



DO NOT BLAME BASEL: THE OUTPUT FLOOR AND EUROPE'S BANKING COMPETITIVENESS DEBATE

Judith Arnal and Karel Lannoo

SUMMARY

The European Commission's 2026 banking competitiveness consultation has put the Output Floor back on the policy table. Industry has called for weakening the floor. The case against doing so rests on four observations.

First, the Output Floor responds to a documented diagnostic. Successive Basel Committee studies have shown that risk-weighted assets vary substantially across banks for reasons that go beyond genuine portfolio differences, undermining the comparability of reported capital ratios. The 72.5 % calibration places a lower bound on the divergence between modelled and standardised outputs without prohibiting internal models.

Second, the March 2026 proposal from the US doesn't provide a precedent for relaxing the floor. It's a structural redesign that removes most advanced approaches for credit and operational risk, narrows the mandatory scope to the largest US banks, and retains model-based measurement mainly for market risk. The absence of a general 72.5 % floor in the new US design is therefore not directly comparable to abolishing the EU floor.

Third, the cost-benefit case for weakening the floor is flimsy. The additional Tier 1 capital impact of full EU Basel III implementation is small at sector level, and there's no evidence that capital requirements have constrained euro area banks' lending capacity. Reopening the core floor would trigger more regulatory uncertainty and supervisory recalibration – the opposite of simplification.

Fourth, reopening the floor is a strategic error: it is the most visible element of the 2017 Basel III compromise, and diluting it unilaterally would forfeit the EU's standing to demand faithful implementation from others.

This CEPS In-Depth Analysis paper proposes a five-question regulatory decision test to discipline the case-by-case proposals that will reach ministers, supervisors and the European Parliament over the coming months: reject changes that undermine the backstop purpose; fix Pillar 2 overlap rather than weaken the floor; use targeted FRTB-style relief only for proven international distortions; complete Banking Union rather than relabel national fragmentation as a Basel burden; and streamline reporting where the cost is purely operational.



Judith Arnal is an Associate Senior Research Fellow in the Financial Markets and Institutions unit at CEPS. Karel Lannoo is CEPS' CEO.

CEPS In-depth Analysis papers offer a deeper and more comprehensive overview of a wide range of key policy questions facing Europe. Unless otherwise indicated, the views expressed are attributable only to the authors in a personal capacity and not to any institution with which they are associated.

© CEPS and ECMI 2026

1. THE RATIONALE BEHIND BASEL'S OUTPUT FLOOR

Bank capital requirements are built around a simple idea. A bank first needs to measure its exposure to a borrower, transaction and risk type. That exposure is then converted into risk-weighted assets (RWAs) by applying a risk weight. The higher the risk weight, the larger the RWA amount and, therefore, the more capital the bank must hold. The logic is intuitive – a well-collateralised residential mortgage is not the same risk as an unsecured loan.

In EU terminology, the relevant denominator is normally called the total risk exposure amount (TREA). Capital ratios are then expressed as capital divided by TREA. For example, the CET1 ratio (the highest-quality capital, essentially common shares and retained earnings) is CET1 capital as a percentage of TREA, the Tier 1 ratio (Tier 1 adds Additional Tier 1 instruments, such as certain perpetual loss-absorbing securities) is Tier 1 capital as a percentage of TREA, and the total capital ratio is own funds as a percentage of TREA.

For credit risk, [the Basel framework](#) has historically allowed two broad ways of calculating those risk-weighted amounts. The first is the *standardised approach*. Under this method, the regulation itself assigns risk weights to broad types of exposures. Depending on the exposure class, the risk weight may depend on features such as the type of borrower, an external credit rating, the loan-to-value ratio for real estate or other regulatory criteria. It's called 'standardised' because the method is largely prescribed by the rulebook and is thus easier to compare across banks. Basel describes it as an approach under which standardised risk weights are assigned to exposures.

The second method is the *internal ratings-based approach*, or IRB. Under IRB, a bank may use its own internal rating systems and risk estimates, but only with supervisory approval. In the foundation IRB approach for corporate and bank exposures, banks estimate the probability of default (PD) but use supervisory estimates for other parameters such as loss given default (LGD), exposure at default (EAD) and maturity. In the advanced IRB approach, banks may estimate PD, LGD and EAD themselves, again subject to regulatory constraints and supervisory approval. For retail exposures, Basel requires banks using IRB to provide their own estimates of PD, LGD and EAD. The 2017 Basel III finalisation also narrowed the scope of IRB itself: advanced IRB is no longer available for large corporates (annual revenue above EUR 500 million), banks and other financial institutions, and IRB has been removed altogether for equity exposures. The Output Floor and these scope restrictions are parallel responses to the same diagnostic.

