ECB FORUM ON CENTRAL BANKING

29 June-1 July 2026

Shaping Europe's future: innovation, growth and stability

The stringency and complexity of post crisis bank regulation

Mariassunta Giannetti

EUROPEAN CENTRAL BANK

EUROSYSTEM

The Stringency and Complexity of Post Crisis Bank Regulation

By Cameron Ellis, Mariassunta Giannetti, and Chotibhak Jotikasthira[1]

Abstract

We propose a distinction between the stringency and the complexity of bank regulation — the bindingness of a regulatory action and the operational burden it imposes on regulated institutions — and construct cumulative country–quarter indices for 21 advanced economies between 1995 and 2025 by feeding every iMaPP and MaPPED event into pairwise large-language-model tournaments. Using an imputation-based generalised synthetic control, we trace the response of systemic risk, stock-market valuations, and private credit to events partitioned by direction and by their stringency and complexity magnitudes. Tightening events that raise both stringency and complexity reduce SRISK by approximately 25 log points over two years; tightenings that raise stringency without complexity produce only a transient reduction, consistent with regulatory arbitrage when the rulebook is operationally simple. Aggregate equity values and private credit do not contract after tightenings. Complexity is harmful only when it is not paired with stringency: bank stocks underperform the broad market by approximately 17 log points in the low-stringency, high-complexity cell. The post-crisis rulebook's complexity therefore is costly if it is associated with lax regulations but appears to make stringent regulation harder to circumvent.

1 Introduction

The post-2008 expansion of bank regulation was unprecedented in scale and scope. Basel III moved prudential regulation beyond a primarily capital-based framework, adding liquidity, leverage, stress-testing, and resolution-planning requirements. Dodd-Frank in the United States, the Capital Requirements Regulation in the European Union, and parallel reforms in the United Kingdom, Switzerland, Japan, and other advanced economies then layered jurisdiction-specific rules on top of the Basel framework. The result is a post-crisis rulebook that is both more stringent and more complex than the regime it replaced.

[1] Cameron Ellis is with the Tippie College of Business, University of Iowa. Email: cameron-ellis@uiowa.edu. Mariassunta Giannetti is with the Stockholm School of Economics, the Centre for Economic Policy Research (CEPR), the European Corporate Governance Institute (ECGI). and Swedish House of Finance. Email: Mariassunta.Giannetti@hhs.se Chotibhak Jotikasthira is with the Cox School of Business, Southern Methodist University. Email: cjotikasthira@smu.edu.

That rulebook is now being reopened across major advanced economies. The debate is not simply about whether banks should hold more or less capital. It is also about the architecture of regulation: whether multiple overlapping constraints are necessary to contain systemic risk, or whether some of the complexity now imposes costs without commensurate benefits. Recent policy choices place major jurisdictions at different points in this stringency-complexity space. The United States has moved toward capital relief, including a lighter treatment of the Basel III endgame and a reduction in the enhanced supplementary leverage ratio for global systemically important banks. The United Kingdom has adopted a broadly capital-neutral implementation of Basel 3.1 while reviewing the ring-fencing regime to reduce compliance costs. The European Union is implementing Capital Requirements Regulation III (CRR III) and Capital Requirements Directive VI (CRD VI), which raise binding requirements for many large banks through the output floor and revised risk weights, while also pursuing simplification of the existing rulebook. Switzerland has moved in the opposite direction from the United States, with proposals that would require UBS to hold substantially more CET1 capital after the acquisition of Credit Suisse.

These divergent reforms raise a common question: can the post-crisis rulebook be made less costly to operate without losing the systemic risk gains it was intended to produce? Put differently, does regulatory complexity make bank regulation more effective, or does it create distortions that weaken the allocation of credit and risk-bearing capacity? This is the central question of the paper.

To address this question, the paper studies the evolution of bank regulation from 1990 to the present, from the consolidation of the Basel I framework and the rise of risk-sensitive capital regulation to the post-2008 layering of capital, liquidity, leverage, stress-testing, and resolution-planning requirements. Over this period, prudential regulation has become both more stringent and more complex. Basel I established an international minimum capital standard; the 1996 Market Risk Amendment and Basel II deepened the reliance on risk weights, internal models, and supervisory review; and Basel III added liquidity and leverage constraints on top of higher and better-quality capital requirements. National reforms, including Dodd-Frank in the United States, the Capital Requirements Regulation in the European Union, and parallel reforms in the United Kingdom, Switzerland, Japan, and other advanced economies, then layered jurisdiction-specific rules on top of the Basel framework.

Basel III is the natural focal point for the current policy debate because it embodies the post-crisis turn toward overlapping constraints. It increased the level and quality of bank capital, but it also multiplied the number of requirements facing large banks. Yet the roots of this complexity predate Basel III. They can be traced to the development of the Basel framework in the 1990s, the 1996 Market Risk Amendment, and Basel II's reliance on internal models, supervisory review, and risk-sensitive capital requirements. Greenwood, Hanson, Stein, and Sunderam (2017) emphasize that post-crisis U.S. capital regulation imposes at least ten overlapping equity constraints on the largest banks, and that the interaction among binding

constraints can generate allocative distortions across activities and institutions. Importantly, their argument is not that regulation is too stringent. If anything, they suggest that a few additional percentage points on risk-based capital ratios for the largest banks would be socially beneficial. The issue, then, is not whether regulation should be strong, but whether strength requires the current degree of complexity.

The extent to which regulators could decrease the costs associated with regulatory complexity without weakening their effectiveness remains an open question. Theoretically, whether regulatory complexity is on net harmful or beneficial is ambiguous. On the one hand, complexity may be a by-product of political economy: in the aftermath of financial crises, politicians and regulators have strong incentives to be seen as acting decisively, and may layer poorly designed rules onto the existing framework to signal competence and responsiveness (Gratton, Guiso, Michelacci, and Morelli, 2021; Almasi, Dagher, and Prato, 2025), even when the new requirements add little to financial stability and mostly increase regulatory uncertainty (Davis, 2017). On the other hand, complexity may be a necessary response to an increasingly complex financial system. In a theoretical model, Asriyan, Foarta, and Vanasco (2023) show that when agents have imperfect information about a product, complex design may either separate types and improve outcomes or obfuscate quality and amplify adverse selection, depending on the underlying information structure. Applied to bank regulation, a layered rulebook could either capture economically distinct sources of risk that simpler rules cannot price and contain or pile up deadweight compliance costs without commensurate stability gains. Whether complex and stringent regulations are desirable or not, they would often be relaxed to favor growth during the next boom because of political opportunism (Dagher, 2018; Müller, 2023), with neglect for financial stability.

The empirical evidence so far has emphasised the cost side. Kalmenovitz, Lowry, and Volkova (2025) document that fragmentation across U.S. regulators raises compliance costs and depresses the activity of regulated entities for corporations. Colliard (2025) develops a quantitative measure of regulatory complexity and provides systematic evidence that the post-crisis rulebook has grown substantially more complex and operationally demanding for regulated financial institutions.

Very high regulatory costs may become unproductive if they incentivize intermediaries to devise new arrangements to avoid them. For instance, the growth of private equity-backed lending is a particularly visible illustration. Leveraged buyouts are increasingly financed by private credit funds whose ultimate investors include life insurers and pension funds; these vehicles fall outside the bank prudential perimeter, but remain interconnected with regulated banks through warehouse facilities, repo financing, and credit-line backstops (Acharya, Cetorelli, and Tuckman, 2024; Chernenko, Erel, and Prilmeier, 2022). Systemic risk thus appears to migrate, rather than disappear, when bank capital and liquidity rules tighten, and may render bank-level regulation progressively less effective at containing aggregate risk.

The same mechanism is visible in fixed-income markets. Post-crisis leverage ratios have reduced the capacity of bank-affiliated dealers to act as intermediaries in corporate bond markets, and a growing share of bond holdings has migrated to

open-end mutual funds. In Giannetti, Jotikasthira, Rapp, and Waibel (2024), we show that, after the introduction of the leverage ratio constraint on bank-affiliated dealers, bond mutual funds began to provide more liquidity in investment-grade corporate bonds, and their performance benefited. However, the liquidity and returns of investment-grade corporate bonds have become more exposed to aggregate outflows from bond funds with liquidity-supplying strategies. When these funds are unable to purchase bonds exposed to leverage constraints, market functioning deteriorates: their missing liquidity provision helps explain which investment-grade bonds experienced the most severe declines in liquidity and returns at the onset of the COVID-19 pandemic, before the Federal Reserve intervened to stabilize the bond market. Regulation has thus relocated fragility to a less-regulated segment rather than eliminating it.

Whether regulation complexity is an unnecessary cost or increases regulations' effectiveness – in the sense of tracking risks that a single ratio cannot summarise – remains an open question. Using bank regulations as a case in point, this paper explores whether an increase in the stringency of bank regulations is indeed associated with a reduction in systemic risk, or whether financial systems adapt in a way that systemic risk falls temporarily and then rebuilds as intermediaries find new ways to increase leverage, as the private credit example illustrates. We also ask whether the complexity of the regulations simply increases their costs, as some argue, or also their effectiveness, and whether it is a necessary cost associated with the financial system's increasing complexity.

These questions are challenging to apply to the data. Regulatory actions vary in scope, granularity, and the degree to which they bind regulated institutions; their effects on systemic risk must be assessed against changes in the very environment the rules attempt to govern. Existing aggregate measures collapse the multiple dimensions of regulation into a single scalar, count tools rather than weight them by intensity, and are silent on the distinction between stringency and complexity.

We aim to provide preliminary evidence on these questions. To achieve this, we construct two indices for the stringency and the complexity of bank regulations. We define stringency as the extent to which a regulatory action constrains the behaviour of a regulated institution — the breadth of institutions affected, the magnitude of the constraint imposed, and how binding the new requirement is in practice — and complexity as the operational and compliance burden the action imposes, whether through new tool types, conditional thresholds and carve-outs, coordination across subsidiaries or jurisdictions, or layering on top of existing rules. The two dimensions are orthogonal in principle: a loan-to-value cap is highly stringent from the borrower's standpoint but simple for the originating bank to implement, whereas the transposition of a Basel III Pillar-2 add-on with multiple conditional thresholds and institution-specific calibrations can be equally stringent yet substantially more complex. We score every dated event in the IMF's integrated Macroprudential Policy (iMaPP) database (Alam et al., 2019) and the Macroprudential Policies Evaluation Database (MaPPED), maintained by the European Central Bank along the stringency and complexity dimensions by feeding the event together with its narrative description into a pairwise large-language-model tournament, proposed by

Asirvatham, Mokshi, and Shleifer (2026). The model asks, for every pair of events of the same direction, which action had the larger impact on the stringency (respectively, complexity) of the regulatory regime. Stringency and complexity are evaluated in separate tournaments under the definitions above, and tightenings and loosenings are ranked in separate tournaments because cross-directional comparisons are not well-defined. Each event receives a z-score along each dimension, and the country–quarter index is the signed cumulative sum of those z-scores; the resulting indices for 21 advanced economies between 1995Q1 and 2025Q3 are event-dated and reversible, weighted by intensity rather than by the simple presence of a tool, and separate two attributes that the literature has typically collapsed into a single scalar.

We use these indices to partition every regulatory event into one of eight cells, defined by the direction of the action (tightening or loosening) and by whether the event-quarter stringency and complexity shocks are above or below the country-specific median on each dimension. This partition lets us read the separate effects of stringency and complexity, holding the direction of the action and the country fixed — an exercise that the single-scalar regulatory measures used elsewhere in the literature do not allow. We trace the response of systemic risk (SRISK), stock-market valuations, and private credit to GDP at horizons up to twelve quarters using an imputation-based generalised synthetic control in the spirit of Dube, Girardi, Jordà, and Taylor (2025), Bai (2009), and Xu (2017). An interactive fixed-effects model is fit on clean-control cells to absorb latent slow-moving country-specific exposures to common shocks; the no-treatment counterfactual is imputed for every country–quarter; and horizon local projections on the imputation residual recover the cell-specific impulse responses.

