From Solvency Selection to Too-Big-to-Break: Architecture, Events and Reversibility in Capital Centralisation

From Solvency Selection to *Too-Big-to-Break*

Architecture, Events and Reversibility in Capital Centralisation

Nicolò Bellanca

Department of Economics and Management, University of Florence

SSRN Working Paper, July 2026

Abstract

Research on capital centralisation has converted a classical proposition of political economy into an empirical programme: monetary and financial conditions select capitals asymmetrically, shape acquisition capacity and target vulnerability, and thereby reorganise ownership and control. This working paper synthesises three recent developments. The first retains the scalar concentration statistic but treats the control operator as the primary measurement object, making the statistic one projection of the structure of command. The second classifies the primitive events that transform that structure - strong merger, accretive acquisition, transfer of an existing stake, pure divestiture and legal split - and derives testable signatures for four frontiers. The third studies the formation, persistence and reversibility of control blocks through a coagulation-fragmentation process. The synthesis separates concentration from centralisation, formation from persistence, and legal break-up from restored autonomy. It yields an integrated empirical programme: replicate the established measure; reconstruct the control operator; classify events; estimate acquisition intensity and size selectivity together with fragmentation intensity and scalability; and follow entities after break-up. Monetary policy, competition policy, resolution and corporate law then jointly shape the extent, form and reversibility of capital centralisation.

Keywords: capital centralisation; solvency rule; corporate control; networks; mergers and acquisitions; reversibility; antitrust; too-big-to-break.

JEL codes: B51; D85; E52; G34; L41; P16.

1. Introduction

Capital centralisation is often described by market shares, the size of the largest firms, or the minimum number of actors required to control a given fraction of economic value. These measures make an otherwise general proposition observable, but a terminal statistic cannot identify the structure of command that generated it, the event that changed it, or the extent to which the change can be reversed.

The difficulty can be formulated through three questions. First, is an economy in which a few actors occupy the leading positions governed by a single pole, several autonomous poles, a pyramid, or a chain of directional dependencies? Second, does the observed change result from a legal merger, the acquisition of previously dispersed control, the transfer of an existing stake, or the removal of links? Third, does a subsequent split restore autonomous entities, or does it merely change their legal identities while coordination, dependence and the capacity to recombine remain intact?

The three questions belong to distinct but cumulative levels. The first concerns the control structure at a point in time. The second concerns the events that alter masses, links and organisational genealogy. The third concerns the persistence and reversibility of the resulting configuration. The sequence is:

monetary and financial conditions → event selection → control structure → persistence or dissolution.

This sequence extends the research programme associated with the solvency rule. In that programme, monetary policy acts through aggregate demand and through the hierarchy of solvency. It favours some capitals over others and affects which firms preserve autonomy, finance acquisitions, or become targets (Brancaccio and Fontana 2013, 2016; Brancaccio et al. 2018, 2019, 2023). Centralisation is therefore a selective outcome rather than a simple by-product of transaction volume.

The three companion working papers develop complementary steps. *From Concentration to Control: Measuring the Architecture of Capital Centralisation* studies measurement at a fixed date and proves that the same scalar statistic is compatible with radically different structures of command (Bellanca 2026a). *When Control Structures Coagulate: Event Signatures and the Four Frontiers of Capital Centralisation* introduces an event taxonomy and distinguishes legal mass, reciprocal integration, hierarchical reach and spectral feedback (Bellanca 2026b). *The Reversibility Frontier of Capital Centralisation: Credit Selectivity, Fragmentation, and Too-Big-to-Break* projects the economy onto control blocks and studies formation, stationary centralisation, finite-horizon memory and local post-break-up races (Bellanca 2026c). Complete proofs remain in the companion papers. The present working paper explains their economic content through elementary algebra, numerical examples and a common empirical design.

The central claim is that selective financial conditions determine who can acquire whom; the resulting events transform several dimensions of command; and their persistence depends on the institutional capacity to fragment both legal masses and the organisational and financial links that hold them together.

This formulation introduces three distinctions. Concentration concerns the distribution of magnitudes; centralisation concerns the organisation of command over those magnitudes. The formation of a large block and its persistence after intervention are different problems. Finally, the legal split of an entity is not the same as the restoration of the economic autonomy of its parts.

The working paper proceeds as follows. Section 2 places the proposal within the solvency-rule programme. Section 3 introduces the control operator and a first numerical example. Section 4 distinguishes events and four frontiers. Section 5 explains why merger and split are not inverse operations. Section 6 presents the dynamics of formation and reversibility. Section 7 links the three levels. Section 8 develops the empirical programme. Section 9 discusses policy implications. Section 10 sets out limits and open questions.

