

Parsimony by Repetition, Not Reduction

CAMS in Relation to Sovereign–Civilian Capacity Dynamics

Kari McKern

18 July 2026

Abstract

Vi Nhat Son’s *Sovereign–Civilian Capacity Dynamics* (SCCD) presents an intellectually attractive compression of political development. Two aggregate capacities—sovereign surplus capacity $S(t)$ and civilian capacity $C(t)$ —generate a differential,

$$\Delta(t) = S(t) - C(t),$$

intended to describe the structural range of action available to a state. At first sight, this appears considerably more parsimonious than CAMS, which records four attributes across eight recurrent social functions. Yet variable count alone is a poor measure of parsimony. SCCD’s two headline variables decompose into ten heterogeneous latent constructs, fourteen proxy formulae, multiple non-linear aggregators, threshold functions, feedback loops and, in the network formulation, a rapidly expanding parameter space. CAMS begins with more observations but uses a smaller repeated grammar: the same four questions are asked of each of eight functionally defined nodes.

This paper argues that SCCD exhibits *scalar parsimony*, while CAMS may exhibit a stronger form of *generative* or *architectural parsimony*. SCCD reduces a polity rapidly to a state–civilian differential. CAMS preserves the topology from which such differentials may emerge. Most SCCD components can plausibly be reconstructed as projections across CAMS nodes, whereas the reverse reconstruction is generally impossible because aggregation into S , C and Δ destroys information about which social functions are carrying capacity, cognition, stress and activation. This does not establish CAMS as the superior model. Its apparent economy may conceal flexibility in archetypal interpretation and complexity within the scoring process. The parsimonious advantage of CAMS therefore remains conditional on ablation, independent scoring, nested-model comparison and genuinely out-of-sample tests.

1 The Problem of Parsimony

Parsimony is often treated as though it meant choosing the theory with the fewest variables. That is too crude. A theory with two headline variables may require dozens of subsidiary assumptions to make those variables observable. A theory with thirty-two observations may be comparatively economical if those observations are generated by repeating the same small measurement grammar.

At least four forms of parsimony should be distinguished:

Ontological parsimony

The number of fundamentally different kinds of entity assumed by a

theory.

Parametric parsimony

The number of adjustable coefficients, thresholds and functional forms.

Measurement parsimony

The number of distinct procedures required to translate observations into model variables.

Explanatory compression

The amount of independently testable structure recovered relative to the model's effective degrees of freedom.

These measures need not agree. SCCD is strikingly parsimonious at the level of its central narrative:

$$S(t), \quad C(t), \quad \Delta(t) = S(t) - C(t).$$

CAMS is less compressed at the surface. It represents a society as an 8×4 matrix:

$$\mathbf{X}_t = \{c_i, k_i, s_i, a_i\}_{i=1}^8,$$

where the eight nodes are Helm, Shield, Lore, Stewards, Craft, Hands, Archive and Flow, and the four recurring attributes represent capacity, available energy, stress and activation or abstraction.

But SCCD's apparent economy is achieved by aggregating early. CAMS aggregates late. The relevant question is therefore not merely which model has fewer cells. It is which model requires fewer independent conceptual rules to produce useful historical discrimination.

2 SCCD's Elegant Compression

SCCD begins from a clear and powerful proposition: the ability of a state to act depends not on its absolute resources but on the balance between the discretionary capacity retained by the state and the autonomous capacity accumulated by society. The model therefore distinguishes sovereign surplus from the resources required merely to maintain order, while civilian capacity represents the economic, organisational, informational and political ability of non-state actors to act independently [1, pp. 1–3].

This is a genuine theoretical contribution. It redirects attention from possession to action. GDP, military expenditure and tax revenue may indicate what resources exist without revealing what can actually be mobilised. SCCD's differential is intended to capture the resulting feasibility envelope.