Three findings emerge. First, tightening events that raise both stringency and complexity reduce systemic risk by approximately 25 log points over the two years following the event, whereas tightenings that raise stringency without complexity produce only a transient reduction that fades within four quarters. This finding suggests that complexity is not unnecessary regulatory baggage: it appears to make stringent regulation harder to circumvent through regulatory arbitrage. Second, neither aggregate equity values nor private credit contract after tightenings, including those that combine high stringency with high complexity; whatever costs the post-crisis rulebook has imposed on regulated banks have not translated into a contraction of aggregate stock-market value or of credit to the private sector. Third, complexity is harmful only when it is not paired with stringency: bank stocks underperform the broad market in the cell that combines low stringency with high complexity, consistent with the theoretical mechanism in Asriyan, Foarta, and Vanasco (2023) under which complexity becomes a cost without a corresponding stability gain.

This paper also contributes to a large literature that evaluates the effects of bank regulations, both in country panels and with microdata within a country. Most of the literature focuses on the supply of credit by regulated financial institutions and real estate prices (Cerutti, Claessens, and Laeven, 2017; Acharya et al. 2022). Unlike

this work, we aim to evaluate the effects on systemic risk and distinguish between stringency and complexity when assessing the costs and benefits of regulation.

More closely related, our paper contributes to a growing empirical literature on regulatory complexity (Kalmenovitz, Lowry, and Volkova, 2025; Colliard, 2025; Greenwood, Hanson, Stein, and Sunderam, 2017) by introducing a measurement that separates complexity from stringency at the event level and by providing the first cross-country evidence on whether complexity is necessary or simply costly. We also extend the theoretical mechanism in Asriyan, Foarta, and Vanasco (2023) to bank regulation, showing that complexity is harmful only when not paired with stringency, consistent with their result that complex product design becomes a cost when it is not aligned with binding constraints on quality.

2 Measuring the stringency and complexity of financial regulation

We construct a country-quarter panel of financial regulation indices for 21 advanced economies over 1990Q1–2025Q4. The underlying event series is the International Monetary Fund's integrated Macroprudential Policy database (iMaPP; Alam et al., 2019; September 2025 vintage), which records every dated tightening or loosening of each of 27 policy tools through December 2024.

For European Union member states, we complement iMaPP with the European Central Bank's Macroprudential Policies Evaluation Database (MaPPED; Budnik and Kleibl, 2018), which records dated activations, recalibrations, and deactivations of macroprudential instruments across EU countries under a tool taxonomy that maps cleanly onto iMaPP's. MaPPED captures national-level events at a finer granularity — particularly for the country-level transposition of EU directives and the recalibrations and deactivations that occur between major reform packages — and provides a richer narrative description of each event than iMaPP. Because the pairwise rankings described in the next subsection take the narrative as input, the additional textual detail in MaPPED materially improves the quality of the resulting z-scores. We harmonise the two databases at the event level by mapping MaPPED tool codes onto the iMaPP taxonomy, deduplicating events that appear in both sources, and retaining the description with the richer narrative whenever the two records agree on the date and direction of the action. For the non-EU countries in our sample, the event panel comes from iMaPP alone.

Although both databases are labelled as macroprudential, their coverage is in fact broader: alongside genuinely macroprudential tools, they record changes in bank capital and liquidity requirements and other essentially microprudential measures that are likely to have system-wide implications (Alam et al., 2019; Budnik and Kleibl, 2018). For this reason, we talk broadly about bank regulation and classify the tools into five categories — *capital-based*, *liquidity-based*, *activity- and exposure-based*, *borrower-based*, and *structural/resolution* — without distinguishing between microprudential and macroprudential intent, because the same tools serve both.

Capital-based tools include systemically-important-financial-institution (SIFI) and global-systemically-important-bank (G-SIB) surcharges, sectoral capital requirements (such as higher risk weights on real estate exposures), and leverage-ratio backstops to risk-weighted measures. Liquidity-based tools include the Basel III Liquidity Coverage Ratio (LCR), the Net Stable Funding Ratio (NSFR), and reserve requirements. Activity- and exposure-based tools include large-exposure limits, concentration limits, and restrictions on foreign-currency lending to unhedged borrowers. Structural and resolution tools include total-loss-absorbing-capacity (TLAC) and minimum-requirement-for-own-funds-and-eligible-liabilities (MREL) rules, resolution planning ("living wills"), and ring-fencing of core banking from trading activities. Borrower-based tools include loan-to-value (LTV) caps, debt-to-income (DTI) and debt-service-to-income (DSTI) limits, and amortisation or maturity requirements on mortgages.

Given our focus on advanced economies (our sample includes Australia, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom, United States, and only one emerging market Turkey), we do not consider capital controls and other balance of payment restrictions that are important for emerging markets (Bruno, Shim, and Shin, 2017).

2.1 Stringency and complexity

Our objective is to capture the effectiveness and the costs of regulation. To achieve this, we develop two distinct measures of the bank regulation in a country: Stringency and complexity. Stringency and complexity are conceptually distinct dimensions of a regulatory action. We define *stringency* as the extent to which the action constrains the behaviour of a regulated institution: the breadth of institutions affected, the magnitude of the constraint imposed, and how binding the new requirement is in practice. A reduction of the mortgage LTV ceiling from 90 to 80 percent, a rise in the countercyclical capital buffer from 0.5 to 2 percent, and the activation of the G-SIB capital surcharge are all stringency-increasing actions of varying magnitude. We define *complexity* as the increase in the operational and compliance burden that the action imposes on the regulated institution: whether it introduces a new tool type that the bank must track separately, adds conditional thresholds or carve-outs, requires coordination across subsidiaries or jurisdictions, or layers on top of existing rules rather than replacing them. The two dimensions are orthogonal in principle. An LTV cap is highly stringent from the borrower's standpoint but simple for the originating bank to implement once the ratio is fixed; in contrast, the transposition of the Basel III Pillar-2 add-on with multiple conditional thresholds and institution-specific calibrations can be equally stringent yet substantially more complex.

Consistent with this intuition, and although our tournament ranks every iMaPP and MaPPED event along both dimensions, the complexity index is driven primarily by capital-based and liquidity-based events, whose compliance burden is borne by the regulated institution rather than by the loan applicant or counterparty. Borrower-

based events, such as loan-to-value and debt-service-to-income caps, receive systematically below-average complexity z-scores in the tournament, consistent with the fact that such rules impose a simple parametric constraint on the originating lender. For activity- and exposure-based tools, the two dimensions are conceptually difficult to separate and, in our data, their complexity z-scores move together with their stringency z-scores.

2.2 Construction of the indices

The central measurement challenge is that a naïve count of policy actions weights a marginal parameter revision and a sweeping Basel III transposition equally. Existing indices in the literature, in particular the Macroprudential Policy Index (MPI) of Cerutti, Claessens and Laeven (2017), aggregate binary “tool in place” flags and therefore capture neither intensity nor reversals. We replace the binary flag with an event-level intensity weight obtained from pairwise comparisons by a large language model. Specifically, we feed every iMaPP event, together with its narrative description where available, to the GABRIEL pairwise-ranking framework proposed by Asirvatham, Mokshi and Shleifer (2026), which conducts a multi-round Elo-style tournament — named after chess statistician Arpad Elo — asking, for every pair of events in the same direction, “which action had a larger impact on the stringency (respectively, complexity) of the regulatory regime?” Stringency and complexity are evaluated in separate tournaments under the definitions above. Tightenings and loosenings are also ranked in separate tournaments because cross-direction comparisons are not well defined. Each event receives a z-score along each dimension; the country-quarter index is the signed cumulative sum of those z-scores, truncated at zero so that an event judged to be of below-average impact cannot reverse the sign of its direction. Every country’s index starts at zero and moves only on quarters with an iMaPP event, so every inflection in the series is traceable to a specific, dated policy action. Our empirical analysis relies on single events, or put differently, on changes of the cumulative index. Therefore, the truncation at the start of the sample is irrelevant for our findings.

2.3 Intuition and comparability

Two interpretive points come first. Our indices are cumulative sums from a zero baseline in 1995Q1, so every number we report below, including the terminal levels cited later in the section and in the narrative that follows, measures accumulated regulatory change over the sample period rather than the absolute level of regulation in any country at any point in time. We do not attempt to measure the level, because we do not have a quantitative starting point against which to calibrate the 1990 baseline, and the empirical tests in this paper accordingly identify the effects of changes in regulation rather than the effects of overall regulatory burden. With that understood, the scale of the event-level z-scores themselves follows the logic of any z-score: the output of the pairwise tournament is first a continuous rank for each event, and we then subtract the pool mean and divide by the pool standard

deviation. A z-score of zero describes an event whose impact on stringency (respectively, complexity) is average among all events of the same direction in the sample. A z-score of +1 describes an event more impactful than roughly five-sixths of comparable events. A z-score of +2 places the event in the top two or three percent of the distribution; only the largest single-tool reforms reach this range. Individual events in the sample rarely exceed +4 on either dimension, and the bulk of events — those describing parameter recalibrations, phase-in continuations, or minor borrower-rule adjustments — sit between –1 and +1. Because the cumulative index adds the z-scores of all events up to quarter t, a country-quarter value of, say, +10 summarises the combined impact of roughly ten events of average magnitude, or of two to three comprehensive reform packages in which several tools moved together.

The z-scores are *comparable across countries and across time* because the standardisation pool is the entire sample. Every tightening ever taken by any of the 21 countries between 1990 and 2024 is benchmarked against the same mean and the same standard deviation; the mirror statement holds for loosenings. The tournament itself is blind to country and date labels: it observes only the tool code, the direction, and the narrative description of each action, and asks the model which of the two descriptions corresponds to a greater impact. A Japanese event from 2005 is therefore ranked against the same global distribution as a United States event from 2014, and a country-quarter level of +10 carries the same interpretation, regardless of the country or year it is attached to. This is the sense in which our indices are directly comparable across jurisdictions and across time, an advantage not shared by country-specific or vintage-specific calibrations used elsewhere in the literature.

Tightenings and loosenings are ranked in *separate tournaments*. Asking whether introducing a loan-to-value cap is larger than removing a reserve requirement is not a well-defined question, because the two actions move the regulatory regime in opposite directions. We therefore run four tournaments in total: tightenings ranked on stringency, tightenings ranked on complexity, and the mirror pair for loosenings. Each event receives a z-score within its own pool, and the sign of its contribution to the cumulative index is set by its direction (+1 for a tightening, –1 for a loosening). A tightening with z = +2 is large relative to other tightenings; a loosening with z = +2 is large relative to other loosenings; the two numbers are conceptually comparable but are not, strictly speaking, measured on the same underlying distribution.

The *complexity* measure is constructed by the same procedure as stringency but with a different comparison question. The stringency tournament asks which of two actions has a larger impact on how binding the regulatory regime is for the regulated institution — that is, on the breadth of institutions affected, the magnitude of the constraint, and the enforceability of the requirement. The complexity tournament asks which of two actions adds more operational and compliance burden — whether the action introduces a new tool type that the bank must track separately, adds conditional thresholds or carve-outs, requires coordination across subsidiaries or jurisdictions, or layers on top of existing rules rather than replacing them. The model reads the iMaPP narrative description of each event and makes these judgments pairwise, and the resulting complexity z-scores are standardised within the tightening

and loosening pools exactly as for stringency. Consistent with the conceptual argument above, although the tournament assigns a complexity z-score to every event, the cumulative complexity index is, in practice, dominated by capital-based and liquidity-based events. Borrower-based events receive below-average complexity z-scores in the sample, and for activity- and exposure-based events, the complexity z-scores move nearly one-for-one with the stringency z-scores. The interpretation of the complexity index is therefore anchored on capital and liquidity tools, where the distinction between stringency and complexity is sharp.