2. From the solvency rule to a multi-layer theory of centralisation

Marx distinguished concentration, the growth of individual capitals through accumulation, from centralisation, the redistribution of already existing capital through merger, acquisition and the reorganisation of command (Marx 1976 [1867]). Credit accelerates the second process because it enables some firms to appropriate assets that others can no longer finance.

The solvency rule translates this intuition into a monetary theory of selection. In a heterogeneous economy, the same change in interest rates, credit constraints or refinancing conditions affects firms differently. Solvency depends on financing costs, profitability, growth and indebtedness. A tightening may reduce aggregate investment and transfer assets or control towards actors with greater liquidity, better access to credit or stronger financial relationships. Evidence on creditor control of corporate acquisitions confirms that financing affects transaction identity as well as volume (Becher, Griffin and Nini 2022). Conversely, support measures may preserve formal firms while protecting already centralised control structures.

The programme formulated solvency selection as an alternative to a uniform representation of monetary policy, reconstructed ownership and control networks to measure global centralisation, and used dynamic models to study the relations among solvency, monetary policy, output and centralisation (Brancaccio et al. 2018, 2019, 2023). This work also connects to the literature on global corporate-control networks (Glattfelder and Battiston 2009; Vitali, Glattfelder and Battiston 2011). The next step sharpens the mechanism without replacing those findings.

A simple representation of the intensity with which actor *i* acquires control of actor *j* is:

$$\lambda_{i\leftarrow j,t} = \kappa_t \times A_{i,t} \times V_{j,t}.$$

The three factors have an immediate interpretation. κ_t is the general market clock: liquidity, aggregate activity and common financing conditions. $A_{i,t}$ is the acquirer's capacity: size, solvency, access to credit, network position and institutional support. $V_{j,t}$ is the target's vulnerability: financial distress, dispersed control, weak owners, sectoral exposure or regulatory constraints.

The formula does not claim that every acquisition can be reduced to three numbers. Its purpose is to distinguish the general volume of opportunities from their selectivity. Two periods can display the same number of transactions and produce opposite structural effects. A period with many deals distributed among similarly sized actors may centralise less than a period with a few deals almost entirely undertaken by very large buyers. Merger volume therefore does not identify the degree of centralisation produced.

The multi-layer theory adds three objects to this formula. First, a **structural state**: masses, control links and the genealogy of blocks. Second, a taxonomy of **events** that transform the state. Third, a set of **fragmentation hazards** describing the capacity to dismantle blocks both legally and organisationally. Observed centralisation is the joint outcome of these levels.

The proposal retains solvency selection as its economic core while changing the level of observation. It examines the control structure behind a concentration series, the microeconomic events behind aggregate correlations, and the conditions governing the reversal of formed blocks.

3. First layer: the architecture of command

3.1 Why size is not enough

Consider four situations. In the first, a large holding company directly controls many firms: there is a dominant pole. In the second, there are two large groups that do not depend on one another: there are autonomous poles. In the third, an upstream group controls an intermediate group, which in turn controls other firms: there is a pyramid. Divergences among legal identity, ultimate ownership and effective control are well established in the literature on business groups and corporate pyramids (Brioschi, Buzzacchi and Colombo 1989; La Porta, Lopez-de-Silanes and Shleifer 1999). In the fourth, the network is a long directional chain: effects can propagate and be temporarily amplified even without a large reciprocal core.

The four economies may require the same number of controllers to reach 80 per cent of total value while distributing power differently. A dominant pole concentrates dependence in one control core; autonomous poles produce an oligopolar but ununified system; a pyramid multiplies upstream control over the base; and a directional chain can accumulate transmission risk that size-based statistics do not reveal.

To make these differences observable, the first working paper represents direct control by a non-negative matrix C. Entry C_{ij} measures the control attributed from actor i (row) to actor j (column) under an explicitly stated rule based on ownership shares, voting rights, relative majority or another institutionally grounded criterion. No single rule is assumed to be natural. Direct attribution and network propagation are separate measurement problems (Bellanca 2026a).

Let $v \geq 0$ be the vector of firms' stand-alone values and assume $\rho(C) < 1$. Integrated value, meaning stand-alone value plus what is transmitted along all control paths, is:

$$\boldsymbol{x = (I - C)^{-1}v = v + Cv + C^2v + C^3v + \ldots}$$

The expression is a path ledger: v records stand-alone value, Cv one-step transmission, C^2v two-step transmission, and later terms longer chains. The assumption $\rho(C) < 1$ is necessary and sufficient for the Neumann series to converge in finite dimension and makes $I - C$ invertible.