The problem is that neither S nor C is primitive in the operational model. Sovereign capacity

is decomposed into:

$$\begin{aligned} S_1 &= \text{fiscal extraction,} \\ S_2 &= \text{coercive monopoly,} \\ S_3 &= \text{legitimacy capital,} \\ S_4 &= \text{administrative penetration,} \\ S_5 &= \text{temporal leverage.} \end{aligned}$$

Civilian capacity is likewise decomposed into private capital concentration, civil organisation, information autonomy, exit options and preference fragmentation [1, pp. 6–11].

The paper explicitly describes the fivefold sovereign decomposition as the minimum set of qualitatively distinct components required by the historical analysis [1, p. 18, n. 2]. The choice is plausible, but it has not yet been demonstrated to be minimal in the statistical sense. No component-ablation result shows that four components are insufficient, that six add no information, or that five outperform alternative decompositions under an information criterion or out-of-sample loss function. “Minimum” here means theoretically selected, not empirically established.

This distinction becomes important because the components are not treated symmetrically. Legitimacy is a Hill-function multiplier; coercive monopoly and temporal leverage are threshold conditions; fiscal extraction and administrative penetration are partially substitutable; civilian fragmentation enters as an inhibitor [1, pp. 20–23]. SCCD therefore requires not only ten latent variables but a theory of the unique mathematical role played by each one.

The full single-state dynamical model contains approximately fifteen parameters before the network extensions are added. Its reduced form is deliberately introduced because the complete model is too demanding for ordinary empirical estimation [1, pp. 34–35]. In the integrated multi-state version, the paper estimates approximately $15N + N(N - 1)$ parameters, reaching roughly 1,680 parameters for twenty major states [1, pp. 80–81].

SCCD is therefore parsimonious as a political narrative, but not necessarily as a complete generative model.

3 The Different Economy of CAMS

CAMS begins with eight functional attractors rather than a binary state–civilian division. Each attractor is scored using the same four-dimensional grammar. In simplified form, the node’s energetic condition can be represented as

$$H_i = K_i - S_i,$$

its cognitive or directed informational activity as

$$Q_i = A_i C_i,$$

and its canonical viability as

$$V_i = C_i + K_i - S_i + \alpha A_i.$$

The exact interpretation of these quantities remains open to refinement, but the underlying economy is clear: CAMS does not invent a different measurement logic for each node. Helm, Shield, Hands and Archive are all asked the same four questions.

This repeated symmetry matters. CAMS contains thirty-two annual readings, but it contains only four metric *types*. SCCD contains ten headline component readings, but each is a different kind of construct requiring a bespoke operational definition.

CAMS is thus not numerically sparse. It is *grammatically sparse*.

The contrast can be expressed as follows:

SCCD: 10 heterogeneous variables $\longrightarrow$ 2 aggregates $\longrightarrow$ 1 differential,

whereas

CAMS: 4 recurring attributes $\times$ 8 functional locations $\longrightarrow$ node and system morphology.

SCCD achieves compression by collapsing social structure into a polarity. CAMS achieves compression by repeating one measurement grammar across a fixed functional topology.

Neither architecture is intrinsically superior. Their relative parsimony depends on the target of explanation. For the narrow question, “How much unilateral discretionary power does a state possess relative to its society?”, SCCD may be sufficient and more intelligible. For questions about the overall metabolic, cognitive and affective organisation of a society, the early reduction to $S - C$ may discard precisely the information that matters.

4 SCCD as a Projection of CAMS

The most consequential relationship between the frameworks is that SCCD’s components appear to cut across CAMS nodes rather than correspond to separate social organs.

A provisional mapping illustrates this:

SCCD construct	Principal CAMS locations
Fiscal extraction S_1	Helm, Stewards, Flow, Archive
Coercive monopoly S_2	Shield and Helm
Legitimacy S_3	Lore, Archive and Helm coordination
Administrative penetration S_4	Archive, Hands, Flow and Helm
Temporal leverage S_5	Archive, Stewards and Lore
Private capital C_1	Stewards and Flow
Civil organisation C_2	Hands, Lore, Craft and Archive
Information autonomy C_3	Lore, Flow and Craft
Exit capacity C_4	Flow, Stewards and Craft
Preference fragmentation C_5	Cross-node incoherence, bond loss and stress shear

Table 1: A provisional mapping from SCCD components to CAMS functional locations.