2.4 Advantages over existing indices

Our indices improve on the Cerutti-Claessens-Laeven MPI and related aggregate measures in four respects. First, they are *event-dated and reversible*: every movement in the index corresponds to a specific iMaPP/MaPPED action and the series falls when policy is loosened, rather than ratcheting upward as in presence-only indices. Second, they *weight by intensity rather than by presence*, so a Basel III transposition and a marginal LTV recalibration receive different weights. Third, they *separate stringency from complexity* at the measurement stage through distinct tournament prompts, whereas the MPI and similar indices collapse both into a single scalar. Fourth, they exploit the full 27-tool iMaPP taxonomy, whereas the MPI maps to twelve instruments, one of which — interbank-exposure limits — has no iMaPP counterpart at all.

2.5 Limitations

Two limitations are worth noting. First, iMaPP does not separately code structural and resolution tools such as TLAC, MREL, resolution planning, or ring-fencing. Second, iMaPP assigns a single Liquidity code to both the LCR and the NSFR, so our liquidity-based index does not distinguish between short-horizon (30-day) and medium-horizon (1-year) liquidity regulation.

2.6 Macroprudential events

Figure 1 plots every dated iMaPP event in the panel by country and quarter. Filled upward triangles are tightenings; open downward triangles are loosenings.

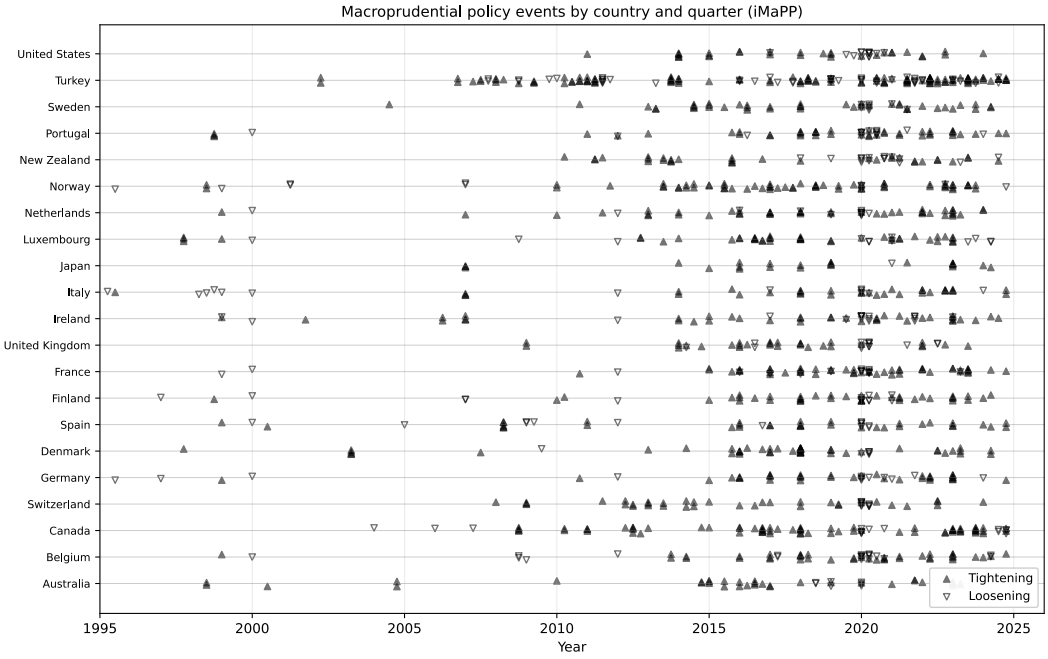

Figure 1. Macroprudential policy events by country and quarter. Each marker is one tool-level entry from the iMaPP database (Alam et al., 2019) for each of the 21 sample countries; filled upward triangles are tightenings and open downward triangles are loosenings. Vertical jitter is added within each country row for readability.

Three patterns are worth flagging. First, before 2010, the events are sparse: bank regulation was a thin and primarily borrower-based affair across most of the sample. Second, tightenings cluster after crises – the post-2008 wave of Basel III transpositions is visible across all 21 countries between 2013 and 2016. Third, loosenings cluster during crises, most visibly in 2000, 2007–2008, 2011–2012 (the European sovereign-debt episode), and 2020 (COVID-19). When we examine how systemic risk, stock returns, and private credit respond around these events, this temporal structure provides important context: tightenings tend to occur in the recovery phase, and loosenings during the downturn itself.

2.7 Distribution of stringency and complexity

Table 1 quantifies the dimension that Figure 1 cannot show. It reports summary statistics for our cumulative stringency and complexity indices, alongside two descriptive comparators: the unweighted iMaPP tool count and our replication of the Cerutti–Claessens–Laeven Macroprudential Policy Index (Cerutti, Claessens, and Laeven, 2017).

Measure	N	Mean	SD	p10	Median	p90
Stringency	2,340	3.212	5.880	-2.520	1.474	11.136
Complexity	2,340	1.251	6.951	-5.483	-0.633	10.250
iMaPP tool count	2,604	5.290	5.041	0.000	4.000	13.000

Cerutti MPI	2,604	2.349	2.316	0.000	2.000	6.000

Table 1. Regulation-measure summary statistics on the merged country-quarter analysis sample, 1995Q1 onwards, 21 countries. Stringency and complexity are the cumulative signed sums of event-level z-scores from our pairwise GABRIEL tournaments on iMaPP and ECB MaPPED events; tool count and Cerutti MPI are descriptive comparators.

Two features stand out. First, both stringency and complexity rise on average over the sample period, with most countries ending in a slightly more stringent and slightly less complex state than they started; the means are below the medians for complexity, indicating that a small number of episodes account for a disproportionate share of the cumulative change. Second, both indices are highly positively skewed: in a few countries and quarters, the regime becomes very stringent and very complex relative to the panel average. This skewness is the empirical analogue of the conceptual argument above. Complexity does not move smoothly with stringency, and a single-scalar regulatory measure would obscure exactly the variation we wish to exploit.

3 Main regulation and deregulation episodes in the sample

This section uses the event-level z-scores underlying our cumulative indices to describe the main regulatory and deregulatory episodes in the 1990Q1–2024Q4 sample. We focus on the largest advanced economies and flag where stringency and complexity diverge.

3.1 The early-1990s reserve-requirement deregulation (1990–1995)

Although the empirical sample begins in 1995Q1 for the credit outcomes and in the early-to-mid 2000s for systemic risk, iMaPP and MaPPED record regulatory events from earlier. We briefly discuss the pre-1995 reserve-requirement deregulations below to provide background on the regulatory baseline against which subsequent events are measured; from 1995Q1 onwards, every event we describe is observed directly in iMaPP, complemented by MaPPED for European Union member states.

The earliest — and across several countries, largest — episode in the sample is a wave of loosenings of reserve requirements and quantitative credit controls that preceded the 1990s liberalisation of bank funding. In Japan, the Bank of Japan's October 1991 cut to reserve requirement ratios across all deposit categories, accompanied by a relaxation of corporate-loan restrictions, produced the single largest negative step in the Japanese series (stringency –4.8, complexity –4.2 in 1991Q4). Germany's elimination of the progressive reserve-requirement schedule in March 1994, followed by its 1995 and 2000 reserve-ratio cuts, accounts for Germany's negative terminal complexity value (–3.9 in 2024Q4): essentially, the

whole German series was offset by these early simplifications before Basel III re-added complexity. Italy's May 1994 reform of the marginal reserve-requirement ratio (–3.1, –3.0), Canada's June 1994 elimination of reserve requirements, and the 2000 ECB-coordinated reserve-framework change in France, Germany, the Netherlands, Spain, and Belgium show the same pattern: stringency and complexity fall together because eliminating a tool both relaxes the constraint and removes the parameter from the rulebook.

3.2 Pre-crisis mortgage-risk tightening in Ireland, Italy, and Spain (2006–2008)

Three countries tightened sharply before the 2008 financial crisis, all via the Basel II transposition interacting with housing-market concerns. Ireland's 2007Q1 Circular, which imposed stricter criteria for preferential risk-weighting of residential mortgages and raised the minimum risk weight on high-LTV commercial exposures, produced a +7.8 stringency jump and a +6.5 complexity jump, making it Ireland's single largest event. Italy's January 2007 application of the 35% risk weight to mortgages secured on owner-occupied residential property (+6.0 / +7.8) and Spain's May 2008 Circular 3/2008 on commercial- and residential-real-estate exposures (+13.7 / +18.5 — the single largest complexity jump in the entire sample) illustrate the same phenomenon. These are capital-based measures and so, consistent with our definition, they load heavily on both dimensions: the risk-weight matrices are both binding and operationally complex.

The contrast between the Italian and Spanish events is illustrative. Italy's measure was a single parametric change — a 35% risk weight applied to mortgages on owner-occupied residential property within an otherwise unchanged Basel I framework — stringent for the affected exposures and modestly complex for banks, who had one new risk-weight bucket to track and report. Spain's Circular 3/2008, in contrast, was the comprehensive Banco de España transposition of Basel II. In a single rule, it introduced the standardised, foundation- Internal Ratings-Based approach (IRB), and advanced-IRB approaches to credit risk, the market- and operational-risk frameworks, a multi-tier real-estate risk-weight matrix conditioning jointly on loan type, LTV band, and seasoning, and a Pillar 2 supervisory add-on regime with bank-specific calibrations. Spanish banks had to operate it in parallel with the pre-existing dynamic-provisioning regime, which carried its own accounting and regulatory-capital implications and effectively doubled the methodological burden. The two events share a regulatory direction — both raise capital requirements on real-estate exposures — but differ sharply in operational footprint, which is precisely what the complexity dimension is designed to capture. Spain accordingly scores +13.7 / +18.5 (the highest complexity jump in the sample) while Italy scores +6.0 / +7.8. The Italy and Spain examples therefore illustrate cleanly what stringency and complexity are designed to capture: stringency tracks how binding the action is for the regulated institution, complexity tracks the operational and compliance burden it imposes, and the two dimensions move together only when the regulatory action both raises a quantitative ceiling and reshapes the methodological infrastructure that supports it.

3.3 Crisis-response borrower protection (2009–2011)

A smaller but distinctive phase follows: borrower-relief measures that loosen stringency without meaningfully simplifying the rulebook. Spain's March 2009 mortgage-payment moratorium under Royal Decree 97/2009 moved stringency by –3.5 but complexity by only –1.9; when the moratorium expired in 2011 and deferred payments resumed, stringency jumped back (+3.3), and a new layer of transitional rules added a further +2.4 in complexity. This is among the clearest examples in the sample of complexity ratcheting up even when stringency round-trips, because temporary regimes leave procedural residue.

