



A Swarm Intelligence-based Model for Organisational Governance: Critical Theoretical Foundations

Edwin Jesu Dass* and Rashad Yazdanifard
Kings University College, Malaysia

Abstract

Hierarchical governance models face persistent structural limitations under volatile, uncertain, complex, and ambiguous (VUCA) conditions, generating slow adaptive response, information bottlenecks, and coordination failures that reduce organisational effectiveness in complex environments. This paper introduces the Swarm Governance Framework (SGF) — the first organisational governance framework derived from swarm intelligence (SI) theory — and the accompanying STRUM model (Signals, Trust, Resonance, Unity, Measure), a five-element design specification that translates the framework's design rules into actionable governance parameters. A critical theoretical review of four SI principles — stigmergy, self-organisation, emergence, and adaptive feedback — operationalises each into a concrete governance design rule, collectively constituting the SGF. The STRUM model simultaneously constitutes a configuration specification for AI coordination agents in agentic enterprise systems, positioning the SGF at the intersection of organisational governance and emerging agentic enterprise design. The paper represents the artefact construction phase of a Design Science Research programme. The dual applicability of STRUM to human governance units and AI agents positions the SGF as a theoretically unified governance framework for the integrated human-AI enterprise.

Keywords: swarm intelligence, organisational governance, STRUM model, design science research, AI coordination

1 Introduction

Contemporary organisations operate in environments that defy the assumptions embedded in hierarchical governance design. Increasing market volatility, technological disruption, geopolitical complexity, and accelerating rates of change have rendered the stable, predictable operating conditions on which traditional command-and-control governance models were premised an increasingly rare exception. Dynamic capabilities theory recognises that sustained organisational performance in complex environments requires governance structures capable of sensing, seizing, and reconfiguring resources rapidly — a capacity that

hierarchical governance architectures, designed for stability, are not equipped to provide (Teece, 2007). The VUCA (volatile, uncertain, complex, and ambiguous) framework, originally formulated to characterise post-Cold War military environments, has been widely adopted in management research to describe the conditions that contemporary organisations must navigate (Bennett & Lemoine, 2014; Lawrence, 2013). Under VUCA conditions, the structural assumptions of hierarchical governance — centralised decision-making authority, sequential information processing, predetermined coordination protocols, and fixed organisational rules — generate persistent limitations: slow adaptive response, information bottlenecks, coordination failures, and an inability to leverage distributed intelligence.

Despite extensive attention to the governance-complexity problem, the organisational theory literature has not yet produced a framework that is both theoretically grounded in the mechanisms of collective adaptive behaviour and capable of generating concrete design rules for organisational architecture. Existing responses include holacracy (Robertson, 2015), sociocracy, and various distributed organisation models (Laloux, 2014). These approaches reduce hierarchical centralisation but remain practitioner-derived and lack a theoretical derivation logic grounded in the mechanisms through which complex adaptive systems produce effective distributed coordination. Research on new forms of organising (Puranam et al., 2014) and sociotechnical systems design (Baxter & Sommerville, 2011) underscores that effective governance in complex environments must be designed with attention to how agent interactions — human and technical — produce emergent system-level behaviour.

This paper addresses this gap by conducting a critical theoretical review of swarm intelligence (SI) — the branch of artificial intelligence and biology concerned with collective behaviour emerging from decentralised, locally informed agents — and applying its core principles to organisational governance design. The review introduces the Swarm Governance Framework (SGF) and the STRUM model (Signals, Trust, Resonance, Unity, Measure) as original theoretical contributions. A secondary contribution positions the STRUM model as a configuration specification for AI coordination agents in agentic enterprise systems.

2 Swarm Intelligence as a Theoretical Foundation

Swarm intelligence refers to the collective behaviour that emerges from populations of locally interacting, decentralised agents following simple rules, without any individual agent possessing global knowledge or directing overall behaviour (Bonabeau et al., 1999; Dorigo & Birattari, 2007). The field draws its theoretical and empirical foundations from the study of social insects — particularly ants, bees, and termites — and has been formalised into computational paradigms underpinning Ant Colony Optimisation (Dorigo & Gambardella, 1997) and Particle Swarm Optimisation (Kennedy & Eberhart, 1995). The defining characteristics of swarm-intelligent systems — decentralisation, local interaction, indirect coordination, and coherent collective behaviour — have been extensively validated across natural and artificial systems (Camazine et al., 2001).