The reason supervisors allowed IRB models reflects a balance between three competing objectives: risk sensitivity, comparability and simplicity. A sophisticated bank may be able to measure risk more precisely than a simple standardised rule. A genuinely low-risk

mortgage book or a portfolio of highly-rated corporate exposures may deserve a lower capital requirement than a much riskier portfolio. In this sense, internal models can improve risk sensitivity. They also create incentives for banks to develop and maintain sound internal risk measurement and control systems – a discipline effect that the standardised approach, by design, doesn't produce.

Comparability, in turn, cuts both ways: standardised RWAs are mechanically more transparent across institutions but uniform risk weights can be economically misleading when applied to banks with very different business models, whereas internal models can, in principle, provide a more meaningful basis to compare across different institutions, at the cost of model risk and dependence on consistent modelling practices.

Model supervision has been intensive, not nominal: reducing RWA variability was a founding rationale of the SSM. The [ECB's TRIM](#) (2016-21), its largest supervisory project, produced over 5 000 findings and a ~EUR 275 billion RWA increase (about 12 % on the models reviewed, ~70bp off CET1). Even after intensive supervision, material variability persisted and this is why an aggregate backstop is a necessary complement for model-by-model review.

Table 1. Differences between standardised and internal model approaches

Same underlying question	Standardised approach	IRB/internal model approach
Who defines the method?	Mainly the regulation.	The bank, according to regulatory parameters and subject to supervisory approval.
What drives the risk weight?	Standard regulatory categories and prescribed risk drivers.	Bank-estimated risk parameters such as PD, LGD and EAD.
Main advantages?	Simplicity and mechanical transparency.	Risk sensitivity; incentives for sound internal risk management.
Main weaknesses?	Less tailored to actual portfolios; uniform treatment can mask differences between structurally heterogeneous banks.	Model risk; comparability depends on the consistency of modelling practices.
Prudential concern	May mechanically overstate or understate some risks.	Need to prove sound data bases and modelling abilities, checked by supervisors. Stringent controls (including MOCs).

Source: Authors' own elaboration.

In a [2013 Basel Committee study](#) of credit-risk RWAs in the banking book, the Committee found that risk weights varied significantly across banks. The top-down analysis covered supervisory data from more than 100 major banks, and a complementary Hypothetical Portfolio Exercise (HPE) was conducted on common wholesale exposures (sovereign,

bank and corporate) across 32 large international banking groups from 13 jurisdictions. Some of the observed variation reflected real differences in asset mix and risk appetite, which is legitimate.

But the Committee also found that the remaining variation was driven by differences in bank and supervisory practices. Translated into capital-ratio terms, the HPE found that 22 of the 32 participating banks would be within one percentage point of a 10 % benchmark capital ratio, while outlier banks could move by as much as two percentage points in either direction. That's a large effect – two banks could appear materially differently capitalised even where part of the difference came from modelling choices rather than from genuinely different risk.

[A later Basel Committee study](#) found similar concerns in retail and SME portfolios. It reported that average risk weights for residential mortgage portfolios ranged from 5 % to 80 % and from 11 % to 82 % for qualifying revolving retail exposures. Again, not all variation was illegitimate; real portfolios differ. But the scale of the dispersion showed why supervisors were concerned that model-based RWAs could become too difficult to compare.

This is the level-playing-field problem. If one bank uses the standardised approach and another uses IRB, the IRB bank may obtain a much lower RWA denominator for broadly similar exposures. Since the capital ratio is capital divided by RWA/TREA, the IRB bank can appear better capitalised even if its actual loss-absorbing capital hasn't changed. The issue isn't simply 'big banks versus small banks'; it's whether reported capital ratios mean the same thing across institutions, business models and jurisdictions.

The Output Floor was Basel's answer to that problem. It doesn't prohibit internal models. Instead, it places a lower bound on their regulatory output. Under the final Basel III framework, a bank's RWAs must be the higher of: first, the RWAs calculated using the approaches the bank has approval to use, including internal models where permitted; and second, 72.5 % of the RWAs calculated using only the standardised approaches. The floor binds on total RWAs, not separately by risk type.