3.4 Basel III transposition: the dominant global episode (2013–2016)

This is the largest coordinated tightening in the sample and accounts for most of the terminal stringency and complexity levels across all advanced countries. The United States' October 2013 final rule on regulatory capital issued by the Federal Reserve Board and the Office of the Comptroller of the Currency, which became effective 2014Q1, is the single largest US event (+7.5 stringency, +8.9 complexity); it combined Basel III capital requirements, leverage-ratio reporting, DSTI-consistent mortgage rules, and the SIFI surcharge phase-in. The 2015Q1 continuation of the phase-in of minimum risk-based capital ratios and the 2016Q1 continuation of the LCR phase-in add a further +6.6 combined. Switzerland — which, at terminal stringency +26.6 and terminal complexity +25.0, is the tightest country in the sample — accumulates most of its index through the same window: the February 2013 activation of the sectoral countercyclical capital buffer at 1%, the June 2012 higher risk weights on high-LTV mortgages, and the 2014Q2 increase of the sectoral buffer to 2% are the four top Swiss moves. France, Italy and Spain show parallel Basel III transposition events in 2014Q1 and 2015Q1 that contribute broadly equally to stringency and complexity. Across Basel III events, the two z-scores move almost in lockstep (within-category correlation of roughly 0.9), consistent with the idea that a richer risk-weight matrix is by construction both more binding and more to track.

3.5 Targeted borrower-based tightenings (2010–2015)

Alongside Basel III, several countries activated housing-specific borrower-based rules. The Netherlands' 2010 maximum loan-to-foreclosure-value ratio (+2.3 / +2.3) and 2011 Code of Conduct revision limiting interest-only mortgages (+3.0 / +2.8) are its two largest events; together with the 2007 Debt-Service-To-Income ratio (DSTI) introduction, they drive the Dutch terminal stringency (+14.4). The United Kingdom's 2009–2011 Mortgage Market Review, which imposed borrower-affordability assessments on lenders (+3.0 / +3.5 in 2009Q1, the UK's largest event), and the 2014 Financial Policy Committee stress-rate requirement, under which mortgage affordability must be tested against a three-percentage-point Bank Rate shock, are tightenings of borrower rules that nonetheless rank high on complexity because they introduce new procedural obligations on lenders. These are the cases where the

stringency–complexity distinction is sharpest: the rule is not especially binding from the bank's capital standpoint but is operationally demanding.

3.6 COVID-era relaxations and the 2022–2024 re-tightening

The panel shows COVID-era loosenings (countercyclical-buffer releases, targeted liquidity relief) in roughly a dozen countries in 2020, but none are among the largest moves in the sample: the pairwise-ranked z-scores place them consistently below crisis-response or Basel III events in magnitude, presumably because the COVID relaxations were mostly parametric adjustments within existing tools rather than structural changes. More visible in the series is the 2022–2023 re-tightening wave: Germany's 2023Q1 combined capital, SIFI, and loan-restriction tightening (+3.5 / +1.8) is the largest German event; France's 2023Q3 formalisation by the Haut Conseil de Stabilité Financière of the 35% DSTI limit (+1.7 / +2.1), the United Kingdom's 2022Q1 restoration of capital buffers (+2.2 / +0.9), and the Netherlands' 2023Q1 multi-tool tightening (+1.8 / +1.3) are close seconds. These events are notable because they rebuild stringency substantially faster than complexity — the stringency-to-complexity ratio in the 2022–2024 episode is often above one, suggesting that the post-COVID rulebook did not add much new procedural infrastructure; it reactivated existing tools.

Within this re-tightening wave, two events stand out as canonical late-sample examples of capital-based, time-varying tools that the indices treat as both stringent and operationally complex. Germany's first-ever activation of the countercyclical capital buffer (CCyB), announced by BaFin in January 2022 at 0.75% and effective from February 2023, is the capital-based component of the German 2023Q1 event noted above; it is bundled in iMaPP with the parallel Systemically Important Financial Institution (SIFI) surcharge calibration and the residential-real-estate systemic-risk buffer that BaFin introduced at the same time. France's reactivation of the countercyclical buffer at 0.5% in April 2023 and its subsequent increase to 1% in October 2024 (effective January 2026) round out the late-sample French capital tightening. Both the German and French CCyB events score as moderately positive on stringency and on complexity, in line with the conceptual argument above: a CCyB action raises a quantitative capital ceiling and adds a conditional layer that interacts with the conservation buffer, the SIFI surcharge, and any sectoral buffers in place, requiring institution-by-institution computation of the geographically weighted exposure that determines the applicable rate.

3.7 Outliers worth flagging

Japan is deliberately light on macroprudential activity — only 14 rated events and a nearly flat series (terminal +1.1 stringency, –0.8 complexity). The only material moves are the 1991 deregulation episode and a handful of Basel III transposition steps. This matches the institutional fact that Japanese macroprudential policy relies more on supervisory guidance than on formal tool activations. Switzerland is the opposite extreme — terminal stringency +26.6 and complexity +25.0, roughly twice

the Euro-area big three. The Swiss index reflects the early 2012 countercyclical-buffer activation, sectoral risk-weight increases targeted at the real-estate-exposed domestic banks, and subsequent supervisory add-ons by the Swiss Financial Market Supervisory Authority — a relatively rare case in which all three tool families (capital, liquidity, borrower) were activated in close succession. Germany is the clearest demonstration of why complexity and stringency should be tracked separately. Germany's terminal stringency is positive (+5.2) but its terminal complexity is negative (–3.9). The country added Basel III requirements like everyone else, but these were more than offset in the complexity dimension by reserve-requirement simplifications enacted by the Bundesbank in the 1990s and during the 1999–2000 transition to the harmonised Eurosystem framework, before reserve-requirement authority moved to the ECB. A single-scalar index of regulatory stance would miss this.

3.8 Stringency versus complexity: four empirical patterns

First, loosenings via reserve requirements and quantitative credit controls move stringency and complexity symmetrically downward. This tool family is, in our sample, the one most prone to being eliminated rather than recalibrated. Second, capital-based Basel III events have a stringency-to-complexity ratio near one — simultaneous and roughly equal increases in both dimensions. Third, borrower-based events (loan-to-value and debt-service-to-income caps) have stringency z-scores that exceed their complexity z-scores on average, consistent with our conjecture that borrower tools are stringent for the borrower but operationally simpler for the originating bank. Fourth, the 2022–2024 re-tightenings restore stringency more than complexity: a regime that was already complex became marginally more binding without much new procedural infrastructure.

4 Sample and outcomes

Our analysis combines a country-quarter panel of macroprudential policy actions with three families of outcomes – systemic risk, stock-market valuations, and private credit – over the period 1995Q1 to 2025Q3 for 21 advanced economies.

4.1 Country panel

Table 2 lists our country panel and the number of dated tightening and loosening events recorded for each country, together with the first and last quarters in which both our systemic-risk measure and the BIS series on private credit to GDP are observed.

Country	ISO3	Tightenings	Loosenings	First SRISK	First credit	Last quarter
Australia	AUS	37	7	2000Q3	1995Q1	2025Q3
Belgium	BEL	41	20	2002Q1	1995Q1	2024Q4
Canada	CAN	71	21	2000Q3	1995Q1	2025Q3

Denmark	DNK	48	6	2002Q1	1995Q1	2025Q3
Finland	FIN	43	18	2003Q2	1995Q1	2025Q3
France	FRA	49	21	2000Q3	1995Q1	2025Q3
Germany	DEU	35	15	2000Q3	1995Q1	2025Q3
Ireland	IRL	38	23	2000Q3	1995Q1	2025Q3
Italy	ITA	34	16	2000Q3	1995Q1	2025Q3
Japan	JPN	26	4	2000Q3	1995Q1	2025Q3
Luxembourg	LUX	43	12	2007Q4	1995Q1	2014Q2
Netherlands	NLD	44	13	2002Q1	1995Q1	2025Q3
New Zealand	NZL	33	16	2007Q1	1995Q1	2007Q3
Norway	NOR	66	12	2000Q3	1995Q1	2025Q3
Portugal	PRT	44	18	2001Q4	1995Q1	2025Q1
Spain	ESP	35	14	2002Q3	1995Q1	2025Q3
Sweden	SWE	43	8	2000Q3	1995Q1	2025Q3
Switzerland	CHE	36	10	2000Q3	1995Q1	2025Q3
Turkey	TUR	121	48	2001Q1	1995Q1	2025Q3
United Kingdom	GBR	30	16	2000Q3	1995Q1	2025Q3
United States	USA	30	16	2000Q3	1995Q1	2025Q3

Table 2. Country coverage. iMaPP event counts and the first and last quarters where both SRISK and the BIS private-credit-to-GDP series are observed; observations are at the country-quarter level. Quarter format YYYYQn.

Our sample covers the United States, the United Kingdom, Japan, Turkey, and 18 European economies. Although the iMaPP regulatory series begins in 1995Q1, the systemic-risk measure SRISK from V-Lab – described below – is available only from the early-to-mid 2000s in most countries, so the systemic-risk regressions operate on a more limited window. Tightenings outnumber loosenings in nearly every country, consistent with the post-2008 ratchet pattern documented in Section 2.

4.2 Outcomes

We focus on four outcomes that summarise systemic risk, the value created by financial intermediation, and the volume of private credit. Our primary measure of systemic risk is the natural log of SRISK, a widely used market-based indicator of systemic risk exposure in the banking sector developed by Brownlees and Engle (2017) and computed weekly by NYU Stern's Volatility Laboratory (V-Lab). Since its introduction, SRISK has been adopted extensively in the academic literature on financial stability — including Acharya, Engle and Richardson (2012), Engle, Jondeau and Rockinger (2015), and Acharya, Pedersen, Philippon and Richardson (2017) — and by central banks and international institutions as a forward-looking gauge of banking-sector fragility that complements supervisory stress tests by relying entirely on publicly observable market data.

SRISK measures the conditional capital shortfall of the banking system in a country in the event of a systemic crisis. The country-level series is constructed as the sum of positive firm-level SRISK across all listed financial institutions domiciled in that

country, where each firm's SRISK is the additional equity capital it would need to raise to maintain a prudential capital ratio (8% of quasi-assets) following a severe market downturn — operationalised as a 40% decline in the broad equity index over a six-month horizon. The country aggregate therefore answers a simple question: if a global financial crisis materialised tomorrow, how many dollars of new equity capital would the domestic financial sector, in aggregate, need to raise to restore adequate capitalisation? Larger values indicate greater systemic vulnerability, reflecting some combination of high leverage, large size, and elevated comovement of bank equity returns with the broader market. We work with the natural logarithm of SRISK to address the strong right skew and the substantial cross-country dispersion in the size of financial systems that mechanically scale the dollar shortfall.

Several considerations justify our reliance on SRISK as the headline outcome. First, by construction, the measure tracks the variable that bank regulation is most directly designed to contain — the conditional capital shortfall of the banking system in a systemic stress — so that the alignment between the policy objective and the outcome measured is unusually tight. Second, SRISK is forward-looking: because the inputs are equity returns and equity-market capitalisation, the measure incorporates investors' expectations of future leverage, asset volatility, and tail comovement, whereas supervisory stress-test outcomes and book-value capital ratios summarise the regulated balance sheet at the reporting date and are subject to known smoothing of credit losses and risk weights. Third, the measure has documented predictive content for systemic episodes — including the 2008 financial crisis (Brownlees and Engle, 2017), the 2011 European sovereign-debt episode (Engle, Jondeau, and Rockinger, 2015), and the 2023 U.S. regional-bank failures (Acharya, Brunnermeier, and Pierret, 2024) — which is the reason it is monitored as a forward-looking gauge by policymakers. Fourth, SRISK relies only on publicly observable market data, which allows us to assemble a cross-country panel of comparable readings without conditioning on supervisory access. Alternative market-based indicators, such as the marginal expected shortfall of Acharya, Pedersen, Philippon, and Richardson (2017) and the CoVaR of Adrian and Brunnermeier (2016), capture related dimensions of systemic risk but do not deliver a dollar-scaled capital shortfall directly comparable to the additions that bank regulation targets.

