification points to a potential narrowing of the public support trigger¹



B. Yes Bank: a dormant sequencing clause was priced only as stress approached²



¹ Identifiers have been restored for presentation; the comparison procedure was run on anonymized documents. For complete underlying text behind the comparison in Panel A, see Online Annex B Graph A.1. CAO = Swiss Capital Adequacy Ordinance; CS = Credit Suisse; CSG = Credit Suisse Group. ² The lines show average yields across the relevant AT1 instruments in each category.

Sources: LSEG, authors' calculations and illustrations.

Markets placed little value on the 2013 bond's stronger contractual protection until Yes Bank came under stress. The 2013 AT1 bond traded until 2019 at yields close to those of the later instruments despite its greater contractual protection (Graph 3.B). The gap widened as the bank's condition deteriorated, then increased further before the write-off on 14 March 2020. The pattern is consistent with investors pricing the sequencing clause only as its application became more likely, in line with broader evidence on the repricing of loss-allocation rules under stress (Di Stefano et al (2026)).

Broader applications and implications for supervisory review

Using this procedure in practice could help to allocate supervisory expertise more efficiently. The ranked queue changes the point at which expert review begins. Supervisors start from recurrent clause pairs linked to their source texts, rather than from an undifferentiated mass of rules and prospectuses. That matters because supervisory expertise is scarce and its allocation shapes which risks receive attention (Eisenbach et al (2022)).

That shift changes the division of labour between the model and supervisory experts. The LLM performs the initial comparison and flags candidate differences; supervisors and legal experts decide

whether a difference is permitted, material or actionable. Supervisors are already using such tools to process large volumes of text, although their move from experimentation to routine supervision remains incomplete (Prenio (2024); Crisanto et al (2025)).

The same division of labour could support other supervisory tasks. For example, under the Basel Committee on Banking Supervision's Regulatory Consistency Assessment Programme, an assessment team conducts a peer review of domestic capital rules against the Basel standards. Under the Financial Sector Assessment Program (FSAP), the International Monetary Fund (IMF) and the World Bank assess national supervisory frameworks against the Basel Committee's Core Principles for Effective Banking Supervision. Mutual evaluations of the Financial Action Task Force (FATF) check national law against the anti-money laundering and combatting the financing of terrorism recommendations.

Operationalising our procedure for routine supervision would first require testing against document pairs independently assessed by supervisory and legal experts. Authorities would need to know how often highly ranked candidates are judged as material and how often material differences fail to enter the queue. Those results would inform the recurrence threshold, the amount of full-document review that remains necessary and the tasks for which screening is suitable. This would establish when the queue provides a reliable starting point for expert judgment. Early screening would also allow supervisors to assess contractual uncertainty before stress exposes its consequences.

References

Coelho, R, Y Lengwiler, K Rishabh and R Vrbaski (2026): "Strengthening the going-concern role of AT1: options and trade-offs", *FSI Briefs*, no 32.

Crisanto, J C, J Ehrentraud, J Prenio and J Yong (2025): "Smart supervision: sound capacity development approaches for tech-savvy supervisors", *FSI Insights on policy implementation*, no 70.

Di Stefano, A, Y Lengwiler and K Rishabh (2026): "The credibility of bail-in", *BIS Working Papers*, no 1356.

Eisenbach, T, D Lucca and R Townsend (2022): "Resource allocation in bank supervision: trade-offs and outcomes", *Journal of Finance*, vol 77, no 3.

Prenio, J (2024): "Peering through the hype – assessing supotech tools' transition from experimentation to supervision", *FSI Insights on policy Implementation*, no 58.

Reserve Bank of India (RBI) (2015): *Master Circular – Basel III capital regulations*, no DBR.No.BP.BC.1/21.06.201/2015-16, 1 July.

Swiss Federal Council (2012): *Ordinance on capital adequacy and risk diversification for banks and securities dealers* (Capital Adequacy Ordinance), no 952.03, 1 June.

Swiss Financial Market Supervisory Authority (FINMA) (2023): "FINMA approves merger of UBS and Credit Suisse", press release, 19 March.