The relevance of swarm intelligence to organisational theory is not analogical but derivational. SI principles do not merely describe patterns resembling organisational dynamics; they constitute a well-validated theoretical account of how distributed adaptive coordination works at a mechanistic level, from which concrete governance design rules can be derived. Sociotechnical systems research establishes that organisational performance in complex environments arises from the interaction of social and technical components rather than from either in isolation (Baxter & Sommerville, 2011) — a framing consistent with the SGF's positioning of governance as an emergent property of distributed agent interactions

rather than a product of central design. What SI contributes beyond complexity theory is a mechanistic, experimentally validated account of the specific principles through which decentralised systems achieve coordination, adaptation, and coherence, enabling the move from descriptive complexity theory to prescriptive governance design.

3 Critical Review: Four Swarm Principles and Governance Translational Relevance

The principles selected for this review were identified through a structured screening of the swarm intelligence literature using three explicit criteria: (1) the principle must be documented across multiple independent empirical studies in both natural and artificial swarm systems; (2) it must address at least one of the four core governance functions — coordination, decision-making, pattern recognition, or adaptation; and (3) it must be directly operationalisable into a concrete organisational design rule without requiring analogical extension that would distort the underlying mechanism. The literature search drew on foundational SI texts (Bonabeau et al., 1999; Camazine et al., 2001; Dorigo & Birattari, 2007) and extended to governance and complexity literatures to assess translational relevance. Four principles satisfied all three criteria: stigmergy, self-organisation, emergence, and adaptive feedback.

3.1 Stigmergy

Stigmergy was first described by Grassé (1959) to explain coordinated construction behaviour in termites without direct inter-individual communication. Agents coordinate indirectly by modifying their shared environment; subsequent agents respond to those traces rather than to direct instruction, producing coordinated behaviour without any agent possessing a global plan. The governance translation of stigmergy is signal-based indirect coordination — designing information environments in which distributed units act in response to shared signals rather than explicit top-down instruction.

3.2 Self-Organisation

Self-organisation is the capacity of a system to develop coherent global patterns through local interactions among its components, without external direction (Camazine et al., 2001). It operates through positive and negative feedback loops, threshold responses, and randomness, producing stable and adaptive structures without central computational design (Bonabeau et al., 1999). Lichtenstein (2000) identifies self-organising transitions as a documented mode of organisational adaptation across multiple empirical contexts. The governance translation is autonomous unit decision-making within shared constraints — authority delegated to the level at which information is available and action is required.

3.3 Emergence

Emergence describes the appearance of novel system-level properties from local interactions that cannot be predicted or derived from individual component properties (Holland, 1998). Anderson (1999) identifies emergence as the hallmark dynamic of complex adaptive systems. The governance translation is system-level pattern recognition without central aggregation — distributed mechanisms that identify patterns as they arise from local interactions, enabling real-time response rather than delayed centralised processing.

3.4 Adaptive Feedback

Adaptive feedback in swarm systems refers to the continuous modification of agent behaviour in response to environmental signals and consequences of prior actions (Ashby, 1956; Holland, 1992). The governance translation is continuous environmental sensing and rule adjustment — governance systems that monitor conditions continuously and adjust governance rules in response, replacing periodic review cycles and fixed policy updates.