Basel's own numerical example, as shown in Table 2, is useful. Here, a bank's pre-floor RWAs are 76, while the standardised calculation gives RWAs of 140. The floor is therefore $72.5\% \times 140 = 101.5$. Since 101.5 is higher than 76, the bank must use 101.5 for its capital requirements. The bank still benefits from internal modelling compared with the full standardised amount of 140, but the benefit is capped.

In plain terms, the bank must hold capital against RWAs of 101.5 instead of 76, so the same capital ratio now requires more capital than under unconstrained internal models, though still less than the 140 of a fully standardised calculation.

Table 2. Numerical example on the impact of the Output Floor.

Basel example	Pre-floor RWA	Standardised RWA	72.5% of standardised RWA
Credit risk	62	124	89.9
Market risk	2	4	2.9
Operational risk	12	12	8.7
Total	76	140	101.5

Source: Authors' own elaboration.

The economic effect is clear. Without the Output Floor, the bank would calculate capital requirements on RWAs of 76. With the Output Floor, it calculates them on RWAs of 101.5. If the relevant total capital requirement were 8 %, for example, the pre-floor requirement would be $8\% \times 76 = 6.08$, while the floored requirement would be $8\% \times 101.5 = 8.12$. The bank doesn't have to move all the way to the full standardised requirement of 11.2, but it can no longer use internal models to reduce the denominator below 72.5 % of the standardised benchmark.

That is why the 72.5 % figure must not be misunderstood. It is *not* a capital ratio. It does not mean that banks must hold capital equal to 72.5 % of their assets. It is a percentage applied to the standardised RWA/TREA benchmark. The ordinary capital ratios and buffers are then applied to the final floored denominator.

2. TRANSPOSING BASEL III INTO EU LEGISLATION

The EU implemented this Basel rule through the final Basel III banking package, particularly CRR III, which amended the Capital Requirements Regulation. In EU legal language, Article 92 CRR now expresses the calculation as:

$$TREA = \max \{ U-TREA ; x * S-TREA \}$$

where U-TREA is the unfloored total risk exposure amount, S-TREA is the standardised total risk exposure amount, and x is 72.5 % once the rule is fully phased in.

Substantively, the EU follows the core Basel design – the Output Floor is an aggregate backstop based on the standardised approaches, calibrated at 72.5 %. The [EBA's 2019 policy advice](#) recommended implementing the Output Floor in the EU to comply with the Basel agreement and calibrating it at 72.5 % of total RWA computed under the standardised approaches, precisely to create a credible backstop to internal models.

There are, however, EU-specific features.

First, the timing differs. Following the Covid-19-related deferral, the Basel Committee's oversight body moved the implementation date for the final Basel III standards to 1 January 2023 and extended the transitional arrangements for the Output Floor to 1 January 2028. The EU implemented the final Basel III package through CRR III, which applies mainly from 2025, and phases in the Output Floor with a two-year delay: 50 % in 2025, 55 % in 2026, 60 % in 2027, 65 % in 2028, 70 % in 2029 and, thereafter, the fully loaded 72.5 % calibration. The EU transition also includes a small number of portfolio-specific temporary treatments. For example, Article 465 CRR allows certain unrated corporate exposures to receive a 65 % risk weight until 31 December 2032, provided that the bank's estimated probability of default is no greater than 0.5 %. It also allows, subject to conditions and Member States' discretion, temporary lower risk weights for certain residential mortgage exposures.

There are also more technical temporary measures for counterparty credit risk and securitisations. The practical point is simple, the EU has two layers of transition. The headline Output Floor percentage reaches the fully loaded Basel calibration from 2030 but some targeted portfolio-specific treatments may continue to impact the calculation until 2032. Article 465 CRR also contemplates, based on an EBA evaluation of each measure, a possible extension of these arrangements for up to four additional years, which means that some portfolio-specific deviations from the fully loaded Basel calibration could, in principle, persist until end-2036.

Second, the EU's implementation has a specific feature on the level of application. In the EU, prudential capital requirements aren't applied only at the top of a banking group.

They also apply, as a starting point, with individual institutions. CRR III follows that logic for the Output Floor: the floor is relevant both at consolidated level and at individual level, so subsidiaries may also have to calculate their TREA by applying the floor. This is an important EU-specific feature because it means that the Output Floor isn't merely a group-level backstop. The rule isn't absolute, however. CRR Article 92 allows a Member State to decide that certain institutions belonging to a group with a parent institution in the same Member State may use the *unfloored* TREA at individual level, provided that the parent institution – or, in the case of certain cooperative or affiliated structures, the central body together with its affiliated institutions – applies the Output Floor on a consolidated basis. In practical terms, this creates a limited domestic derogation from solo-level application. It doesn't remove the Output Floor at consolidated level and it doesn't create a general cross-border waiver.