This mapping suggests that SCCD’s variables may be second-order composites. Legitimacy, for example, is unlikely to reside in a single institutional sector. It is produced by the relation between narrative production, institutional memory, executive coordination and lived social ex-

perience. Administrative penetration likewise requires not only bureaucracy but communication, labour, implementation and institutional memory.

CAMS can therefore represent a society first and derive sovereign and civilian capacity afterwards:

$$\mathcal{S}_t = F_S(\mathbf{X}_t, \boldsymbol{\omega}_t),$$

$$\mathcal{C}_t = F_C(\mathbf{X}_t, \boldsymbol{\omega}_t),$$

$$\Delta_t^{\text{CAMS}} = \mathcal{S}_t - \mathcal{C}_t,$$

where $\boldsymbol{\omega}_t$ describes the historically changing degree to which each functional node is incorporated into the sovereign apparatus.

This last qualification is essential. A bank, university, media organisation, defence contractor or digital platform may be formally civilian while functionally integrated into sovereign power. Conversely, a nominally state institution may be captured by private interests. The sovereign–civilian boundary is therefore not stable across countries or historical periods.

CAMS does not need to decide permanently whether Stewards, Lore or Flow belong to the state or to civil society. It observes their functional condition first. Their sovereign alignment can be estimated as a secondary relation.

The reverse transformation is not generally possible. Once the eight-node structure has been compressed into S and C , one cannot determine whether a declining differential arose from Shield exhaustion, Archive decay, Stewards capture, Hands alienation, Flow fragmentation or some combination of them. SCCD itself acknowledges that identical values of Δ may describe very different political realities and that the scalar differential cannot preserve compositional distinctions [1, p. 12].

In information terms, CAMS can potentially generate SCCD as a projection. SCCD cannot reconstruct CAMS from its aggregate outputs. That asymmetry is evidence—not proof—of greater explanatory compression in CAMS.

5 The Hidden Complexity of Operationalisation

The strongest argument for CAMS parsimony lies not in the number of model variables but in the operational burden required to populate them.

SCCD concedes that none of its ten subcomponents is directly observable. Every component is a latent variable requiring multiple proxies, validity judgements, scaling rules and error propagation [1, pp. 38–39]. Its operational architecture eventually contains fourteen proxy formulae for the ten components, a multi-stage composite-construction procedure, reliability weighting, missing-data imputation and Monte Carlo uncertainty propagation [1, p. 72].

This is methodologically conscientious, but it means that much of the model’s complexity has migrated from the equations into the measurement layer. Two analysts may agree on the theoretical definition of legitimacy capital while disagreeing over surveys, bond spreads, tax compliance, voluntary enlistment or elite cohesion. Different proxy selections may generate different values of S_3 , which then propagate non-linearly because legitimacy is a multiplier.

Country-specific normalisation introduces another layer. SCCD scales components relative to each polity's historical range and sometimes uses a rolling thirty-year maximum [1, p. 20]. That may assist within-country trajectory analysis, but it makes the scale itself historically endogenous. A state can appear to change partly because the denominator has changed.

CAMS faces a different measurement problem. Its 0–10 node scores are semantically broad and currently depend substantially on structured expert or language-model assessment. The operation is uniform, but the scorer embodies a large amount of historical knowledge and interpretive judgement. Apparent simplicity in the rubric can therefore conceal complexity in the assessor.

This is the most serious challenge to a CAMS parsimony claim. A flexible scorer may silently accommodate anomalies, reinterpret node boundaries or infer the intended historical morphology. An eight-node framework can become unfalsifiably elastic if every event is allowed to express itself in several nodes simultaneously.