3.2 A three-firm example

Consider firms A, B and C with stand-alone values:

$$\boldsymbol{v = (100, 80, 60).}$$

A controls 40 per cent of B, while B controls 50 per cent of C. There are no other links. The first round transmits 32 units associated with B to A and 30 units associated with C to B:

$$\boldsymbol{Cv = (32, 30, 0).}$$

The second round transmits part of C's value to A through B:

$$\boldsymbol{C^2v = (12, 0, 0),}$$

because 0.40 × 0.50 × 60 = 12. The chain then ends, and integrated value is:

$$\boldsymbol{x = (144, 110, 60).}$$

Without control links, the first two actors hold 180 out of 240, or 75 per cent of value, so all three are needed to reach the 80 per cent threshold. With the control network, the first two account for 254 out of 314, approximately 80.9 per cent, and two actors suffice. Stand-alone values are unchanged; the control structure has changed the attribution of integrated value.

The example also shows why a threshold statistic is informative but incomplete. It signals that the final ranking has become more concentrated. It does not reveal whether the movement was generated by a cycle, a pyramid, a crucial single edge or a fragile chain. Moreover, the threshold is exceeded by only 0.9 percentage points. A small data perturbation could move the required number of controllers back from two to three without a major structural change. Threshold margins above and below the crossing should therefore also be reported.

3.3 A profile, not a super-index

A compact diagnostic profile supplements the threshold measure with:

1. **amplification:** how much indirect paths increase control;
2. **localisation:** whether integrated control is concentrated on a few actors or remains diffuse;
3. **autonomous poles:** how many reciprocally connected cores remain mutually independent;
4. **hierarchy:** how much mass lies downstream from a control core;
5. **transient stress:** whether a directional network can amplify disturbances for several periods even while remaining far from asymptotic criticality, as suggested by the theory of non-normal matrices and pseudospectra (Trefethen and Embree 2005).

The first working paper establishes non-identification: the same threshold statistic can coexist with different amplification strengths, numbers of poles, degrees of localisation and hierarchical structures. Any empirical correlation among these properties is institutional rather than mathematically necessary.

This step makes the classical distinction between concentration and centralisation operational. A change in stand-alone values **v** is growth or concentration of units. A change in the operator **C** is a reorganisation of command. The observed vector **x** combines both. Two counterfactuals can separate them: change C while holding initial values fixed, and change v while holding the initial network fixed. If the movement in the threshold statistic is reproduced by changes in v alone, evidence of centralisation in the stronger sense is weak. If it requires a change in C, there is evidence of a transformation of command.

4. Second layer: events and the four frontiers

4.1 Operations that conventional statistics conflate

Between two dates a network may change for different reasons. The second working paper distinguishes five primitive events (Bellanca 2026b).

A **strong merger** replaces two legal centres with one. Their masses are added and the network is compressed. An **accretive acquisition** assigns to a buyer some control that was previously dispersed or unassigned, while the target remains legally distinct. A **transfer** moves an existing stake from a seller to a buyer: total control over the target may remain unchanged, but its holder changes. A **pure divestiture** removes one or more links without reallocating them. A **legal split** divides one block into descendants and must determine how internal and external links are redistributed.

These operations are not scalar variants of one phenomenon. A merger changes the number of legal entities; an accretive acquisition adds control; a transfer redistributes it; a divestiture removes it; and a split enlarges the organisational state space without determining a unique network outcome.

4.2 Four frontiers of centralisation

To describe event effects, define four intuitive frontiers:

- **B — legal mass:** the share of the system held by the largest legal block;
- **I — reciprocal integration:** the share of total mass contained in the largest non-trivial strongly connected component, with I = 0 if no such component exists;
- **H — hierarchical reach:** the largest share of total mass lying strictly downstream from a strongly connected component in the condensation graph, excluding the source component itself;
- **R — spectral feedback:** $R = \rho(C)$, the spectral radius of the non-negative control operator. It measures asymptotic feedback through cycles and is distinct from finite-horizon transient stress.

For an isolated primitive event at fixed total mass, the definitions imply the weak sign restrictions below. Accretive acquisition adds a non-negative edge, so strongly connected components cannot split and Perron monotonicity prevents R from falling. Pure divestiture removes edges, so components can only refine and R cannot rise. Hierarchical reach remains ambiguous because adding an edge can close a cycle, while removing one can open a cycle into a chain.

Table 1. Weak event-frontier signatures

Event	B	I	H	R
Strong merger	weakly ↑	ambiguous	ambiguous	ambiguous
Accretive acquisition	unchanged	weakly ↑	ambiguous	weakly ↑
Transfer	unchanged	ambiguous	ambiguous	ambiguous
Pure divestiture	unchanged	weakly ↓	ambiguous	weakly ↓
Legal split	weakly ↓	ambiguous	ambiguous	ambiguous

These weak sign restrictions apply to an isolated, correctly measured primitive event at fixed total mass. A fall in R after an event classified as an accretive acquisition therefore signals misclassification, a change in the control-attribution rule, measurement error or simultaneous restructuring. The theory leaves some cells genuinely ambiguous but rules out the monotonicity violations shown in the table.