A natural concern with a market-based measure built on listed entities is that risk may relocate to unlisted parts of the financial system that SRISK does not capture. Two features mitigate this concern. First, SRISK is computed on the consolidated balance sheet of each listed bank holding company, so within-firm relocation of risk from a regulated bank subsidiary to a lightly-regulated nonbank affiliate — a channel that Cetorelli and Kundu (2026) document for the 2015 U.S. Basel III phase-in — is captured at the holding-company level rather than missed. Second, the listed-financials perimeter is wide: it includes not only banks but listed insurers, broker-dealers, and asset managers, so non-bank intermediaries that operate as listed entities are also covered. The remaining measurement gap is restricted to fully independent unlisted intermediaries — private credit funds, private-equity-backed direct lenders, securitisation special-purpose vehicles, and similar — to which credit and risk-taking may flow without being captured by SRISK.

Stock-market valuations offer a forward-looking, market-based summary of the stability-growth trade-off that bank regulation must manage. A tightening that reduces systemic risk without commensurate restraint on profitable intermediation should increase aggregate equity values; a tightening that imposes operational and compliance costs without a corresponding stability gain should depress bank equity values relative to the broader market. Because equity prices integrate expected cash flows, compliance costs, and risk premia into a single number, they provide a compact assessment of the net effect of regulation on the regulated sector and on the rest of the economy, complementing the SRISK-based stability measure on the cost side of the stability-growth trade-off.

To evaluate the effects of bank regulations on economic performance, we therefore estimate the effects of regulatory events on cumulative country level stock market returns. While the evidence is descriptive and country level performance has many drivers, market capitalization has a long history in the finance-and-growth literature, where it has been used to capture changes in the value created by the financial system. Stock market returns integrate expected cash flows and discount rates into a single number and allow for an analysis at quarterly frequency for every country in our sample. A regulatory tightening that constrains profitable real-side activity should affect returns negatively.

The change in the log of bank-sector capitalisation is the cumulative bank-sector equity return and captures the market-implied franchise value of regulated banks — the present value of expected future profits from financial intermediation, net of expected compliance and adjustment costs. A tightening that imposes operational and compliance costs without commensurate stability gains should depress it. The change in the log ratio of bank to total capitalisation is the bank-relative-to-market return: it strips out the common market component and isolates the bank-specific response, asking whether bank equity values fall relative to the broader market after a regulatory event.

Finally, to capture the volume of credit extended within and outside the banking system, we use the BIS measure of private credit to GDP. The credit aggregate captures all credit extended to the private non-financial sector from all sectors (banks, domestic non-bank financial intermediaries, non-financial creditors, and cross-border lenders, including holders of debt securities issued by the sector), expressed as a percentage of GDP. As in the earlier tests, we examine how this variable changes in response to changes in regulation.

Prices and quantities provide complementary readings of intermediation. The equity returns we examine integrate expected future cash flows and risk premia forward-looking, whereas the credit aggregate captures realised intermediation contemporaneously. The BIS Private credit / GDP series also has a wider perimeter than SRISK or the equity measures: it includes non-bank lenders, cross-border creditors, and debt-securities holders, so that credit migration to the non-bank perimeter shows up in the credit total even when it may not be directly measured by listed-financials' equity values.

Outcome	N	Mean	SD	p10	Median	p90
log SRISK	1,852	10.148	2.253	7.221	10.598	12.681
log total stock	2,520	6.466	1.691	5.123	6.215	7.523
log bank stock	2,142	5.147	1.305	4.173	5.028	6.620
log(bank/total)	2,142	-1.503	1.409	-2.713	-1.166	-0.264
Private credit/GDP (pp)	2,583	177.987	67.076	109.640	169.900	259.780

Table 3. Outcome summary statistics on the merged country-quarter analysis sample, 1995Q1 onwards, 21 countries.

Table 3 reports cross-country, cross-quarter summary statistics for the four outcomes. The most informative single entry is the distribution of private credit to GDP, which averages roughly 178% of GDP in the panel, with 10th and 90th percentiles of 110% and 260% – a substantial spread reflecting both cross-country differences in financial development and the sustained credit expansion of the past two decades. The remaining rows describe the dispersion of the (log) outcome levels; in the empirical analysis below, we work with within-country changes around dated regulatory events, which absorb level differences across countries and isolate the responses to policy. We have also estimated Table 1 in changes rather than levels and find no qualitative impact on the summary picture.

5 Methodology

Evaluating the causal effect of bank regulation on financial outcomes is challenging because the relevant counterfactual — what would have happened in a country in the absence of a particular regulatory action — is never observed. This difficulty is common to the entire empirical literature on financial cycles: regulatory actions are typically introduced in response to identifiable financial conditions, are bundled with other policy and macroeconomic developments, and often affect outcomes through channels that operate over horizons longer than the typical policy review window. Our empirical strategy is designed to approximate the missing counterfactual as closely as we can, by combining a long panel of clean-control country–quarters, a flexible factor structure for country-level outcomes, and an ex-ante partition of events along the stringency and complexity dimensions of our indices. The estimates that follow should be interpreted as suggestive and intended to motivate in-depth analyses of specific interventions affecting the stringency and complexity of regulations.

Our empirical question is whether tightening or loosening the stringency or complexity of bank regulation affects banking-system risk, asset prices, and the credit aggregates over a multi-year horizon. We answer it with country–quarter local projections built on imputation-based generalised synthetic control: an interactive fixed-effects model is fit on a sample of clean-control cells, the no-treatment counterfactual is imputed for every cell, and horizon regressions on the imputation residual recover the impulse response. The factor-loading counterfactual generalises

the additive two-way fixed effects of Dube, Girardi, Jordà, and Taylor (2025) along the lines of Bai (2009) and Xu (2017), allowing each country to load heterogeneously on a small number of latent common shocks rather than absorbing only their average.

Each event-quarter is classified into one of eight mutually exclusive cells, the product of three binary partitions: the event direction (tighten or loosen), the magnitude of the event-quarter stringency shock (above or below the direction-specific median, "high" or "low" stringency), and the magnitude of the event-quarter complexity shock (above or below the direction-specific median, "high" or "low" complexity). The horizon regression then estimates a separate impulse response for each cell relative to a common pool of clean controls.

5.1 Interactive Fixed Effects (IFE)-imputation specification

For outcome $y_{c,t}$ in country c and quarter t, we posit the no-treatment data-generating process

$$y^0_{c,t} = \alpha_c + \delta_t + \lambda'_c f_t + \beta' X_{c,t-1} + \varepsilon_{c,t}. \qquad (1)$$

where $f_t \in \mathbb{R}^r$ is a vector of $r = 3$ latent factors common across countries, $\lambda_c \in \mathbb{R}^r$ is a country-specific loading, and $X_{c,t-1}$ stacks the seven Global Macro Database (Müller, Xu, Lehbib, and Chen, 2025) controls, including inflation, the short-term interest rate, the current-account-to-GDP ratio, the government-debt-to-GDP ratio, the unemployment rate, the real-GDP growth rate, and the change in the bilateral USD exchange rate, evaluated at $t - 1$. The Global Macro Database is annual; we obtain a quarterly panel by linear interpolation between annual observations. The factor structure absorbs latent dynamics with country-specific exposure that the additive country fixed effect α_c alone cannot capture; two countries with similar $\hat{\lambda}$ act as donors for each other in the imputation, the same role they would play under synthetic control with unit weights.

Stage 1: fit the IFE model on clean cells. We estimate $(\hat{\alpha}_c, \hat{\delta}_t, \hat{\lambda}_c, \hat{f}_t, \hat{\beta})$ by Bai (2009) iteration on clean-control cells, withholding treated and recently treated cells from the fit. The procedure alternates a singular-value-decomposition step on the residual matrix to update factors and loadings with an OLS step on the fixed effects and covariates; missing or withheld cells are imputed in expectation by the current model fit at each iteration.

Stage 2: impute the counterfactual. For every country–quarter we form $\hat{y}^0_{c,t} = \hat{\alpha}_c + \hat{\delta}_t + \hat{\lambda}'_c \hat{f}_t + \hat{\beta}' X_{c,t-1}$ and define the imputation residual $\tilde{y}_{c,t} \equiv y_{c,t} - \hat{y}^0_{c,t}$.

Stage 3: horizon LP on residuals. At each forecast horizon h, we estimate

$$\tilde{y}_{c,t+h} - \tilde{y}_{c,t-3} = \alpha_h + \sum_{j=1}^{8} \theta^h_j D_j(c,t) + \varepsilon^h_{c,t+h}. \qquad (2)$$

where $D_j(c,t)$ is the indicator that the event in country–quarter (c,t) falls in cell j (defined below), and the intercept α_h absorbs the mean clean-control residual at horizon h. The left-hand side is the change in the imputation residual between the

pre-event base period $t - 3$ and the forecast period $t + h$. The base period $t - 3$ rather than $t - 1$ allows up to two quarters of anticipation. We estimate the response at seventeen horizons, $h \in \{-6, -5, -4, -2, -1, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$; $h = -3$ is excluded because the left-hand side is mechanically zero. Negative horizons supply a parallel-trends placebo. The factor structure of stage 1 absorbs the latent slow-moving heterogeneity that under additive two-way fixed effects would otherwise load onto the event-quarter shock.

5.2 Cell construction

Each event-quarter is assigned to exactly one cell based on three binary partitions.

Event direction. A country–quarter has direction $d \in \{\text{tighten}, \text{loosen}\}$ if it contains at least one iMaPP event of that direction. Quarters that contain events of both directions are assigned to whichever direction has the larger total absolute tournament-weighted contribution.

Stringency and complexity magnitudes. For each event-quarter, we form the absolute, z-scored magnitudes $|\Delta S^z_{c,t}|$ and $|\Delta C^z_{c,t}|$ from the direction-signed event-quarter shocks

$$\Delta S_{c,t} = \sum_{e \in \mathcal{E}_{c,t}} \text{sign}(e)\, \max\{z^S_e, 0\}, \quad \Delta C_{c,t} = \sum_{e \in \mathcal{E}_{c,t}} \text{sign}(e)\, \max\{z^C_e, 0\}. \qquad (3)$$

where $\mathcal{E}_{c,t}$ is the set of iMaPP events in country c in quarter t, $\text{sign}(e) = +1$ for a tightening and -1 for a loosening, and z^S_e and z^C_e are event-level tournament z-scores for stringency and complexity, respectively. After construction, we winsorise each series at the 1%/99% quantiles over event cells and standardise by the nonzero standard deviation.

Direction-specific medians. We define "high stringency" (HS) for an event-quarter as $|\Delta S^z|$ above the median of $|\Delta S^z|$ within the direction's event-quarters; "low stringency" (LS) is the complement. "High complexity" (HC) and "low complexity" (LC) are defined analogously. The eight cells $D_j \in \{\text{T·HS·HC}, \text{T·HS·LC}, \text{T·LS·HC}, \text{T·LS·LC}, \text{L·HS·HC}, \text{L·HS·LC}, \text{L·LS·HC}, \text{L·LS·LC}\}$, with T = tighten and L = loosen, partition the event-quarters. The coefficient θ^h_j in equation (2) is the average h-period change in the imputation residual for events of cell j relative to the clean-control mean.

All impulse responses reported in the figures and in the headline ATT table use the direction-specific median split described above: an event-quarter is classified as high (low) stringency if its absolute, z-scored stringency shock lies above (below) the median of the corresponding direction-specific magnitude pool, and analogously for complexity. We choose the median split for the headline because it yields the largest treated samples per cell and the most balanced comparison across the four cells.

5.3 Sample: clean controls and treated cells

The estimation sample at horizon h retains a country–quarter cell (c,t) if either it is treated (at least one event in quarter t, assigned to one of the eight cells) or it is a clean control (no event in the past four quarters and no event in the next two quarters). The clean-control window of one year matches the canonical specification used to identify the impulse response from short-run cumulative changes in regulatory state, and the forward window of two quarters excludes cells that are likely to be anticipating a near-future event from the control pool.

