On Capital as Constraint, Not Description

Regulatory capital cannot simultaneously function as a forward constraint and a backward mirror. The moment realised dispersion becomes the implicit anchor of calibration, the architecture shifts from constraining uncertainty to describing it. Description may be informative, but it is not binding. A prudential system that derives its structural confidence from observed outcomes risks embedding procyclicality into its own design, because outcomes are state-contingent while constraint must be state-robust.

Capital therefore belongs to the counterfactual domain. It exists to absorb uncertainty not yet realised, including regime transitions that have no historical frequency representation. If calibration logic reacts to realised variability rather than exposure structure, the framework begins to internalise the cycle instead of buffering it. This is not a technical nuance; it is a structural inversion.

The credibility of capital depends less on granularity than on interpretive coherence. Where economically similar exposures produce materially divergent capital requirements due to methodological latitude, the dispersion ceases to communicate risk and begins to communicate calibration freedom. At that point, the burden of interpretation migrates from architecture to explanation, and supervisory confidence becomes conditional on narrative justification. Capital that requires explanation to be trusted has already lost part of its binding force.

Risk sensitivity is therefore subordinate to comparability. Sensitivity is meaningful only within a bounded measurement architecture. Where model discretion introduces output divergence that cannot be unambiguously attributed to exposure heterogeneity, variability becomes noise. The question is not whether dispersion exists, but whether it is risk-justified or design-induced. A framework that tolerates design-induced dispersion in the name of flexibility ultimately weakens its own constraint function.

Internal coherence is equally decisive. Expected credit loss, risk weights, stress testing and internal capital assessment cannot operate on implicitly divergent probability structures without forcing reconciliation into governance space. When different supervisory engines embed different state assumptions, overlays proliferate not because institutions are opportunistic, but because architecture is inconsistent. Reconciliation then becomes operational rather than structural. Month-end adjustments substitute for probabilistic unity. In such a system, modelling complexity increases while constraint reliability decreases.

Representativeness is often invoked as a safeguard against miscalibration. Yet representativeness without consequence remains descriptive. If deviation from benchmark does not alter calibration logic or supervisory posture, it expands reporting volume without strengthening prudence. A prudential requirement that cannot change behaviour functions as disclosure, not constraint. The effectiveness of representativeness therefore depends not on data granularity but on consequence definition.

Proportionality, similarly, is frequently framed as an objective in itself. It is not. Simplification enhances prudence only where incremental modelling complexity increases model risk more than informational content. Where simplification reduces dispersion while preserving conservative constraint, it strengthens credibility. Where it reduces constraint while expanding optionality, it weakens the architecture regardless of intent. The legitimacy of proportionality depends entirely on whether it reduces unwarranted variability without eroding forward robustness.

Variability in capital outcomes must be examined through this lens. Dispersion grounded in exposure heterogeneity is inevitable and legitimate. Dispersion arising from heterogeneous facility definitions, inconsistent observation units, discretionary MoC layering or opportunistic fallback usage is architectural incoherence. Supervisory tolerance for incoherence erodes cross-sectional trust and introduces interpretive latitude where constraint should be uniform.

Governance credibility is tested not when losses materialise, but when calibration remains binding in the absence of stress. A framework that intervenes only after dispersion becomes observable has shifted from ex-ante constraint to ex-post correction. Prudence requires calibration discipline embedded in design, conservative hierarchy embedded in fallback logic, and revocation capacity embedded in governance architecture. Constraint must precede variance; otherwise variance governs constraint.

Capital ratios function as public trust instruments. Their systemic value lies in their comparability and forward robustness. Where methodological heterogeneity produces unexplained divergence across institutions or jurisdictions, the signalling function weakens and supervisory authority becomes interpretive. The erosion is gradual but cumulative.

The supervisory task is therefore singular: preserve capital as a forward-consistent, comparable and binding constraint across states of the world. Simplification, modelling sophistication, representativeness and fallback calibration are subordinate design choices. They acquire prudential legitimacy only insofar as they reinforce that constraint. Where they expand discretion without enhancing coherence, they undermine the very objective they claim to serve.