4.3 A four-firm example

Consider blocks A, B, C and D with masses (40, 30, 20, 10). Rows are controllers and columns are targets. The initial operator has $A \to B = 0.40$, $A \to D = 0.20$, $B \to C = 0.50$ and $C \to A = 0.30$; all other entries are zero. Thus A, B and C form a control cycle and D lies downstream.

$$C = [0 \;\; 0.40 \;\; 0 \;\; 0.20; \;\; 0 \;\; 0 \;\; 0.50 \;\; 0; \;\; 0.30 \;\; 0 \;\; 0 \;\; 0; \;\; 0 \;\; 0 \;\; 0 \;\; 0].$$

Apply three alternatives to this state: (i) the mass-weighted strong merger of B and C; (ii) an accretive edge $D \to A$ of weight 0.35, equal to one half of A's unassigned control capacity; and (iii) a transfer of 0.10 control over A from C to D, which reduces $C \to A$ from 0.30 to 0.20 and adds $D \to A = 0.10$.

For the merger, $C^M = [0 \;\; 0.24 \;\; 0.20; \;\; 0.30 \;\; 0.20 \;\; 0; \;\; 0 \;\; 0 \;\; 0]$.

Table 2. Four-firm event comparison

State or event	B	I	H	R
Initial state	0.40	0.90	0.10	0.3915
Strong merger B + C	0.50	0.90	0.10	0.3864
Accretive acquisition $D \to A$	0.40	1.00	0.00	0.4507
Transfer $C \to D$ on target A	0.40	1.00	0.00	0.3615

The merger raises the largest legal mass but slightly reduces spectral feedback. The accretive acquisition and the transfer both place all actors in one cycle, so I rises to 1 and H falls to 0. The accretive event strengthens the cycle and raises R; the transfer weakens an existing edge and lowers R. Identical movements in two indicators can therefore conceal opposite mechanisms in a third.

The empirical record must distinguish a stake acquired from dispersed ownership from one acquired from an identified seller. The first operation adds net control; the second redistributes existing control. Pooling them erases one of the theory's clearest restrictions.

5. Compression, split and organisational memory

A merger does not merely combine masses. It also compresses represented relationships. Before merger, entities may have different creditors, boards, contracts, cross-holdings, information channels and production dependencies. After consolidation, many of those differences are represented as internal relations within one block or disappear from the observed data.

The second working paper uses the fixed mass-weighted compression $\mathrm{Comp}(Q, m; C) = Q^T C B_m$, where Q aggregates n pre-merger blocks into r blocks and $Q^T B_m = I_r$. Vectorisation gives $\mathrm{rank}[\mathrm{Comp}(Q, m; \cdot)] = r^2$ and nullity $n^2 - r^2$. A binary merger has $r = n - 1$ and therefore erases exactly $2n - 1$ linear directions. The result concerns this specified linear compression rather than every conceivable aggregation rule.

Two pre-merger networks with different cycles, hierarchies and spectral feedback can generate the same aggregate operator. Aggregation is therefore many-to-one, and no function of the consolidated operator alone can recover every admissible pre-merger network.

A split is an admissible reconstruction rather than the mechanical inverse of merger. Allocating assets does not settle how debts, contracts, stakes, strategic employees, data, suppliers and external relationships follow each descendant. Proportional splitting is one admissible lift among others; law, policy, creditors, management and organisational history select the outcome.

This observation distinguishes two institutional capacities:

- the **capacity to fragment legal mass**;
- the **capacity to separate control and dependence links**.

The first reduces B. The second affects I, H and R. One may succeed while the other fails. A group may be divided into several companies that continue to coordinate through ownership, debt, exclusive supply, shared infrastructure or overlapping management. Conversely, an intervention may dismantle control links without changing legal firm size.

Scalar reversibility and organisational reversibility therefore differ. The mass vector may return close to its initial distribution while the network does not return to its previous class of configurations. Memory need not reside in exceptional behaviour or psychological attachment. It may arise from the elementary fact that aggregation erased information and that later disaggregation must choose among many reconstructions.

6. Third layer: formation, persistence and the reversibility frontier

6.1 Coagulation and fragmentation

To study aggregate dynamics, the third working paper represents the economy as a distribution of control blocks (Bellanca 2026c). When two blocks merge, their masses are added: this is coagulation. When a block is effectively divided, its mass is partitioned: this is fragmentation.

The model distinguishes four features:

- **κ:** the general speed of acquisitions;
- **β:** the size selectivity of acquisitions;
- **φ:** the general intensity of effective fragmentation;
- **δ:** the scalability of fragmentation with block size.

To avoid a collision with $R = \rho(C)$, this synthesis writes δ for the fragmentation exponent denoted γ in Bellanca (2026c), and η for the ratio φ/κ denoted ρ there.