3. WHY THE OUTPUT FLOOR IS BACK ON THE POLICY AGENDA

The Output Floor is back for two related but different reasons. The first is internal to the EU – the Commission's 2026 banking competitiveness consultation explicitly asks whether elements of the EU's implementation of the Output Floor should be revised. The second is international: the US has proposed a different route for completing its Basel III implementation, which creates level-playing-field concerns but hasn't led to the same type of Output Floor debate as in the EU.

3.1. THE EU COMPETITIVENESS DEBATE

The clearest institutional reason why the Output Floor is back on the EU agenda is the Commission's 2026 [targeted consultation](#) on the competitiveness of the EU banking sector. The consultation is not a legislative proposal and doesn't prejudge a final Commission position, but it is politically important because it frames the issues that may feed into the Commission's 2026 report on banking competitiveness.

The consultation didn't emerge in a vacuum. In early 2026, the banking industry was publicly pressing the Commission to simplify the EU regulatory framework and reduce what it described as cumulative capital and supervisory burdens. That's why the European Commission made the decision to anticipate a CRR3 mandate to assess the competitiveness of the EU banking sector in 2028. The Output Floor entered the debate because the Commission's consultation asks about it directly. The document recalls that the EU introduced the Output Floor as part of the banking package applying from January 2025, that the measure aims to limit unwarranted variability in own-funds requirements produced by internal models, and that the fully loaded floor is calibrated at 72.5 % of the own-funds requirements that would apply under standardised approaches. It also notes one of the key EU-specific features: while Basel envisages application at the highest level of consolidation, the EU applies the floor both at consolidated level and at the individual subsidiaries level.

The Commission then asks stakeholders whether the current CRR III rules strike the right balance, whether some or all transitional derogations linked to the Output Floor should be prolonged or made permanent, whether the Output Floor should apply only at consolidated level, and whether the 72.5 % calibration should be increased, made more risk-sensitive or reduced. That is the strongest public basis for saying that the Output Floor has been put back on the EU policy table.

Industry responses confirm that this isn't a theoretical question. The European Banking Federation (EBF) [argues](#) that the current framework doesn't strike the right balance between risk-weighted requirements, the Output Floor, the leverage ratio and Pillar 2. It supports prolonging the Output Floor transitional derogations, applying the floor only at

consolidated level, making the calibration more risk-sensitive and, to a lesser extent, reducing the 72.5 % calibration. The EBF also supports introducing these transitional approaches' changes into the Standard Approach, given that these are issues that go beyond the output floor and that affect all banks. AFME similarly [asks](#) the EU to bring forward the review of Output Floor transitional arrangements for unrated corporates, low-risk mortgages and SA-CCR, to embed them permanently into a more risk-based standardised approach.

There is, however, a strong supervisory counter-position. The ECB's supervisory leadership has warned governments not to reduce bank capital requirements simply in the name of competitiveness, arguing that solid capital and liquidity positions remain necessary in a risky environment. The policy debate is thus not just about whether banks should have more balance-sheet capacity. It's about where to draw the line between legitimate simplification and prudential weakening. Part of the dynamic can be named plainly – the Output Floor and intrusive model supervision removed a margin of regulatory arbitrage that some banks had treated as structural. Much of the complaint is less about lending capacity than about losing that margin.

3.2. THE US ANGLE – NOT A REOPENING OF THE OUTPUT FLOOR BUT A DIFFERENT BASEL III IMPLEMENTATION PATH

The [US angle](#) should be framed with some care. The March 2026 US proposal should not be described simply as an attempt to 'abolish the Output Floor'. It's a broader redesign of the capital framework for the largest and most internationally active US banks. The proposal would apply mainly to Category I and II banking organisations, introduce a single expanded risk-based approach, and remove the existing advanced approaches from the regulatory capital framework, except for advanced market-risk models under the Basel FRTB framework.

This structural choice is important. The traditional rationale for an Output Floor is to prevent banks from gaming the rules, using internal models to generate risk-weighted assets that fall materially below those produced by the standardised approach. If most internal-model approaches are removed for credit risk and operational risk, the case for an aggregate Output Floor becomes less direct in those areas. The US proposal is therefore not a case of banks retaining internal models while the backstop is removed. Rather, it moves the framework towards a more standardised calculation of risk-weighted assets, while introducing certain recalibrations within that framework.