Accordingly, CAMS cannot claim parsimony merely because it has a repeated rubric. It must show that:

1. different scorers recover similar structures;
2. the structures survive blinded identities and years;
3. node removal causes predictable information loss;
4. alternative node systems perform worse; and
5. failed predictions are retained rather than repaired retrospectively.

The existing CAMS/JUNO programme has begun to meet this standard through ensemble scoring, identity blinding, family classification and explicit recording of failed hypotheses. The rejected morphospace prediction, failed crisis-scaling rule, uncertain abstraction weight and largely negative Pulse precursor results are scientifically important. They show that the framework has permitted evidence to reduce its claims. That kind of empirical subtraction is itself a form of parsimony.

6 Scalar Sufficiency Versus Morphological Sufficiency

The real theoretical disagreement is not between ten variables and thirty-two variables. It is between two views of what must be preserved.

SCCD assumes that the state–civilian capacity balance is the privileged macroscopic coordinate. The ten components matter because they determine S and C ; their ultimate explanatory role is to locate the polity on the Δ axis.

CAMS assumes that a society's condition is morphological. Several macroscopic quantities may be derived from the eight-node configuration, but no single one is presumed sufficient in advance. A society can have a stable mean while developing severe internal shear. High average capacity may coexist with a collapse of the Archive–Hands relation, excessive Shield activation,

cognitively intense but energetically depleted Lore, or a decoupling between executive direction and productive implementation.

The SCCD paper recognises part of this problem when it notes that the same differential can arise from different compositions. Yet its solution is to retain the ten-dimensional decomposition beneath the scalar. CAMS goes further: it treats relational configuration, not only component magnitude, as constitutive of the system.

This difference matters near transitions. Critical change may begin not as a fall in total capacity but as altered coupling among functions. A state may possess fiscal, coercive and administrative resources while losing the ability to convert them into coordinated action. Another may possess modest resources but maintain unusually effective alignment among narrative, archive, executive direction and labour.

SCCD represents some of this through legitimacy multipliers, fragmentation discounts and coherence thresholds. CAMS asks whether these apparently distinct correction terms are symptoms of a more general phenomenon: the topology of functional coordination.

If so, CAMS could be more parsimonious because several SCCD mechanisms would become manifestations of one underlying architecture. Regulatory capture would appear as Stewards–Helm domination and bond distortion. Legitimacy collapse would appear as divergence among Lore, Archive, Helm and Hands. Administrative failure would appear as breakdown between Archive, Flow and Hands. Preference fragmentation would appear as system-wide phase dispersion rather than as an isolated civilian variable.

The theory would then replace several specially formulated loops with recurrent changes in node condition and coupling.

7 Where SCCD Remains More Parsimonious

A fair comparison must also identify where SCCD has the advantage.

First, its explanandum is narrower. SCCD is specifically a theory of state action, civilian accumulation and hegemonic decline. CAMS aims to describe a wider range of social cognition, metabolism, stress and coordination. A narrower model should not be penalised for omitting distinctions irrelevant to its target.

Second, the state–civilian differential is politically interpretable. Policymakers can understand the distinction between discretionary sovereign capacity and autonomous civilian power. CAMS archetypes require more translation before they can guide institutional action.

Third, SCCD contains explicit causal hypotheses. It identifies productive drain, distributional drain, capture, reconstitution, legitimacy multiplication and exit dynamics. CAMS currently has stronger evidence as a system of structural description than as a fully identified causal mechanism. Its node trajectories can recover historically coherent morphology, but that does not by itself establish why the trajectories changed.

Fourth, SCCD has specified a proper prospective standard. It proposes that a model trained on 1970–2000 should predict regime transitions in 2001–2020, achieve an AUC of at least 0.70, keep level error below a specified threshold and outperform an autoregressive baseline [1, p. 58]. The paper has not yet completed this test, but the criterion is clear.

Finally, SCCD is appropriately restrained about its present status. Its case-study validation is explicitly qualitative-structural rather than quantitatively predictive [1, p. 87]. It also acknowledges that threshold parameters come from too few historical cases, that authoritarian data are especially uncertain and that pre-1990 reconstructions require manual coding and structural imputation [1, pp. 102–103].