The distinction between κ and β is crucial. κ states how fast the acquisition clock runs. β states which pairs are relatively favoured. When β is high, large actors have a disproportionate probability of acquiring and participating in transactions that create still larger blocks. An economy may therefore display fewer transactions but stronger selection.

Similarly, φ and δ capture different aspects of deconcentration capacity. A high φ means that many fragmentation events occur relative to acquisitions. A high δ means that this capacity reaches large blocks at least as strongly as small ones. Counting remedies without estimating their size profile can be misleading: many divisions of small entities do not demonstrate the ability to dismantle a dominant group.

Define the ratio:

$$\boldsymbol{\eta = \varphi / \kappa.}$$

For fixed β, δ and split law, multiplying κ and φ by the same positive factor multiplies the generator and changes only physical time. The ratio $\eta = \varphi/\kappa$ governs the relative balance between acquisition and fragmentation. This clock-composition separation is testable and prevents transaction volume from being treated automatically as a change in long-run structure.

6.2 Formation does not imply irreversibility

In the pure product-kernel benchmark and under the stated solution concepts, $\beta = 1/2$ is the ordinary formation boundary: above it, size selectivity can generate a macroscopic block in the large-system limit. This is neither a universal policy target nor a result about the capacity to dismantle an existing block.

A model that permits mergers but forbids splits creates false mechanical irreversibility: once formed, the largest block cannot shrink because the process contains no reverse transitions. Serious analysis of reversibility must introduce bankruptcy separations, spin-offs, divestitures, resolution, dissolution of common control and other effective fragmentation events.

Assume that every pair of positive block sizes can merge and that every admissible binary split has positive conditional probability. At fixed total mass, the resulting continuous-time chain is finite and irreducible. It therefore has a unique invariant distribution, and its law converges to that distribution in total variation. If baseline parameters are permanently restored, every post-shock scalar state converges to the same invariant law. Permanent finite-system hysteresis is excluded, but economically relevant memory may outlast a regulatory mandate, a managerial generation or a political cycle. Four cases should be distinguished:

1. mechanical irreversibility, when reverse transitions are absent;
2. finite-horizon memory, when return exists but is slow;
3. metastability, when the system remains for a very long time in high- or low-centralisation regions;
4. structural irreversibility, when endogenous institutions or organisational information prevent the full state from returning.

6.3 Local races after break-up

Suppose a large block is divided into two parts. Immediately after intervention, three events may occur:

- the two fragments remerge;
- one fragment merges with a third block;
- one fragment is divided again.

The first outcome directly restores the previous identity. The second need not recreate the old company but may preserve macroscopic concentration through a new alliance or acquisition. The third continues dissolution.

The relative strength of direct repair is summarised by:

$$\chi^D = 2\beta - \delta - 1.$$

For two macroscopic fragments in the restricted first-event race, $\chi^D > 0$ means that the direct-remerger rate grows faster with system size than their combined renewed-fragmentation rate. The boundary is:

$$\delta = 2\beta - 1.$$

Attachment to external actors also depends on the rest of the economy. If the relevant background β-moment scales as N^s, the attachment-only exponent is:

$$\chi^T = \beta + s - 1 - \delta.$$

Writing χ^T for the third-party attachment channel, the aggregate local escape index - denoted $\chi^A(s)$ in Bellanca (2026c) - compares the strongest coagulative first-event channel with renewed fragmentation:

$$\chi^A(s) = \max\{2\beta - 1, \beta + s - 1\} - \delta = \max\{\chi^D, \chi^T\}.$$

In a diffuse environment, where a positive mass share remains distributed across small blocks, s equals 1 and the boundary becomes:

$$\boldsymbol{\delta = \beta}.$$

A numerical example makes the meaning transparent. Let $\beta = 0.8$.

- With $\delta = 0.4$, $2\beta - 1 = 0.6$ and $\delta < 0.6$: direct remerger dominates locally.
- With $\delta = 0.7$, direct remerger does not dominate, but $\delta < \beta$: attachment of one fragment to a third block tends to occur first.
- With $\delta = 0.9$, fragmentation scales more rapidly than both aggregative routes: dissolution has a local advantage.

The intermediate regime separates identity repair from continued dominance. Break-up may prevent restoration of the old entity while one fragment rapidly attaches to a third block. The large block changes form rather than disappearing, so evaluation must follow the fragments beyond the legal date of the split.

The channel exponents and the aggregate index are local first-event results: they compare which event tends to arrive first in a temporarily frozen post-split state. They do not, by themselves, prove global persistence, metastability or permanent hysteresis. The distinction is essential if a testable prediction is not to be inflated into a stronger claim than the theory warrants (Bellanca 2026c).

7. Continuity across the three levels

The three models cannot be superimposed without conditions. The operator model contains more information than the mass model. Moving from one to the other requires an empirical rule **P** that turns a control network into a partition of blocks under common command. The rule must state how it treats pyramids, cross-holdings, indirect control and legally distinct but coordinated entities.