Table 1. Swarm Intelligence Principles and Governance Design Rule Translations

Swarm Principle	SI Foundation	Governance Design Rule	STRUM Element
Stigmergy	Indirect coordination through environmental traces — agents modify their shared environment and subsequent agents respond to those traces rather than to direct instruction (Grassé, 1959; Heylighen, 2016)	Signal-based indirect coordination: shared information traces that distributed units read and respond to, eliminating the need for top-down instruction	Signals
Self-Organisation	Coherent global patterns emerge from local interactions among agents within shared constraints, without external direction (Camazine et al., 2001; Bonabeau et al., 1999)	Autonomous unit decision-making within shared constraints: authority distributed to the point at which information is available; constraint architecture maintains coherence	Trust / Unity
Emergence	Novel system-level properties arise from local interactions; cannot be predicted from individual component properties; no global blueprint required (Holland, 1998; Anderson, 1999)	System-level pattern recognition without central aggregation: distributed sensing identifies emergent patterns in real time, replacing hierarchical reporting chains	Resonance
Adaptive Feedback	Continuous modification of agent behaviour in response to environmental signals and consequences of prior actions, maintaining effectiveness across variability (Ashby, 1956; Holland, 1992)	Continuous environmental sensing and rule adjustment: governance parameters updated in real time in response to signals, replacing periodic review cycles	Measure

Table 2 synthesises the derivation logic of the SGF: from each SI principle, a governance design rule is derived, which is in turn formalised as a STRUM element. The four principles collectively produce five STRUM elements (self-organisation generates both Trust and Unity, addressing the constraint architecture at the unit level and system level respectively).

Table 2. Synthesising Framework: *SI Principles* → *Governance Design Rules* → *STRUM Model* → *SGF*

<i>SI Principles</i>		<i>Governance Design Rules</i>		<i>STRUM Elements</i>		<i>SGF</i>
Stigmergy	→	Signal-based indirect coordination	→	Signals		Swarm Governance Framework (SGF)
Self-Organisation	→	Autonomous decision-making within shared constraints	→	Trust / Unity		
Emergence	→	System-level pattern recognition	→	Resonance		
Adaptive Feedback	→	Continuous sensing and rule adjustment	→	Measure		

4 The Swarm Governance Framework and the STRUM Model

4.1 Framework Overview

The Swarm Governance Framework (SGF) is introduced as an original theoretical contribution: the first governance framework derived from swarm intelligence principles and designed to address the limitations of hierarchical governance under conditions of organisational complexity. The SGF is not a modification of existing practitioner frameworks but a framework built from first principles — from the mechanisms through which complex adaptive systems generate distributed, adaptive coordination. The SGF is organised around the four governance design rules identified in Section 3, each operationalising a swarm principle into a concrete organisational design specification.

4.2 The STRUM Model

The four design rules of the SGF are formalised through the STRUM model — a five-element design specification that translates governance rules into actionable organisational design parameters. STRUM stands for Signals, Trust, Resonance, Unity, and Measure.

Signals formalises the stigmergy principle. It specifies that the governance architecture must include explicit mechanisms for generating, propagating, and reading shared signals across the organisation — environmental traces that allow distributed units to adjust behaviour in response to the collective state without explicit instruction. Trust formalises the self-organisation principle at the unit level. It specifies that autonomous unit decision-making requires a trust architecture — shared constraints, norms, values, and accountability mechanisms that define the space within which units may act independently. Resonance formalises the emergence principle. It specifies mechanisms for detecting and amplifying significant patterns from distributed interactions without routing them through centralised processing. Unity formalises the self-organisation principle at the system level — the mechanisms through which the constraint architecture is maintained and updated as a coherent shared property of the organisation as a whole. Measure formalises the adaptive feedback principle: continuous measurement systems that monitor environmental conditions and governance action consequences, feeding back into rule adjustment in real time.

4.3 Unity, Trust, and the Governance of Governance

Trust and Unity are the two STRUM elements most directly concerned with the constraint architecture and are closely interdependent in practice. Trust specifies the constraint architecture as experienced by individual units — the scope of autonomous action permitted to each unit and the conditions under which decisions must be escalated. Unity specifies the mechanisms through which that constraint architecture is maintained as a coherent shared property of the system as a whole. The relationship between them over time constitutes the governance of the governance: as the organisation evolves — through unit expansion, environmental change, or the progressive integration of AI agents — Trust-level constraints applied to individual units require updating, and it is the Unity mechanisms that determine how those updates are formulated and propagated consistently across the system. Without effective Unity, Trust degrades into locally customised constraint sets, producing constraint drift — a failure mode in which individual units operate under constraints that have diverged from a shared standard, generating incoherent collective behaviour even where each unit is functioning correctly within its own local rules. Without a defined Trust architecture, Unity has no coherent system to maintain. Their interaction defines the long-run adaptive capacity of the SGF.