Consequently, the Output Floor would in practice become largely irrelevant under the proposed US structure. Since risk-weighted assets would mostly be calculated under a

standardised or standardised-like approach, remaining generally above the 72.5 % benchmark that the Output Floor is designed to enforce.

That said, the proposal still raises prudential concerns. The most important area is market risk, where the US proposal retains the Basel logic of model-based measurement. From an Output Floor perspective, this is also the area where the proposal is most problematic. The assessment notes that, in some cases, the standardised measure could operate as a cap rather than a floor for model-derived market-risk capital requirements. In practical terms, a bank could be allowed, subject to supervisory approval, to use the standardised measure where it is lower than the model-based requirement. This reverses a standardised benchmark's usual prudential function, which is intended to act as a backstop rather than a ceiling.

The implications are particularly relevant for banks with large trading books. For a diversified banking group, the absence of a general standardised floor may have limited effects. For a bank more heavily concentrated in trading activity, however, it could create a material capital advantage. The concern is especially acute for derivatives-related exposures, given that the proposal may also reduce requirements for counterparty credit risk and CVA risk in certain areas.

Three further points of asymmetry should be kept in mind when comparing the two regimes.

First, the scope of application differs materially: the US proposal applies only to Category I and II banking organisations, whereas CRR III applies to all EU credit institutions. The two regimes are therefore not addressing comparable populations, and any conclusion about relative regulatory tightness that doesn't control for this scope difference will systematically mismeasure the actual divergence.

Second, the pre-reform US baseline for large banks was not equivalent to the EU baseline. [ECB counterfactual analysis](#) estimates that, under current US rules, the largest EU banks would face stricter capital requirements than under CRR III – driven mainly by the stricter US calibration of the GSIB surcharge and by legal limitations on the use of internal models – while mid-sized EU banks would face somewhat less stringent requirements. Anyhow, this estimate should be read with caution: the published analysis reveals few details on the underlying data and methodology, so it is best read as an indication of the gap's direction rather than a precise measure of it. Part of the US proposal can therefore be read as a partial recalibration aimed at narrowing that pre-existing differential for the largest banks, rather than as a movement below the Basel benchmark from a neutral baseline.

Third, the regimes rest on different accounting bases (US GAAP versus IFRS) which diverge on derivative and repo netting; comparing reported ratios at face value, without adjusting for this, mismeasures the true divergence.

Table 3 compares the EU and US approaches. It should, however, be read with an important caveat. The US proposal isn't simply a lighter version of the EU regime. It changes the structure of the framework by removing most advanced approaches for credit and operational risk, narrowing the mandatory scope of Basel III implementation, and retaining model-based measurement mainly for market risk. For this reason, the absence of a general 72.5 % Output Floor in the US proposal is not directly comparable to abolishing the EU Output Floor.

The more meaningful comparison is in the areas where internal models remain relevant, particularly market risk and trading-related exposures. The table therefore distinguishes between two issues: broader Basel III divergence, which matters for competitiveness, and the more specific prudential question of whether the standardised benchmark continues to operate as a genuine backstop where internal models are still used.

Table 3. An EU-US Output Floor comparison

Issue	Current / proposed EU position	US March 2026 proposal	Why it matters
Credit-risk internal models	IRB models remain possible in the EU, subject to restrictions, supervisory approval and the Output Floor.	Advanced approaches for credit risk would be removed from the regulatory capital framework.	The Output Floor remains more directly relevant in the EU for credit risk. The US is not keeping IRB credit models and removing the floor.
Operational risk	Advanced operational-risk models are removed and replaced by a more standardised framework (Internal Loss Multiplier (ILM = 1, i.e. loss history excluded).	Advanced operational-risk models are removed and replaced by a more standardised framework. (Internal Loss Multiplier (ILM = 1, i.e. loss history excluded).	Again, the classic 'internal model output too low' problem is less central in the US proposal for operational risk.
General Output Floor	The EU has implemented the Basel Output Floor, reaching 72.5 % when fully phased in.	No equivalent overall 72.5 % Basel-style Output Floor in the new single-stack design.	This looks divergent but it's not fully comparable because the US proposal removes most advanced approaches outside market risk.
Market risk / FRTB design	Binding FRTB own-funds requirements have been postponed until 1 January 2027. During 2026, pre-FRTB methodologies continue for market-risk own-funds requirements, while FRTB-SA remains relevant for Output Floor	The proposal introduces a revised market risk framework and retains internal models for market risk, subject to supervisory approval. However, in some cases the standardised measure could operate as a cap	Market risk is the most meaningful comparison point because both systems still allow model-based measurement. The key prudential concern is not merely the absence of a general Output Floor but whether the standardised benchmark