Once blocks have been constructed, exact scalar dynamics require strong lumpability: any two full operator states with the same mass partition must have the same total jump intensity into every new mass partition. Merger and split hazards that depend only on masses are a transparent sufficient condition, not the necessary condition itself. If equal-mass states have different aggregate transition intensities because of network position, creditor structure or organisational centrality, the scalar projection contains hidden state.

This condition is empirically testable. Researchers can compare blocks with the same size distribution but different control structures. Systematic differences in transition hazards reject an autonomous size-only model; their absence supports the scalar representation as a parsimonious approximation.

The integrated theory can therefore be described in five steps:

1. monetary and financial conditions determine the general pace and selectivity of opportunities;
2. acquirer capacity and target vulnerability select events;

3. events move B, I, H and R in distinct ways;
4. a projection rule converts the control structure into mass blocks;
5. acquisition and fragmentation determine formation, memory and possible routes of reconstruction.

The scheme also locates possible irreversibility. Under the finite-chain support conditions, restored parameters imply return in distribution for the scalar mass process. At the operator level, merger is non-invertible and a split selects one reconstruction. Structural memory may therefore lie in the gap between the return of masses and the failure of the control structure to return.

8. An integrated empirical programme

The synthesis generates an agenda for ownership, control, balance-sheet and transaction data. Its purpose is to convert general propositions into falsifiable restrictions. The programme has seven stages.

8.1 Replicate before extending

The first stage reproduces the established centralisation measure on the same universe of firms and under the same coverage rules. Direct and indirect ownership, integrated values, rankings and the minimum share of controllers needed to reach 80 per cent are reconstructed. Margins around the threshold are also reported, so that structural changes can be distinguished from crossings caused by near ties.

This discipline is essential. The proposal does not replace the existing series with an incompatible measure. It preserves it as a benchmark and tests whether the added dimensions contain new information.

8.2 Reconstruct the operator and the blocks

For each date, the rule mapping ownership shares, voting rights or other sources into control links must be declared. Robustness should be assessed under transparent alternative rules. The operator is then used to compute reciprocal cores, hierarchy among cores, amplification, localisation and transient stress.

The same reconstruction defines blocks under common command. Aggregation by legal identity alone is insufficient when pyramids or cross-holdings exist. Mass can be measured by assets, capital, employment, revenues or integrated values; robustness to that choice is part of the analysis.

8.3 Construct an event panel

An annual index is too coarse to identify mechanisms. The minimum dataset should record:

- legal identities and masses before and after the event;
- buyer, seller and target;
- the source of the acquired stake;
- control links and their weights;
- solvency, profitability, debt and financing conditions;

- regulatory interventions;
- the genealogy of consolidated blocks.

Each event is classified as a strong merger, accretive acquisition, transfer, pure divestiture or split. Researchers then compute changes across the four frontiers. The monotone cells of the signature table become over-identifying restrictions.

8.4 Estimate acquisition intensity and selectivity

For every potential acquirer-target pair, a risk set is constructed. Acquisition probability is explained using pre-event size, solvency, liquidity, debt costs, sector, country, network position and common monetary conditions. The time baseline identifies the general pace; size elasticities identify selectivity.

The roles of buyer and target should be distinguished. A regime in which large buyers acquire targets of every size is not equivalent to a regime in which deals predominantly involve pairs of large actors. A symmetric specification with one β is a benchmark to be tested, not a restriction always to be imposed.

8.5 Estimate fragmentation intensity and scalability

Fragmentation events must be separated by cause: antitrust remedies, voluntary spin-offs, bankruptcy separations, forced sales, resolution, dissolution of common control or loss of a decisive stake. For each cause, researchers estimate how the hazard varies with block size.

The exponent δ should not automatically receive a political interpretation. Low scalability may reflect public protection, but also legal complexity, asset specificity, technological integration or political influence. The theory identifies the profile; historical and institutional research must identify its cause.

8.6 Follow fragments through competing risks

After every major split, at least three outcomes should be recorded: direct remerger, attachment of one fragment to a third block, and renewed fragmentation. The empirical design compares the timing and probability of these events in relation to β, δ, fragment size and background structure.

The channel exponents χ^D and χ^T, together with the aggregate index $\chi^A(s)$, can be used as regressors or stratification variables. The prediction is that direct remerger becomes relatively more frequent when χ^D turns positive, while attachment to third parties becomes relatively more frequent when χ^T is positive; $\chi^A(s) > 0$ indicates that at least one coagulative first-event channel dominates renewed fragmentation locally. Subsequent persistence must be measured separately: one attachment event does not by itself demonstrate durable macroscopic control.