4.4 Boundary Conditions

The SGF is most applicable in environments characterised by high complexity, rapid change, and informational distribution — where centralised governance fails because the information required for effective decision-making is distributed and the rate of change exceeds hierarchical adaptive speed. The framework's applicability diminishes in stable, compliance-intensive contexts where standardised governance and adherence to fixed protocols are the primary coordination requirements. The SGF is positioned as a governance architecture for complexity, not a universal replacement for all governance forms.

5 Theoretical Differentiation

Holacracy (Robertson, 2015) distributes authority across self-organising circles and formalises governance through constitutionally defined roles. While sharing the SGF's commitment to distributed authority, holacracy is practitioner-derived and does not claim theoretical derivation from the mechanisms of biological collective behaviour. Its design rules are based on practice-based intuition rather than a validated theoretical account of how distributed systems produce adaptive coordination. The SGF, by deriving its design rules from swarm intelligence theory, provides explanatory and generative power that holacracy lacks: it can explain why signal-based coordination outperforms instruction-based coordination under complexity, predict conditions under which autonomous decision-making produces system-level coherence, and generate design rules for novel governance contexts by applying the underlying principles.

Complexity-based management approaches (Anderson, 1999; Stacey, 1996; Snowden & Boone, 2007) correctly identify that distributed, adaptive coordination is required under complexity but operate at the level of strategic orientation rather than concrete design rules. The SGF shifts from the descriptive and strategic level of complexity theory to the prescriptive and design level, providing a concrete specification — the STRUM model — for how organisations should structure their governance architecture for complex environments. Research on new forms of organising (Puranam et al., 2014) identifies the design of coordination mechanisms as a central challenge for distributed organisations; the STRUM model addresses this challenge with a theoretically grounded design specification.

The relationship between the SGF and dynamic capabilities theory (Teece, 2007) is one of operationalisation rather than mere alignment. Teece identifies three microfoundations through which organisations sustain performance under volatility: sensing — scanning the environment for threats and opportunities; seizing — mobilising resources rapidly in response; and reconfiguring — transforming and realigning assets as conditions change. The STRUM model operationalises each of these microfoundations as a governance design parameter. Resonance institutionalises the sensing capability: rather than relying on a central scanner, emergent patterns are detected and surfaced from distributed interactions across the organisation in real time. Trust and Unity together institutionalise the seizing capability: a clear autonomous action architecture (Trust) supported by coherent constraint maintenance (Unity) removes the escalation delays that prevent rapid resource deployment in hierarchical systems. Measure institutionalises the reconfiguration capability: continuous feedback-driven parameter adjustment replaces the periodic review cycles through which hierarchical organisations attempt, and frequently fail, to reconfigure. Where dynamic capabilities theory identifies what organisations must be able to do, STRUM specifies how the governance architecture must be designed to make them possible.

6 Governing the Agentic Enterprise

6.1 The AI Coordination Challenge

The emergence of agentic AI systems — agents capable of autonomous action, tool use, and multi-step task execution — poses a governance challenge structurally analogous to the problem the SGF is designed to address (Wooldridge & Jennings, 1995). As organisations adopt AI agents for complex operational functions, they face the same core challenge: how do distributed, autonomous agents produce coherent collective behaviour without central direction? Growing bodies of AI governance research identify this coordination problem as a primary challenge for agentic enterprise deployment (Dafoe, 2018; Floridi et al., 2018), but existing frameworks address configuration parameters primarily at the infrastructure or compliance level, without a theoretical account of why specific configurations produce coherent collective behaviour.

6.2 STRUM as AI Agent Configuration Specification

The five elements of the STRUM model map directly onto the configuration parameters required for multi-agent AI systems. Signals specifies the shared environmental state architecture — the logging and observability layer through which agents record action traces that other agents and human operators can read. Trust specifies the permission and constraint architecture — the scope within which agents may act autonomously without human approval. Resonance specifies cross-agent pattern detection and escalation mechanisms — observing agent population behaviour as a whole to detect emergent patterns that per-agent monitoring cannot identify. Unity addresses constraint drift across the agent population as agents are updated, expanded, or reconfigured over time. Measure specifies the performance feedback architecture through which agent behavioural parameters are updated in response to outcome signals.