Issue	Current / proposed EU position	US March 2026 proposal	Why it matters
	purposes. The Commission has also proposed targeted temporary amendments to mitigate the FRTB capital impact from 2027.	rather than a floor for model-based market risk requirements.	continues to operate as a genuine backstop where internal models remain in use.
Market-risk scope and trading-book impact	The EU threshold referred to in the assessment is much lower, at EUR 50 million.	The US proposal would exempt banks with less than USD 5 billion in trading activity from the complex market-risk requirements. It may also reduce requirements for counterparty credit risk and CVA risk in certain areas.	The US framework would be narrower in scope and could be less demanding for some firms, especially banks with large or concentrated trading books and derivative-related exposures. This is where the level-playing-field concern is strongest.
Policy conclusion	The EU has already used targeted delay and temporary relief for FRTB to preserve the level playing field with the US and while undergoing more structural reforms.	The US is taking a narrower and structurally different route.	The US proposal is not a clean argument for weakening the EU Output Floor generally; it's an argument for looking carefully at market-risk level playing field issues.

Source: Authors' own elaboration.

A final methodological caveat concerns the impact estimates that have circulated in the policy debate around the US proposal. The headline capital-savings figures most often cited do not derive primarily from the standardised-only redesign of credit and operational risk: a substantial share originates from the recalibration of the GSIB surcharge and from changes to the stress-test-based buffer framework, both of which sit outside the Output Floor debate. Any quantitative comparison of the EU and US regimes thus requires the careful disentangling of which component of capital requirements is driving the result, and a recognition that the figures most often cited in industry materials capture a broader policy package than the one directly relevant to the Output Floor question.

4. DEFENDING THE FLOOR – AND SIMPLIFYING AROUND IT

The Commission consultation invites stakeholders to assess whether the Output Floor should be prolonged, scoped down or recalibrated. The intuitive case for doing so doesn't survive close examination, and the EU has stronger alternatives available. This section sets out, first, why weakening the floor is the wrong response; second, what should be done instead; and third, a decision test for handling individual proposals as they arrive.

4.1. WHY A WEAKENED FLOOR IS THE WRONG RESPONSE

The most basic reason is *conceptual*. The Output Floor is not regulatory paperwork. It is a prudential backstop. Its job is to stop internal models from pushing capital requirements too far below a common standardised reference, so that two banks reporting similar capital ratios remain meaningfully comparable. Removing or lowering it wouldn't simplify a form or eliminate a duplicated process. It would weaken the rule that keeps internal model capital ratios comparable and credible. If policymakers want simplification, they should target duplication, reporting friction and overlapping capital layers. They should not weaken the rule that prevents model outputs from becoming too favourable.

The 2008-09 financial crisis showed how far reported capital ratios could diverge from real risk – Dexia passed the July 2011 EU stress test only months before being rescued. The [fiscal cost of the bank rescues](#) reached roughly 10-13 % of the EU's GDP, with long-lasting effects on countries' public finances. Limiting the understatement of risk through internal models is precisely what the Output Floor was designed to do.

A second reason is that the EU also isn't starting from a pristine Basel implementation. CRR III already includes a gradual phase-in and several portfolio-specific transitional treatments – the 65 % risk weight for certain unrated corporates until the end of 2032, the discretionary mortgage relief and technical transitional measures for counterparty credit risk and securitisations. The debate is not whether to introduce flexibility but whether to entrench it. Making temporary measures permanent without any evidence of a Basel-compatible miscalibration wouldn't clarify the current rules. It would turn an adjustment period into a structural lowering of the backstop by another name.

The quantitative case for doing so is also weak: the EBA's Basel III monitoring work finds that the additional Tier 1 capital shortfall for full EU-specific Basel III implementation is small at sector level, even though the Output Floor and operational risk are the main drivers of the increase in minimum required capital. Individual model-using banks experience the floor more strongly, but the aggregate claim that the EU banking system cannot finance the economy unless the floor is diluted isn't supported by the data.

A third reason is that reopening the core floor would itself create the kind of regulatory cost that the simplification agenda is supposed to reduce. Banks have already invested in CRR III systems, reporting templates, model governance and implementation planning. A broad reopening would invite another lobbying cycle, another phase of technical uncertainty and another round of supervisory recalibration. That's the opposite of simplification. A rule can be demanding and still be clear. A permanently unstable rule is expensive even when it is less stringent.