8.7 Link micro-events to macroeconomic dynamics

The final stage returns the new variables to the macroeconomic framework. Instead of one centralisation series, researchers may use acquisition selectivity, the fragmentation-acquisition ratio, fragmentation scalability, the number of autonomous poles, hierarchy, amplification and transient stress. These variables can be related to interest rates, credit, solvency and activity through dynamic models without asking aggregate models to identify the microeconomic structure on their own.

The decisive test is non-redundancy. If operator profiles and event hazards add no explanatory or predictive information beyond the established statistic and conventional controls, the extension is not empirically progressive. If they diverge in interpretable episodes, the programme acquires new content.

9. Policy implications

9.1 Monetary policy and competition policy are structurally coupled

If financial conditions raise β, large buyers gain a more than proportional advantage. In that situation, a competition authority that keeps the number of interventions or the ratio $\eta = \varphi/\kappa$ unchanged does not necessarily preserve its structural capacity. The local post-break-up boundaries imply that fragmentation scalability must rise with acquisition selectivity if the balance of first-event races is to be maintained; more generally, authorities may also need to increase η or improve the split law and the quality of remedies.

Monetary policy therefore creates a structural externality for competition policy. Central banks, antitrust authorities and resolution agencies operate with different mandates and information sets, but their effects converge on the same process of selection and recombination of capital.

9.2 From the number of remedies to their architecture

The number of splits or divestitures does not determine remedy effectiveness. Three questions remain:

1. does it reduce the mass of the legal block?
2. does it separate control and dependence links?
3. does it prevent fragments from recombining directly or through third parties?

A policy may succeed on the first dimension and fail on the other two. The distinction among B, I, H and R makes these outcomes separately measurable. Evidence on merger control also shows that effectiveness depends on remedy design rather than formal intervention alone (Duso, Gugler and Yurtoglu 2011). Post-remedy monitoring is part of the economic definition of success.

9.3 From too-big-to-fail to too-big-to-break

Too-big-to-fail describes a situation in which the costs of an entity's failure induce public intervention. *Too-big-to-break* is broader. A block is too-big-to-break when the effective capacity to divide it and make its parts autonomous scales too slowly relative to its size and its routes of recombination.

Possible causes include bailouts, but also asset complementarities, common debt, indivisible infrastructure, intellectual property, political power or the absence of independent buyers. The concept does not imply that the block cannot be divided legally. It implies that legal division need not produce a proportional reduction in command.

9.4 Evaluation must use institutional horizons

The fact that a finite model eventually returns to its long-run distribution does not justify waiting. If relaxation takes decades, mathematical reversibility is irrelevant for current policy. Appropriate horizons include the life of a remedy, a technological cycle, a regulatory mandate or an investment generation.

Policy should therefore compare repair, reattachment and dissolution times with institutional horizons, not with infinity. The question shifts from the abstract existence of a path back to its practical availability.

10. Limits and research agenda

The synthesis preserves the differences among the three models. The first concerns measurement at a fixed date, the second finite event dynamics on a changing-dimensional operator, and the third aggregate mass dynamics. Complete integration still requires a stochastic process on the full space of control operators.

Four problems are priorities.

The first is the **projection rule**. Researchers must determine when legally separate entities should be treated as one strategic block and test the robustness of that decision.

The second is the **network dependence of hazards**. Hierarchical position, cycle membership, creditor centrality and amplification may affect acquisitions and splits. If so, the mass-only model is not closed.

The third is **selection among post-break-up reconstructions**. The theory characterises the admissible set but does not determine which lift is selected. Institutional, contractual and strategic models of separation are needed.

The fourth is global metastability. The channel exponents χ^D and χ^T and the aggregate index $\chi^A(s)$ describe local first-event races. Results are still needed on transition times between high- and low-centralisation regions and on how those times grow with system size.

Empirically, the greatest difficulty is data quality. Ownership, control, debt, contracts and supply relations are rarely observed in one source. A defensible design declares the scope of measurement, uses several attribution rules and treats robustness as part of the theory.

Causal inference also requires caution. A monetary-policy change may act simultaneously on stand-alone values, credit availability, target vulnerability and the probability of public intervention. The proposed decomposition does not automatically solve causal identification, but it at least prevents distinct channels from being collapsed into one series.

11. Conclusion

The capital-centralisation programme has shown that money, credit and solvency select capitals and contribute to the reorganisation of ownership and control. The three developments synthesised in this working paper broaden that programme without abandoning its core.

The first development moves from a ranking to the control structure that generates it: the same concentration statistic may correspond to a dominant pole, autonomous poles, a pyramid or a fragile chain. The second classifies changes between snapshots by primitive event and derives their distinct effects on legal mass, integration, hierarchy and spectral feedback. The third separates formation from reversibility. Deconcentration depends on the scalability and design of remedies and on the routes through which fragments may recombine or redirect control.