Table 3. STRUM Elements and AI Agent Configuration Parameters

STRUM Element	Governance Function	AI Agent Configuration Parameter
Signals	Signal-based indirect coordination: distributed units read and respond to shared environmental traces	Logging and observability architecture: mechanisms through which agents record action traces that other agents and human operators can read and respond to
Trust	Autonomous unit decision-making within a shared constraint architecture	Permission and constraint architecture: the scope and boundaries of autonomous agent action, defining what agents may and may not do independently
Resonance	System-level pattern recognition without central aggregation	Pattern-surfacing and escalation logic: mechanisms for detecting and surfacing significant emergent patterns across the agent population
Unity	Shared constraint maintenance — governance of the governance rules across the system	Governance rule consistency mechanisms: ensuring the constraint architecture is maintained coherently and updated consistently across the agent population
Measure	Continuous environmental sensing and adaptive rule adjustment	Performance feedback and behavioural parameter update: mechanisms through which agents modify parameters in response to outcome signals and environmental change

6.3 STRUM as a Configuration Framework: Novelty and Application

The novelty of STRUM as an AI agent configuration framework lies in its derivation logic. Existing AI agent configuration approaches are developed from engineering requirements — the parameters technical infrastructure demands — or from safety and compliance requirements — the behaviours that must be constrained. Neither provides a theoretically complete account of what governance parameters a distributed agentic system requires to produce coherent collective coordination (Floridi et al., 2018). STRUM approaches the problem from a validated theoretical account of how distributed systems achieve collective coordination in nature, producing a configuration framework that is both complete — the five elements together address all dimensions of the distributed coordination problem — and principled — each element is grounded in a theoretical account of why that dimension is necessary.

STRUM provides practitioners with a three-stage design tool: a completeness specification at the design stage (an agentic system that has not addressed all five elements has an incomplete governance design); a structured evaluation framework at the audit stage; and a set of measurable governance effectiveness dimensions at the evaluation stage. The structural equivalence between STRUM's application to human governance and AI agent governance suggests that the governance challenge of agentic AI is an instance of the broader problem of governing distributed autonomous agents — one the SGF addresses prescriptively, positioning it as a unified governance framework for the integrated human-AI enterprise.

7 Operational Application of the STRUM Model

7.1 STRUM Operational Indicators

Operationalising the STRUM model requires organisations to identify concrete instantiations of each element within their governance architecture and to define at least one measurable operational indicator per element. Table 4 presents indicative metrics for each STRUM element.

Table 4. STRUM Operational Indicators and Example Metrics

STRUM Element	Governance Function	Operational Indicator	Example Metric
Signals	Signal-based indirect coordination	Quality and coverage of the shared information environment	Signal propagation latency; information trace uptake rate across units
Trust	Autonomous unit decision-making	Boundary clarity and appropriateness of autonomous action scope	Decision escalation rate; constraint violation frequency
Resonance	System-level pattern recognition	Sensitivity and accuracy of emergent pattern detection	Time-to-pattern-detection; cross-unit anomaly surfacing rate
Unity	Constraint architecture coherence	Consistency of governance rules across units and over time	Constraint divergence index; governance rule synchronisation lag
Measure	Adaptive feedback fidelity	Accuracy and latency of performance attribution and parameter adjustment	Feedback loop latency; behavioural parameter update frequency

As the indicators in Table 3 suggest, the five elements together constitute a diagnostic framework for governance architecture health. A system in which escalation rates are persistently high (Trust element underdeveloped), pattern detection is slow (Resonance element underdeveloped), or feedback loops are long (Measure element underdeveloped) can trace these symptoms directly to specific architectural gaps, enabling targeted governance redesign rather than undifferentiated governance reform.