The fourth and final reason is international. The Output Floor is the most visible single element of the 2017 Basel III compromise and the EU was among those that pushed hardest for a credible backstop to internal models. Unilaterally reopening it would undermine the authority of the Basel Committee itself and leave the EU exposed as an outlier on global capital standards, the very position it has spent years urging other jurisdictions not to take.

The timing compounds the risk: it would come just as international cooperation on banking standards is already fraying and large banks have been retrenching towards their home markets. Faithful implementation is, in this context, the least the EU can offer if it expects reciprocal discipline from others; abandoning the floor the moment a peer diverges forfeits exactly that leverage.

4.2. WHAT TO DO INSTEAD

If the diagnosis is that the EU banking sector needs more competitive conditions, the right response is to defend the floor and simplify the framework around it. This position is pro-competitiveness without being deregulatory. Two considerations sit underneath it. First, as the Eurosystem response to the Commission consultation makes clear, competitiveness comes from harmonisation, integration and scale, and not from deregulation – there's no evidence that capital requirements have hampered euro area banks' lending capacity. Second, the EU has a strong interest in other jurisdictions faithfully implementing Basel III. If it responds to softer foreign implementation by weakening one of Basel III's most visible safeguards, it loses standing in the Basel Committee when arguing for global consistency. The Output Floor is an internationally agreed minimum, not unilateral EU gold-plating: unpicking it would weaken the Basel framework itself and the EU's authority to hold other jurisdictions to account over their own implementation. The credible position is to maintain the 72.5 % backstop, monitor divergence and act surgically where a genuine distortion is identified.

This translates into four practical priorities.

The first is the interaction between the Output Floor and Pillar 2. Where model risk is already constrained by the floor, supervisors should not replicate the same concern

automatically through Pillar 2 add-ons. Pillar 2 must remain available for bank-specific risks not captured by Pillar 1, but its methodology should be transparent enough to show when it is addressing something genuinely additional. Addressing this overlap is a better response to industry complaints than weakening the floor, because it targets duplication without removing the common backstop.

The second is to treat Banking Union as the real competitiveness reform. If the underlying problem is that capital and liquidity cannot move freely inside cross-border groups, the answer isn't a lower floor. The answer is to make the euro area work more like a single jurisdiction, with adequate safeguards for subsidiaries and host authorities. Progress on EDIS, crisis management, national insolvency barriers, supervisory harmonisation and group waivers would do far more for bank scale and efficiency than any reduction in floor calibration.

Responsibility here is shared across two layers. The structural enablers (EDIS, insolvency harmonisation and the statutory conditions for cross-border waivers) sit with the co-legislators and Member States; the ECB cannot deliver them. But within the existing framework the ECB, as supervisor, retains its own margin. The relevant instrument is the cross-border capital and liquidity waiver, which allows a group to be managed as a single entity rather than as a set of separately capitalised subsidiaries. A waiver granted on its own would strip the host subsidiary of its standalone buffer – which is precisely why host authorities resist it and why such waivers are rarely used across borders. The constructive route is thus not a looser waiver policy but a structured one: waivers paired with enforceable, collateralised intragroup guarantees that pre-position protection for the host, substituting a ring-fenced claim for the solo buffer the waiver removes. On this basis, group-level treatment can be efficient for the bank and acceptable to the host. This is also where the EU-specific issue of solo-level application of the Output Floor should be reconsidered – not as a stand-alone capital relief measure but as part of the broader integration agenda.

The third is to use targeted instruments for trading-book divergence. The Commission's April 2026 draft delegated act on FRTB is the right type of instrument: narrow, time-limited and linked to an identified international level-playing-field issue in trading activities. As argued in Section 3.2, the prudentially significant US-EU divergence sits in market risk, not in credit or operational risk, and that is where the case for targeted, reversible relief is strongest. The approach shouldn't be generalised into a broad capital holiday. It should remain conditional on the evolution of US and UK implementation.

The fourth is to simplify reporting and disclosure. The Output Floor requires a standardised benchmark calculation even for banks that use internal models, so it inevitably creates operational work. Some of that work is inherent to the backstop. Some

may be avoidable if templates, definitions and supervisory data requests are better integrated. Reporting simplification should be pursued aggressively, because it reduces costs without reducing resilience.