A unified proposition follows. Centralisation is best represented as a causal and institutional sequence: solvency selection, event selection, transformation of command, and persistence or dissolution. Measurement requires a profile, explanation requires transition hazards, and governance requires coordination among institutions that currently operate separately.

The final question extends beyond whether a giant block forms. It concerns the form of command it assumes, the event that strengthens it, and whether dismantling capacity scales more rapidly than the forces of recombination. This is the reversibility frontier of capital centralisation.

Acknowledgments

The author thanks Mauro Maltagliati for careful proof-checking of the formal arguments in the three companion working papers. The usual disclaimer applies.

Use of generative AI

OpenAI's ChatGPT was used as a research-and-editing aid to organise and stress-test the synthesis, assist with mathematical and terminological consistency checks, and support stylistic and bibliographic revision. It was not used to generate research data. All theoretical choices, numerical claims, references and final wording were reviewed by the author, who takes full responsibility for the manuscript.

Data and code statement

This synthesis reports no new observational data and no new numerical simulation. Its examples and formal results are drawn from the three cited companion working papers and their associated appendices and replication materials.

References

Becher, D. A., Griffin, T. P., and Nini, G. (2022). "Creditor control of corporate acquisitions." *Review of Financial Studies*, 35(4), 1897–1932. https://doi.org/10.1093/rfs/hhab075

Bellanca, N. (2026a). *From Concentration to Control: Measuring the Architecture of Capital Centralisation*. SSRN Working Paper, Abstract ID 7161739, 22 July 2026. https://ssrn.com/abstract=7161739

Bellanca, N. (2026b). *When Control Structures Coagulate: Event Signatures and the Four Frontiers of Capital Centralisation*. SSRN Working Paper, Abstract ID 7167758, 23 July 2026. https://ssrn.com/abstract=7167758

Bellanca, N. (2026c). *The Reversibility Frontier of Capital Centralisation: Credit Selectivity, Fragmentation, and Too-Big-to-Break*. SSRN Working Paper, Abstract ID 7167604, 23 July 2026. https://ssrn.com/abstract=7167604

Brancaccio, E., and Fontana, G. (2013). "'Solvency rule' versus 'Taylor rule': An alternative interpretation of the relation between monetary policy and the economic crisis." *Cambridge Journal of Economics*, 37(1), 17–33. https://doi.org/10.1093/cje/bes028

Brancaccio, E., and Fontana, G. (2016). "'Solvency rule' and capital centralisation in a monetary union." *Cambridge Journal of Economics*, 40(4), 1055–1075. https://doi.org/10.1093/cje/bev068

Brancaccio, E., Giammetti, R., Lopreite, M., and Puliga, M. (2018). "Centralization of capital and financial crisis: A global network analysis of corporate control." *Structural Change and Economic Dynamics*, 45, 94–104. https://doi.org/10.1016/j.strueco.2018.03.001

Brancaccio, E., Giammetti, R., Lopreite, M., and Puliga, M. (2019). "Monetary policy, crisis and capital centralization in corporate ownership and control networks: A B-VAR analysis." *Structural Change and Economic Dynamics*, 51, 55–66. https://doi.org/10.1016/j.strueco.2019.08.005

Brancaccio, E., Giammetti, R., Lopreite, M., and Puliga, M. (2023). "Convergence in solvency and capital centralization: A B-VAR analysis for high-income and euro area countries." *Metroeconomica*, 74(1), 40–73. https://doi.org/10.1111/meca.12401

Brioschi, F., Buzzacchi, L., and Colombo, M. G. (1989). "Risk capital financing and the separation of ownership and control in business groups." *Journal of Banking & Finance*, 13(4–5), 747–772. https://doi.org/10.1016/0378-4266(89)90040-X

Duso, T., Gugler, K., and Yurtoglu, B. B. (2011). "How effective is European merger control?" *European Economic Review*, 55(7), 980–1006. https://doi.org/10.1016/j.euroecorev.2011.04.003

Glattfelder, J. B., and Battiston, S. (2009). "Backbone of complex networks of corporations: The flow of control." *Physical Review E*, 80(3), 036104. https://doi.org/10.1103/PhysRevE.80.036104

La Porta, R., Lopez-de-Silanes, F., and Shleifer, A. (1999). "Corporate ownership around the world." *Journal of Finance*, 54(2), 471–517. https://doi.org/10.1111/0022-1082.00115

Marx, K. (1976 [1867]). *Capital: A Critique of Political Economy, Volume I*, trans. B. Fowkes. Harmondsworth: Penguin.

Trefethen, L. N., and Embree, M. (2005). *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton, NJ: Princeton University Press.

Vitali, S., Glattfelder, J. B., and Battiston, S. (2011). "The network of global corporate control." *PLoS ONE*, 6(10), e25995. https://doi.org/10.1371/journal.pone.0025995