7.2 Illustrative Case: Hypothetical Distributed Technology Organisation

To illustrate the SGF's practical application and surface characteristic trade-offs, consider a hypothetical distributed technology firm operating five autonomous product teams across geographically distributed locations in a high-VUCA environment. Under a STRUM-aligned governance architecture: the Signals element is instantiated through a continuously updated shared operational dashboard that all teams read and write, recording decision outcomes, resource states, and dependency flags in real time — enabling indirect coordination without inter-team meetings. The Trust element is operationalised through a published set of decision rights defining each team's autonomous action scope: teams may independently allocate internal resources and modify sprint priorities; cross-team resource allocation and customer-facing commitments require a defined escalation protocol.

The Resonance element is activated through a cross-team anomaly detection layer that monitors the shared dashboard for unusual patterns — declining adoption rates, repeated escalations from specific teams, unusual dependency concentrations — and surfaces them to a governance steering layer for collective attention. The Unity element is maintained through quarterly cross-team governance reviews in which the decision rights framework itself is

assessed and updated collectively, ensuring constraint coherence across teams. The Measure element operates through bi-weekly outcome monitoring tied to each team's operational parameters, enabling continuous fine-tuning of the constraint architecture.

The hypothetical case surfaces a key architectural trade-off: a Signals-heavy environment with rich information traces but underdeveloped Resonance mechanisms produces information overload rather than coordination, as teams are exposed to more environmental traces than they can process without pattern-surfacing support. This illustrates the interdependence of STRUM elements: selective adoption without attending to element interactions produces governance failure modes that are traceable to specific architectural gaps rather than to the framework as a whole.

8 Risk and Ethical Considerations

The SGF and STRUM model carry implementation risks and ethical considerations that practitioners must address explicitly in design. Implementation failure modes arise when STRUM elements are adopted partially. Signals without Resonance generates information noise rather than coordination: traces accumulate in the shared environment without pattern-detection mechanisms to surface meaningful signals for action. Trust without Unity produces constraint drift: individually defined decision rights diverge across units over time, producing incoherent collective behaviour even when individual units are functioning correctly within their local constraints (Dafoe, 2018). Measure without Trust produces performance surveillance rather than adaptive governance: continuous monitoring without genuine autonomy reduces to performance management, undermining the self-organisation principle on which the framework depends.

In human-AI enterprise contexts, the ethical risks require explicit governance design. Floridi et al. (2018) establish explicability and accountability as non-negotiable foundational requirements for ethical AI governance — requirements that apply not only to individual AI agents but to the governance system that configures and constrains them. The STRUM model must therefore be designed to ensure that each governance action, whether by a human unit or an AI agent, is traceable to a specific Trust-level constraint and attributable to a specific actor; without this traceability, the distributed architecture produces accountability gaps that are structurally more difficult to audit than the failures of hierarchical systems it replaces.

A second ethical risk concerns the governance of Unity itself. Dafoe (2018) identifies concentration of decision-making power as a primary AI governance risk: when the mechanisms that set and update rules for AI systems are held by a small group, the result is a form of structural power concentration that can be more durable and less visible than traditional authority hierarchies. The same risk applies to the Unity layer of the SGF: if the mechanisms through which the constraint architecture is updated are controlled by a narrow governance group, the framework reproduces central authority under a distributed label. Addressing both risks requires that the Unity layer itself be governed through inclusive, transparent, and auditable processes — making the governance of the governance a first-order design requirement, not an afterthought.

9 Design Science Research Positioning

This paper is situated within a Design Science Research (DSR) methodology, following Hevner et al. (2004). DSR is an established research paradigm focused on the creation and evaluation of purposeful artefacts — models, frameworks, methods — designed to address specific organisational problems. This paper represents the artefact construction phase of a

multi-stage DSR programme: the SGF and STRUM model are the constructed artefact, derived from a critical theoretical review of existing theoretical foundations. As an artefact construction phase, empirical validation of the SGF is outside the scope of this paper and constitutes the primary agenda for subsequent research phases — consistent with established DSR practice in which artefact construction and evaluation represent distinct and separately publishable research contributions (Hevner et al., 2004). Subsequent phases will involve empirical testing of the SGF and STRUM model across organisational contexts of varying complexity, and iterative refinement of the design rules and indicators identified in Section 7.