4.3. A REGULATORY DECISION TEST

The four priorities above describe the policy line. But policy isn't made on a policy line. It's made one proposal at a time. Over the coming months, every change presented to ministers, supervisors and the European Parliament will carry its own competitiveness or simplification justification: a Pillar 2 recalibration here, an extension of a transitional measure there, a delegated act, a reporting waiver, a clarification on subsidiary application. Each will look defensible on its own. The cumulative effect of approving them one by one, without a common test, would lead to the slow erosion of the floor that this CEPS In-Depth Analysis paper has argued against.

A regulatory decision test is thus necessary. It distils the above into a short sequence of questions to be applied case by case. The questions must be asked in the order shown below in Table 4. A proposal that fails on one question doesn't move to the next. Any proposal that touches the Output Floor or the framework around it should pass all five before being adopted.

Table 4. A regulatory decision test for the Output Floor

Test	If yes...	If no...
Does the proposal undermine the floor's backstop purpose?	Reject it. Preserve the 72.5 % floor.	Proceed to the next test.
Is there double counting between the floor and Pillar 2 or buffers?	Fix the overlap by clarifying Pillar 2 methodology.	Do not use double counting as an argument against the floor.
Is there a proven international distortion in a specific market?	Use targeted, time-limited relief, FRTB-style.	Avoid broad relief.
Is the problem caused by national fragmentation?	Fix Banking Union barriers and improve capital and liquidity mobility with safeguards.	Do not relabel it as a Basel burden.
Is the cost mainly an operational reporting burden?	Streamline templates and data requests.	Keep the prudential rule intact.

Source: Authors' own elaboration.

5. CONCLUSIONS: KEEP THE BACKSTOP – BUT SIMPLIFY AROUND IT

The Output Floor is not a marginal technical rule. It is one of the mechanisms through which the final Basel III framework restores credibility to risk-weighted capital ratios. Its purpose is simple: banks may use internal models, but those models should not be able to push regulatory capital requirements too far below a common standardised benchmark. That's why the 72.5 % floor should be understood as a prudential backstop – not an administrative burden.

The EU's implementation follows the core Basel logic. It applies an aggregate floor, calibrated at 72.5 % when fully phased in, while adding EU-specific features on timing, transitional treatments and the level of application. Those choices can legitimately be reviewed. The EU framework should be efficient, proportionate and internationally competitive. But the distinction between reviewing implementation details and weakening the core backstop is essential. A temporary transitional measure is not the same as a permanent recalibration. A targeted FRTB adjustment is not the same as lowering the Output Floor across the banking book. And avoiding double counting in Pillar 2 is not the same as removing the common floor.

All that is why the 2026 debate should be framed carefully. The Commission's competitiveness consultation has clearly put the Output Floor back on the policy table and industry responses show that some stakeholders want meaningful changes. At the same time, the US example doesn't provide a straightforward precedent for weakening the EU floor. The March 2026 US proposal changes the architecture of the capital framework by removing the most advanced approaches for credit and operational risk, narrowing the mandatory scope of application and retaining internal models, mainly in market risk.

The wrong response would be to turn the competitiveness debate into a dilution of Basel III. Lowering the 72.5 % calibration, making transitional relief permanent without strong evidence or removing the floor's backstop function would revive the very problem Basel III was designed to address: capital ratios that depend too heavily on bank-specific modelling choices and are thus harder to compare across institutions and jurisdictions. It would also weaken the EU's credibility in pressing other jurisdictions to faithfully implement Basel III.

The right response is more disciplined. The EU should keep the fully phased-in 72.5 % Output Floor, preserve the backstop's purpose and simplify the framework around it. That means addressing genuine double counting between the floor and Pillar 2, improving transparency in supervisory add-ons, simplifying reporting and disclosure, and using targeted, temporary measures where there's a proven international distortion in a

specific market segment. It also means treating Banking Union as the real competitiveness reform: if capital and liquidity are trapped by national fragmentation, the answer isn't a weaker floor but a more integrated European banking market.

The policy test should be straightforward. Any proposed change should be rejected if it undermines the floor's backstop purpose. If the problem is double counting, fix Pillar 2. If the problem is international divergence in a specific market, use targeted and temporary relief. If the problem is national fragmentation, complete the Banking Union. If the problem is operational burden, streamline reporting. This approach is pro-competitiveness but without being deregulatory.

The central message is simple – Europe will not become more competitive by making bank capital ratios less credible. Sustainable competitiveness comes from strong, comparable and trusted banks operating in a genuinely integrated market.

The Output Floor is part of that credibility. It should be defended, not dismantled.



**CEPS
PLACE DU CONGRES 1
B-1000 BRUSSELS**