5.4 Inference

All standard errors, confidence intervals, and p-values come from a single procedure: a within-quarter full-reshuffle randomization test. It asks whether a cell's estimated response reflects the stringency–complexity labeling we measure or could instead be produced by chance. For cell j the statistic is the coefficient θ_j of equation (2), equal to the post-event-window mean of the cell-j average imputation residual relative to the clean controls. Holding the calendar quarter fixed, we reassign the cell-j treatment label at random across all country–quarters in that quarter — drawing a pseudo-cell of the same size from the whole panel — and recompute the statistic against the same clean-control reference. Across B = 2,000 relabelings, this traces the null distribution of θ_j under random labeling; the exact two-sided p-value is the share of relabelings whose permuted statistic is at least as large in absolute value as the one we estimate, and the 90% confidence interval inverts the test.

Permuting within the calendar quarter holds fixed any common shock the factor structure leaves in the residual; drawing the pseudo-cell from the entire panel rather than from the clean controls alone keeps the reshuffle pool well populated in every quarter, where a within-clean-control permutation would be degenerate in quarters with few treated units. The point estimates are the LP-IFE coefficients of equation (2) and are unchanged by the choice of inference; the randomization test only assesses them. The headline event-study figures (the quarterly, median-split tightening responses) apply the same test one horizon at a time, plotting the per-horizon coefficient with its pointwise 90% randomization band over the pre-event placebo horizons and the per-outcome average treatment of the treated (ATT) window: $h \in \{-6, -5, -4, -2, -1, 1, \ldots, 8\}$ for the forward-looking outcomes (SRISK, equity) and $h \in \{-6, -5, -4, -2, -1, 1, \ldots, 12\}$ for the credit aggregates.

We use this within-quarter full reshuffle, applied to the per-outcome ATT windows reported in Table 4.

6 Results

This section reports the empirical evidence on whether stringency and complexity reduce systemic risk, and whether complexity is necessary or simply costly. We emphasise that the findings reported below are not directly informative about the

ongoing regulatory changes. The sample is limited to advanced economies over a single post-Basel regulatory cycle, the identification rests on the LP-IFE design absorbing latent country-specific dynamics that the regulatory shock would otherwise inherit, and the estimates speak to the average direction of the response within each stringency-complexity cell rather than to the causal effect of any single regulatory action. We accordingly read the magnitudes that follow as suggestive rather than as definitive point estimates.

Two findings are specified ex ante by the theoretical mechanism, and both cross the 90% within-quarter randomization-inference interval at the cell-ATT level. First, tightening events that raise both stringency and complexity (the High S, High C cell) reduce systemic risk by an average of approximately 41 log points over the eight quarters following the event — roughly a one-third reduction in the level of country-aggregate SRISK — and the reduction is sustained across the window. Tightenings that raise stringency without raising complexity (the High S, Low C cell) produce a smaller and statistically insignificant decline that fades within roughly four quarters. The persistence of the systemic-risk reduction in High S, High C, contrasted with the transience in High S, Low C, is the central finding of the paper. We interpret it as evidence of regulatory arbitrage: when a stringent rule is operationally simple, banks can take on equivalent risk in adjacent activities or vehicles while remaining technically compliant, and the systemic-risk gain decays. When the rule is operationally demanding, such substitution is harder, and the gain is preserved.

A natural micro-foundation for this interpretation is provided by Cetorelli and Kundu (2026), who document that the 2015 U.S. Basel III phase-in induced bank holding companies to extract equity from lightly regulated nonbank subsidiaries and redeploy it inside their depository institutions, raising bank capital ratios by 5 to 8 percentage points while leaving consolidated equity, assets, and lending unchanged. Risk was relocated to the less regulated side of the firm — nonbank affiliates expanded into consumer lending, operated at lower capital ratios, and saw credit quality deteriorate — rather than eliminated. The same logic implies that a stringent capital rule applied narrowly to the bank entity, without parallel constraints on the wider organisational perimeter, is operationally simple enough to be arbitraged within the holding company. Complexity in the rulebook — consolidated requirements, cross-subsidiary thresholds, and reporting depth across affiliates — is precisely what closes off this within-firm reallocation, and its absence is consistent with the transience we observe in the High S, Low C cell.

This contrast operates without aggregate equity or credit costs: the broad-market equity index is positive in all four tightening cells, and private credit to GDP shows no statistically detectable contraction, including in cells that combine high stringency with high complexity. The complexity that the post-crisis rulebook adds, therefore, appears to be neither a free lunch nor pure deadweight; it makes stringent regulation harder to circumvent, with a measurable cost only when it is not accompanied by commensurate stringency.

The remainder of this section reports the cell-by-cell evidence in detail; Table 4 reports average treatment-on-treated effects across the eight cells, and Figures 2–6 plot the underlying horizon-by-horizon impulse responses.

Cell	log SRISK	log Total	log Bank	log Bk/Tot	Priv c/Y
ATT window	**h = 1, ..., 8**	**h = 1, ..., 8**	**h = 1, ..., 8**	**h = 1, ..., 8**	**h = 4, ..., 12**
Tighten					
High S, High C	-0.406**	0.146	-0.053	-0.070	-11.169
	[-0.677, -0.116]	[-0.083, 0.315]	[-0.242, 0.096]	[-0.170, 0.066]	[-21.004, 2.000]
High S, Low C	-0.191	0.183	0.155	-0.073	-6.561
	[-0.437, 0.112]	[-0.010, 0.330]	[-0.058, 0.316]	[-0.142, 0.006]	[-16.443, 6.518]
Low S, High C	-0.141	0.144**	-0.026	-0.236***	-6.232
	[-0.339, 0.061]	[0.015, 0.239]	[-0.262, 0.149]	[-0.356, -0.072]	[-15.705, 6.932]
Low S, Low C	0.017	0.049	0.288**	-0.001	-1.833
	[-0.193, 0.220]	[-0.130, 0.172]	[0.086, 0.433]	[-0.143, 0.208]	[-11.220, 11.763]
Loosen					
High S, High C	0.720	0.003	-0.081	-0.012	1.651
	[-0.182, 1.516]	[-0.175, 0.172]	[-0.234, 0.078]	[-0.157, 0.158]	[-11.198, 14.211]
High S, Low C	1.147**	0.629**	-0.066	-0.199*	-4.010
	[0.367, 1.887]	[0.216, 0.946]	[-0.201, 0.087]	[-0.363, -0.018]	[-19.991, 9.009]
Low S, High C	-0.273	0.224	0.162	-0.137	-9.570
	[-0.601, 0.080]	[-0.002, 0.431]	[-0.084, 0.383]	[-0.328, 0.081]	[-24.775, 7.501]
Low S, Low C	-0.115	0.143	0.237**	-0.063	-2.738
	[-0.302, 0.132]	[-0.060, 0.255]	[0.045, 0.344]	[-0.166, 0.121]	[-13.143, 16.013]
N	7,296	9,931	8,715	8,715	11,123

Table 4. Average treatment-on-treated effects from the LP-IFE design. For each outcome and each of the eight cells, the entry is the simple average of the horizon-specific coefficients over the post-event quarters listed in the “ATT window” row. The window is h = 1–8 for the forward-looking outcomes (SRISK, equity) and h = 4–12 for the credit aggregates, reflecting the slower transmission of regulation through bank balance sheets. Significance stars and the 90% confidence interval in brackets below each coefficient are from a within-quarter full-reshuffle randomization test (B = 2,000): within each calendar quarter the treated-cell label is reassigned at random across the whole panel and the windowed coefficient is recomputed against the fixed clean-control reference, so the test asks whether the cell as measured does more than random labeling. The point estimates are the LP-IFE coefficients and are unchanged by the choice of inference. * p < 0.10, ** p < 0.05, *** p < 0.01.

6.1 Effects on systemic risk

Column 1 of Table 4 evaluates the response of log SRISK. Tightening events that increase both stringency and complexity (the High S, High C cell) reduce systemic risk by an average of approximately 41 log points over the eight quarters following the event (coefficient of -0.406, within-quarter randomization-inference p = 0.009), equivalent to a sustained reduction of roughly one-third in the level of country-level SRISK relative to the no-treatment counterfactual. Tightenings that raise stringency but not complexity produce a smaller, statistically insignificant decline of about 19 log

points; the two low-stringency cells show no meaningful response. The contrast between the High S, High C and High S, Low C cells suggests that complexity is not unnecessary regulatory baggage: it appears to make stringent regulation more effective at containing systemic risk. The temporary decrease in systemic risk evident from Figure 2 explains the statistically insignificant effect in Table 4.

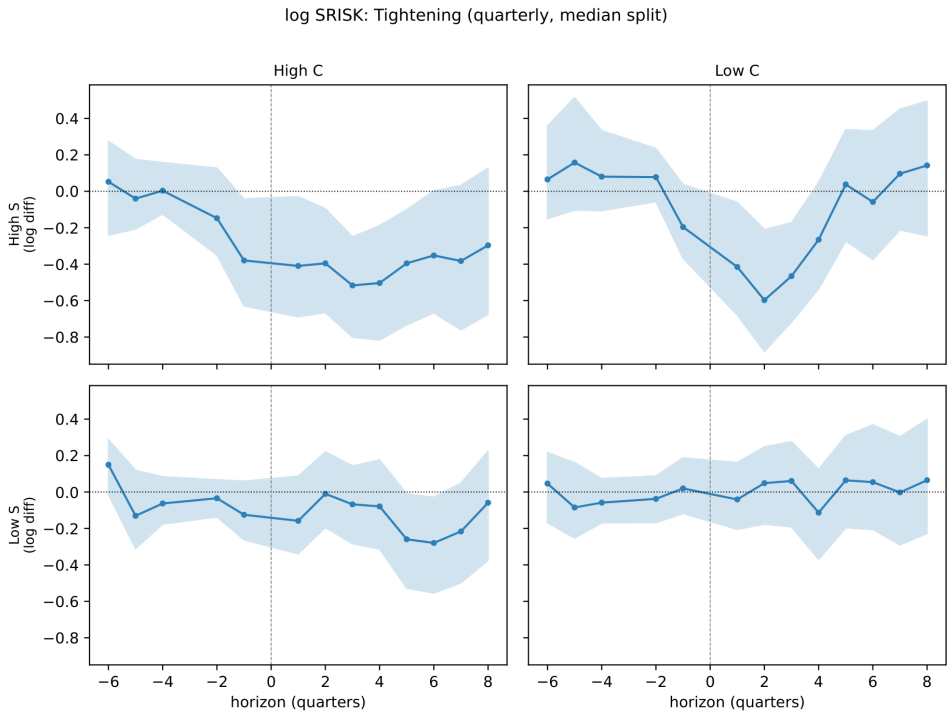

Figure 2. Effect of tightening events on log SRISK by stringency-complexity cell. Each subpanel plots the per-horizon cell coefficient with its pointwise 90% within-quarter full-reshuffle randomization band, under the LP-IFE design described in Section 4; rows are stringency (high / low) and columns are complexity (high / low).

Figure 2 plots the per-horizon impulse responses behind column 1 of Table 4. The reduction in SRISK associated with High S, High C tightenings is sustained across the full eight-quarter window. Tightenings with high stringency but low complexity, in contrast, produce a smaller and shorter-lived reduction. We interpret the differential persistence as consistent with regulatory arbitrage: when a stringent rule is operationally simple, banks can take on equivalent risk in adjacent activities or vehicles while remaining technically compliant, and the systemic-risk gain decays; when the rule is operationally demanding, such substitution is harder, and the gain is preserved.