10 Conclusion

This paper has introduced the Swarm Governance Framework (SGF) and the STRUM model as an original theoretical contribution to organisational governance design. By conducting a critical theoretical review of swarm intelligence and applying its four core principles — stigmergy, self-organisation, emergence, and adaptive feedback — to the governance design problem, the paper has produced the first governance framework derived from the mechanisms of biological collective intelligence, accompanied by a formal five-element design specification with operational indicators and an illustrative application.

The SGF addresses a specific and persistent gap: the absence of a governance framework that is both theoretically grounded in the mechanisms of distributed adaptive coordination and capable of generating concrete design rules for organisational architecture. The STRUM model — Signals, Trust, Resonance, Unity, Measure — translates the SGF's design rules into actionable governance parameters applicable to both human organisations and AI agent coordination systems. The Trust-Unity interaction identified in Section 4.3 provides the dynamic through which the framework maintains coherence over time, while the risk and ethical considerations in Section 8 identify the conditions under which partial or unreflective implementation produces governance failure modes. Future research will proceed through the DSR evaluation and refinement phases, subjecting the SGF and STRUM model to empirical testing, and developing the measurement instruments identified in Table 3 for use in organisational audit and governance design.

References

- Anderson, P. (1999). Complexity theory and organization science. *Organization Science*, 10(3), 216–232. <https://doi.org/10.1287/orsc.10.3.216>
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, 23(1), 4–17. <https://doi.org/10.1016/j.intcom.2010.07.003>
- Bennett, N., & Lemoine, G. J. (2014). What a difference a word makes: Understanding threats to performance in a VUCA world. *Business Horizons*, 57(3), 311–317. <https://doi.org/10.1016/j.bushor.2014.01.001>
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press.
- Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-Organization in Biological Systems*. Princeton University Press.
- Dafoe, A. (2018). *AI governance: A research agenda*. Future of Humanity Institute, University of Oxford.

- Dorigo, M., & Birattari, M. (2007). Swarm intelligence. *Scholarpedia*, 2(9), 1462. <https://doi.org/10.4249/scholarpedia.1462>
- Dorigo, M., & Gambardella, L. M. (1997). Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1), 53–66. <https://doi.org/10.1109/4235.585892>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People — An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Grassé, P.-P. (1959). La reconstruction du nid et les coordinations inter-individuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. La théorie de la stigmergie [The reconstruction of the nest and interindividual coordination in *Bellicositermes natalensis* and *Cubitermes* sp. The theory of stigmergy]. *Insectes Sociaux*, 6(1), 41–80.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- Heylighen, F. (2016). Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research*, 38, 4–13. <https://doi.org/10.1016/j.cogsys.2015.12.006>
- Holland, J. H. (1992). *Adaptation in Natural and Artificial Systems*. MIT Press.
- Holland, J. H. (1998). *Emergence: From Chaos to Order*. Oxford University Press.
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *Proceedings of the IEEE International Conference on Neural Networks*, 4, 1942–1948. <https://doi.org/10.1109/ICNN.1995.488968>
- Laloux, F. (2014). *Reinventing Organizations*. Nelson Parker.
- Lawrence, K. (2013). *Developing Leaders in a VUCA Environment*. UNC Executive Development.
- Lichtenstein, B. B. (2000). Self-organized transitions: A pattern amid the chaos of transformative change. *Academy of Management Perspectives*, 14(4), 128–141. <https://doi.org/10.5465/ame.2000.3979820>
- Puranam, P., Alexy, O., & Reitzig, M. (2014). What's "new" about new forms of organizing? *Academy of Management Review*, 39(2), 162–180. <https://doi.org/10.5465/amr.2011.0436>
- Robertson, B. J. (2015). *Holacracy: The New Management System for a Rapidly Changing World*. Henry Holt.
- Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76.
- Stacey, R. D. (1996). *Complexity and Creativity in Organizations*. Berrett-Koehler.
- Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350. <https://doi.org/10.1002/smj.640>
- Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152. <https://doi.org/10.1017/s0269888900008122>