CAMS should adopt the same clarity. Its stronger parsimony claim is presently a hypothesis, not a settled result.

8 A Decisive Nested-Model Test

The relationship between the frameworks can be tested rather than argued indefinitely.

A common corpus should be scored independently under both systems. It should include multiple political forms, blinded society identities, annual observations and holdout periods not used during model construction.

Five model classes would then be compared:

- M_0 : autoregressive and conventional-indicator baselines,
- M_1 : SCCD reduced form,
- M_2 : SCCD non-linear single-state model,
- M_3 : full CAMS eight-node model,
- M_4 : SCCD quantities derived as projections of CAMS.

The target outcomes should not be defined solely by either framework. They might include independently coded regime transitions, mobilisation capacity, fiscal discretion, policy implementation failure, state fragmentation, recovery duration and blinded society identification.

Parsimony should then be assessed using held-out likelihood, calibration, effective parameter count, minimum description length and robustness to rescaling.

The critical comparisons are straightforward. CAMS would receive strong support if the eight-node configuration predicted held-out outcomes materially better than SCCD, if deriving S , C and Δ from CAMS retained SCCD’s useful predictions, and if removing nodes or collapsing the topology caused reproducible losses.

SCCD would receive strong support if its reduced form matched or exceeded CAMS on external outcomes with substantially fewer effective degrees of freedom, or if the CAMS topology added descriptive richness without additional predictive information.

A third result is possible and perhaps most likely: SCCD may prove to be a useful low-dimensional projection for state-action questions, while CAMS supplies the higher-dimensional morphology required to explain why societies with similar differentials behave differently.

Under that result, the models would not be rivals at the same level. SCCD would stand to CAMS as a gauge stands to a field: highly useful, interpretable and economical, but deliberately incomplete.

9 Conclusion

The central achievement of SCCD is to identify a compelling political differential: the changing balance between the discretionary capacity of the state and the autonomous capacity of the society it governs. Its conceptual top layer is elegant and parsimonious.

Its full architecture, however, is not a two-variable theory. It is a ten-latent-variable measurement system with heterogeneous aggregation rules, approximately fifteen parameters in the single-state dynamical formulation, numerous thresholds and feedback loops, and a network extension whose parameter count grows rapidly with the number of states. The paper itself recognises that the integrated model requires a hierarchy of simpler approximations for practical use [1, pp. 80–81].

CAMS presents the inverse profile. It is high-dimensional on the surface but comparatively uniform underneath. Eight recurrent functions are described through four recurring attributes. Its possible parsimony resides not in having fewer observations but in generating many social distinctions from a repeated architecture.

The strongest defensible conclusion is therefore conditional:

SCCD is more parsimonious as a scalar theory of sovereign action. CAMS may be more parsimonious as a generative theory of whole-society organisation.

Most of SCCD’s components can plausibly be represented as cross-sections through the CAMS topology. The reverse reconstruction is not possible after aggregation because information about functional location and coordination has been lost. That gives CAMS a potentially deeper form of explanatory compression: it may recover the SCCD differential while retaining the morphology from which the differential emerged.

But potential is not proof. CAMS must still demonstrate that eight nodes are necessary, that four metrics are sufficient, that its scoring process does not hide uncontrolled flexibility and that its additional resolution improves performance on genuinely independent targets. The decisive evidence will come from nested ablation and out-of-sample comparison, not from elegance alone.

Parsimony is not the art of describing the world with the smallest number of symbols. It is the art of preserving the largest amount of recoverable structure with the smallest number of independent rules. On that stricter criterion, CAMS has a serious claim—but one that should now be tested directly against SCCD rather than merely asserted.

References

- [1] Vi Nhat Son. *Sovereign–Civilian Capacity Dynamics: A Formal Theory of State Power, Civilian Accumulation, and Imperial Decline*. SSRN working paper no. 6499959, 2026.