6.2 Effects on stock-market valuations

Columns 2 to 4 of Table 4 evaluate equity-market responses. Column 2 reports the response of the log of total country-level market capitalisation, that is, stock market returns at different maturities. Across all four tightening cells, the total market is positive in the eight quarters following a tightening event, by between roughly 5 and

18 log points; under within-quarter randomization inference none of the four cells is individually significant, with the High S, Low C cell closest (p = 0.10). The positive sign across every cell is the relevant pattern. Tightening events appear to occur in the recovery phase of the cycle, and the regulatory regime, however stringent or complex, does not appear to impede the recovery.

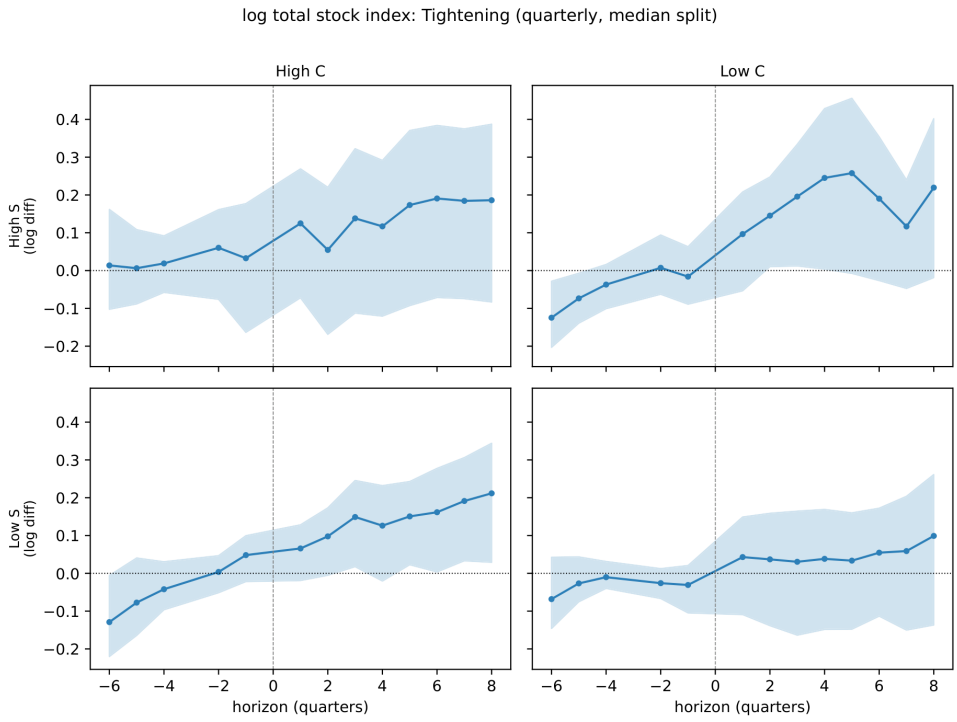

Figure 3. Effect of tightening events on log total stock-market capitalisation by stringency-complexity cell. See Figure 2 for layout and inference details.

Figure 3 confirms this pattern visually. Trajectories across the four tightening cells are broadly positive, including in the High S, High C cell. Whatever costs complexity imposes on regulated banks, they do not show up at the level of aggregate equity values.

Columns 3 and 4 turn to the bank sector. Column 3 reports the change in the level of bank-sector capitalisation; column 4 reports the change in the log ratio of bank to total capitalisation, which strips out the common market-wide component. The bank-to-total ratio falls in the Low S, High C cell by about 24 log points (-0.236, within-quarter randomization-inference p = 0.072), equivalent to a reduction of roughly one-fifth in the bank-to-total equity-index ratio, by far the largest tightening-cell underperformance. A much smaller decline in the High S, Low C cell (-0.073) also clears the 90% interval but is an order of magnitude smaller and economically minor. The High S, High C cell shows no significant bank underperformance. Stringent and complex regulation does not destroy bank franchise value relative to the rest of the market; it is lax-but-complex regulation that does so. This pattern echoes the theoretical mechanism in Asriyan, Foarta, and Vanasco (2023): complexity is harmful when it is not paired with commensurate stringency, because in that case it imposes operational burden without disciplining the activities that generate the underlying risk.

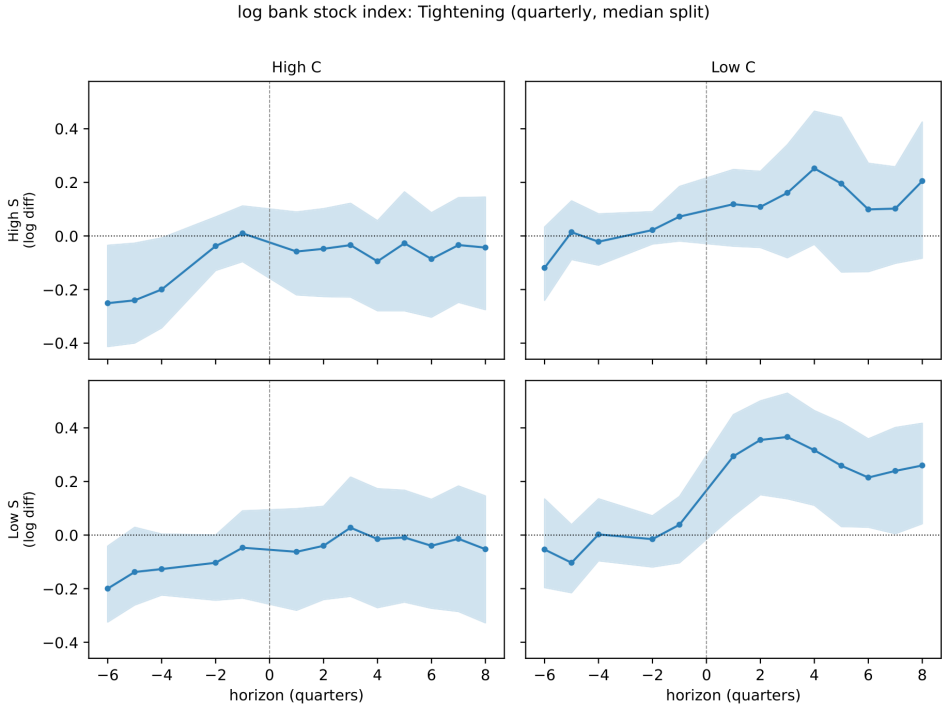

Figure 4. Effect of tightening events on log bank-sector capitalisation by stringency-complexity cell. See Figure 2 for layout and inference details.

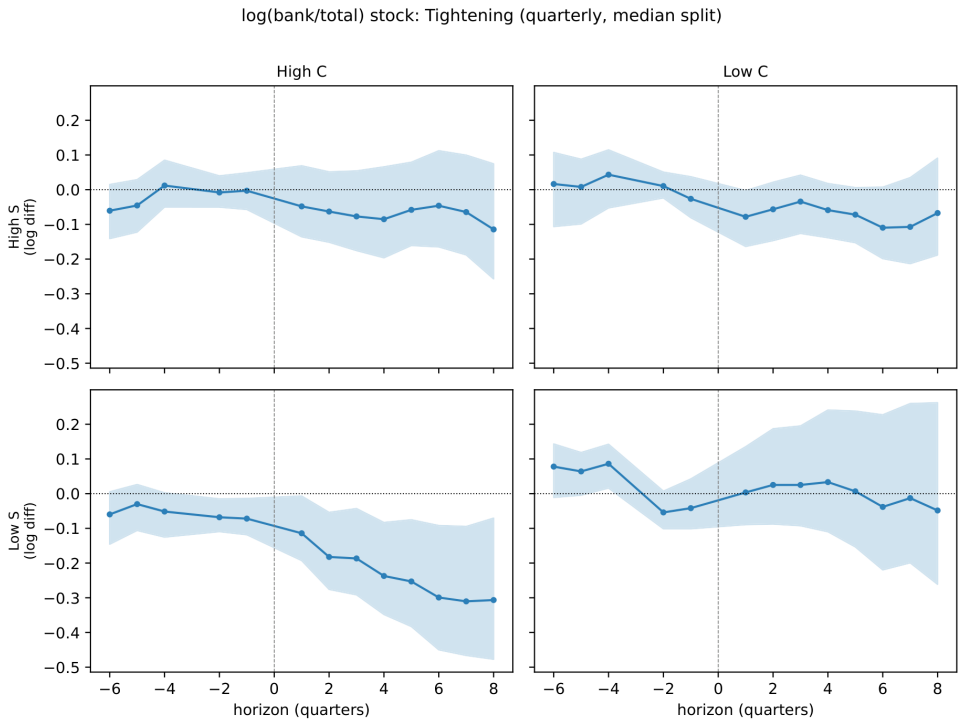

Figure 5. Effect of tightening events on the log ratio of bank to total stock-market capitalisation by stringency-complexity cell. See Figure 2 for layout and inference details.

Figures 4 and 5 plot the bank-stock and bank-to-total responses by horizon. The largest bank-to-total underperformance occurs in the Low S, High C cell, precisely where complexity is high and stringency is low. Complexity can be harmful for the value of banks subject to regulation, but only when it is not accompanied by commensurate stringency.

6.3 Effects on private and bank credit

Column 5 of Table 4 evaluates private credit to GDP at horizons of one to three years. All four tightening cells are mildly negative in point estimate (between -11.17 and -1.83 percentage points of GDP), but none is statistically distinguishable from zero; the High S, High C estimate of -11.17 pp is economically substantial in isolation but its within-quarter randomization-inference interval spans roughly -17 to +12 pp. We therefore read the credit column as a null — the data do not detect an aggregate credit-to-GDP response to tightenings — and as a robustness check on the SRISK and equity headlines rather than as a free-standing finding.

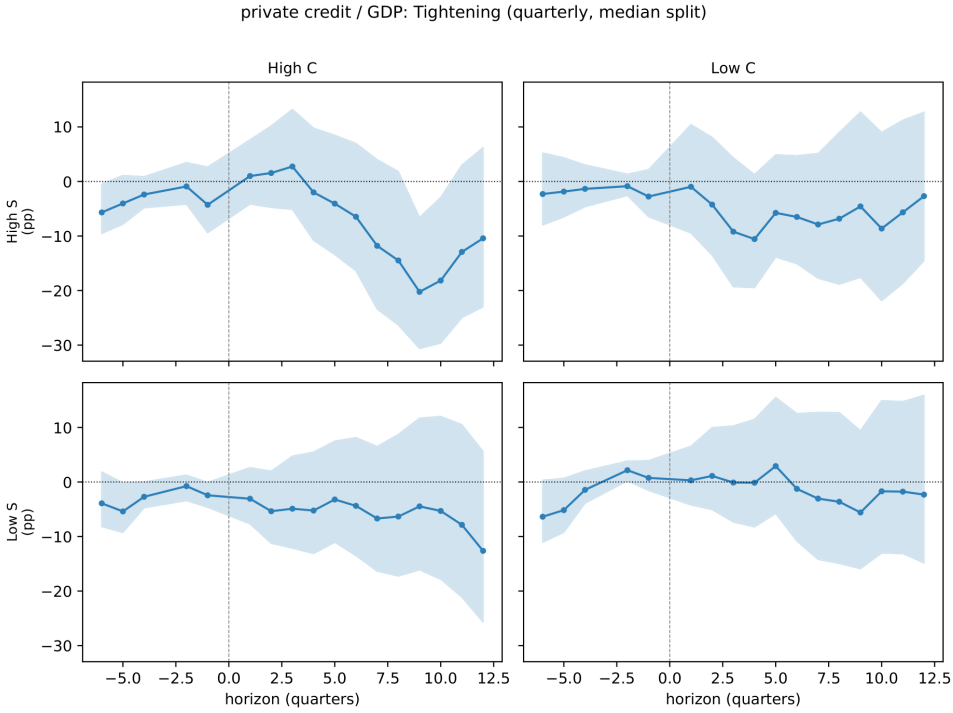

Figure 6. Effect of tightening events on private credit to GDP by stringency-complexity cell. See Figure 2 for layout and inference details.

Figure 6 plots the horizon-specific private-credit responses; the four tightening trajectories sit close to zero across the full window. One natural reading of the all-sectors credit measure is that it could mask a composition change: a tightening that contracts bank lending while reallocating credit to non-bank intermediaries would leave the all-sectors aggregate unchanged.

6.4 Loosening events

We do not draw strong inferences from the loosening panel of Table 4. The estimates are noisy, and the treated samples are thin: many loosening events are concentrated in crisis quarters when other variables are moving, and only a small fraction of loosenings register meaningful complexity changes (the loosening complexity median is close to zero). We note, however, that loosening events in the high-stringency, low-complexity cell increase systemic risk, consistent with what we find for the strengthening events.

Overall, the evidence in Table 4 and Figures 2–6 is consistent with the theoretical mechanism of Asriyan, Foarta, and Vanasco (2023): complexity is neither uniformly harmful nor uniformly beneficial. When paired with stringency, it is associated with a larger and more persistent reduction in systemic risk, without imposing detectable costs on bank or aggregate equity values. When stringency is low, complexity becomes a cost that depresses bank franchise value without a commensurate

stability gain. These results provide preliminary evidence that the post-crisis rulebook's complexity is not a pure deadweight loss.

7 Conclusion

We have provided preliminary evidence on whether the post-crisis tightening of bank regulation has reduced systemic risk, and on the role of regulatory complexity in that process. Our central conceptual point is that stringency and complexity are distinct attributes of a regulatory regime, and that the empirical question is not how much regulation a country has but how stringent and how complex it is. We construct two indices that separate these dimensions through pairwise rankings of every dated iMaPP event in 21 advanced economies between 1995Q1 and 2025Q3, and we read each index alongside outcomes for systemic risk, equity-market valuations, and private credit.

Three findings emerge. First, tightening events that increase both stringency and complexity reduce systemic risk by approximately 41 log points over two years — roughly a one-third reduction in the level — while tightenings that raise stringency without raising complexity produce only a smaller, transient reduction in SRISK. Complexity is, on this evidence, not unnecessary regulatory baggage. It appears to make stringent regulation harder to circumvent through regulatory arbitrage. Second, neither aggregate equity values nor private credit decline after tightenings, including those that combine high stringency with high complexity. Whatever costs the post-crisis rulebook has imposed on regulated banks have not translated into a contraction of aggregate stock-market value or of credit to the private sector. Third, complexity is harmful when it is not paired with stringency: bank stocks underperform the broad market in the cell that combines low stringency with high complexity, consistent with the theoretical mechanism in Asriyan, Foarta, and Vanasco (2023) under which complexity becomes a cost without a corresponding stability gain.

These results have implications for the deregulatory initiatives currently under way across major advanced economies. We caution that the implications we draw are preliminary: our identification rests on regulatory events through 2025, and the post-2024 actions have only begun to take effect, so the inferences below are necessarily qualitative. The United States has lowered stringency without commensurate complexity reduction: the November 2025 cut to the enhanced supplementary leverage ratio, the dilution of the Basel III endgame, and the broader rollback of post-crisis supervision move the U.S. regime away from the High S, High C combination that, on our estimates, produced the most persistent reduction in systemic risk. The United Kingdom is moving in the same direction more slowly, with a capital-neutral transposition of Basel 3.1 and an active review of the ring-fencing regime. In the European Union, by contrast, CRR3 and CRD VI raise the binding capital requirements of the largest banks while a separate simplification agenda — pursued through the Omnibus packages, the Commission's 2026 banking-competitiveness consultation, and the ECB-led High-Level Task Force on Simplification — seeks to lower complexity that is not paired with stringency. Read through the lens of our estimates, this two-track European approach moves in a defensible direction: raising

stringency on the dimension that contains systemic risk, while removing the kind of procedural depth that, in our Low S, High C cell, imposes a cost on bank valuations without a corresponding stability gain. The risk is that the simplification effort does not distinguish complexity that supports binding constraints from complexity that is mere residue; stripping the former would erode the systemic-risk gains that the capital-side tightening is intended to deliver. Switzerland's response to the Credit Suisse collapse — the Federal Council's proposed additional CET1 surcharge on UBS — sits in the High S, High C cell that our estimates associate with the most durable systemic-risk reduction, and the absence of an offsetting effect on bank valuations in that cell is consistent with our findings. Our evidence accordingly suggests that the U.S. rollback risks going too far where stringency is being lowered, that the U.K. trajectory inherits the same risk on a slower timetable, that the EU agenda is broadly well-targeted provided it preserves complexity that is load-bearing, and that the Swiss tightening, far from being excessive on financial-stability grounds, is the kind of policy our paper would predict to deliver a persistent reduction in systemic risk. The challenge throughout is not to minimise complexity per se but to ensure that operational depth is matched by bindingness; a less-complex but equally stringent regulation may be more easily circumvented.

Two limitations shape the agenda for future work. SRISK covers all publicly listed financial institutions — banks and non-bank financial firms such as insurers, broker-dealers, and listed asset managers — but not unlisted vehicles. The migration of credit and risk-taking to private credit funds, private-equity-backed direct lenders, securitisation special-purpose vehicles, and other unlisted intermediaries is a parallel margin that our outcomes measure imperfectly: these vehicles fall outside the listed-financials perimeter that V-Lab's SRISK panel observes, and the BIS all-sectors credit measure aggregates them with bank lending without identifying the channel separately. Whether systemic risk is genuinely lower after a tightening, or simply rebuilds in vehicles outside our measurement perimeter, is a question that would benefit from a complementary measurement effort on the unlisted-vehicle side. In addition, SRISK is available only from the early-to-mid 2000s in most countries, so the systemic-risk regressions identify the effects only from post-Basel-II events and may not generalise to earlier regulatory cycles.

References

Acharya, V. V., Engle, R. and Richardson, M. (2012), "Capital shortfall: A new approach to ranking and regulating systemic risks", American Economic Review, Vol. 102, No. 3, pp. 59-64.

Acharya, V. V., Pedersen, L. H., Philippon, T. and Richardson, M. (2017), "Measuring systemic risk", Review of Financial Studies, Vol. 30, No. 1, pp. 2-47.

Acharya, V.V., Bergant, K., Crosignani, M., Eisert, T. and McCann, F. (2022), "The Anatomy of the Transmission of Macroprudential Policies", The Journal of Finance, Vol. 77, No. 5, pp. 2533–2575.

Acharya, V. V., Brunnermeier, M. K. and Pierret, D. (2024), “Systemic Risk Measures: Taking Stock from 1927 to 2023”, *Annual Review of Financial Economics*, Vol. 17, forthcoming.

Acharya, V.V., Cetorelli, N. and Tuckman, B. (2024), “Where Do Banks End and NBFIs Begin?”, *NBER Working Paper*, No. 32316.

Adrian, T. and Brunnermeier, M. K. (2016), “CoVaR”, *American Economic Review*, Vol. 106, No. 7, pp. 1705–1741.

Alam, Z., Alter, A., Eiseman, J., Gelos, G., Kang, H., Narita, M., Nier, E. and Wang, N. (2019), "Digging Deeper — Evidence on the Effects of Macroprudential Policies from a New Database", IMF Working Paper, No. 19/66.

Almasi, P., Dagher, J. and Prato, C. (2025), “Financial Regulatory Cycles: A Political Economy Model”, *Journal of Financial Intermediation*, Vol. 63, Article 101164.

Asirvatham, A., Mokshi, V. and Shleifer, A. (2026), "GPT as a Measurement Tool", NBER Working Paper, No. 34834.

Asriyan, V., Foarta, D. and Vanasco, V. (2023), “The Good, the Bad, and the Complex: Product Design with Imperfect Information”, *American Economic Journal: Microeconomics*, Vol. 15, No. 2, pp. 187–226.

Bai, J. (2009), “Panel Data Models with Interactive Fixed Effects”, *Econometrica*, Vol. 77, No. 4, pp. 1229–1279.

Brownlees, C. and Engle, R. F. (2017), “SRISK: A Conditional Capital Shortfall Measure of Systemic Risk”, *The Review of Financial Studies*, Vol. 30, No. 1, pp. 48–79.

Budnik, K. and Kleibl, J. (2018), "Macroprudential Regulation in the European Union in 1995–2014: Introducing a New Data Set on Policy Actions of a Macroprudential Nature", ECB Working Paper Series, No. 2123.

Bruno, Valentina, Ilhyock Shim, and Hyun Song Shin (2017). "Comparative Assessment of Macroprudential Policies." Journal of Financial Stability, 28, 183–202.

Cerutti, E., Claessens, S. and Laeven, L. (2017), "The use and effectiveness of macroprudential policies: New evidence", Journal of Financial Stability, Vol. 28, pp. 203–224.

Cetorelli, N. and Kundu, S. (2026), “Regulatory Arbitrage Within the Firm”, *Federal Reserve Bank of New York Staff Reports*, No. 1196.

Colliard, J.-E. (2025), “Measuring Regulatory Complexity”, *Journal of Financial Economics*, forthcoming.

Chernenko, S., Erel, I. and Prilmeier, R. (2022), “Why Do Firms Borrow Directly from Nonbanks?”, *The Review of Financial Studies*, Vol. 35, No. 11, pp. 4902–4947.

Dagher, J. (2018), “Regulatory Cycles: Revisiting the Political Economy of Financial Crises”, *IMF Working Paper*, No. 18/8.

Davis, S. J. (2017), “Regulatory Complexity and Policy Uncertainty: Headwinds of Our Own Making”, *SSRN Electronic Journal*, doi:10.2139/ssrn.2723980.

Dube, A., Girardi, D., Jordà, Ò. and Taylor, A. M. (2025), “A Local Projections Approach to Difference-in-Differences”, *Journal of Applied Econometrics*, Vol. 40, No. 7, pp. 741–758.

Engle, R., Jondeau, E. and Rockinger, M. (2015), "Systemic risk in Europe", *Review of Finance*, Vol. 19, No. 1, pp. 145-190.

Giannetti, M., Jotikasthira, C., Rapp, A.C. and Waibel, M. (2024), “Intermediary Balance Sheet Constraints, Bond Mutual Funds’ Strategies, and Bond Returns”, *Swedish House of Finance Research Paper*, No. 23-10.

Gratton, G., Guiso, L., Michelacci, C. and Morelli, M. (2021), “From Weber to Kafka: Political Instability and the Overproduction of Laws”, *American Economic Review*, Vol. 111, No. 7, pp. 2964–3003.

Greenwood, R., Hanson, S. G., Stein, J. C. and Sunderam, A. (2017), “Strengthening and Streamlining Bank Capital Regulation”, *Brookings Papers on Economic Activity*, Fall 2017, pp. 479–565.

Kalmenovitz, J., Lowry, M. and Volkova, E. (2025), “Regulatory Fragmentation”, *The Journal of Finance*, Vol. 80, No. 2, pp. 1081–1126.

Müller, K. (2023), “Electoral Cycles in Macroprudential Regulation”, *American Economic Journal: Economic Policy*, Vol. 15, No. 4, pp. 295–322.

Müller, K., Xu, C., Lehbib, M. and Chen, Z. (2025), “The Global Macro Database: A New International Macroeconomic Dataset”, *NBER Working Paper*, No. 33714.

Xu, Y. (2017), “Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models”, *Political Analysis*, Vol. 25, No. 1, pp. 57–76.