

GOVERNING AI-ASSISTED STRATEGY TOOLS: A SUBSTITUTABILITY– COMMITMENT FRAMEWORK AND THE ROLE OF EPISTEMIC CUSTODY

Michael Mncedisi Willie

Policy, Research and Monitoring, Council for Medical Schemes, South Africa

e-mail: m.willie@medicalschemes.co.za
(Corresponding Author indicated by an asterisk *)

ABSTRACT

Artificial intelligence now performs much of the analytical work that established strategy tools were designed to structure. Existing research asks which tools can be automated. This article asks a prior question: whether the authority and organisational value of a strategic instrument survive when its analytical core is machine-generated. Approach. The article is conceptual and follows a theory-construction design, synthesising the strategy-as-practice, AI-enabled decision-making and corporate governance literatures through a documented search and inclusion procedure. Framework. The Substitutability–Commitment Framework classifies strategic instruments along two independent dimensions: analytical substitutability, the extent to which a tool’s cognitive work is machine-executable, and commitment durability, the extent to which its output binds organisational resources. Crossing the dimensions yields four classes, each implying a distinct governance response. The framework is complemented by epistemic custody, defined as the explicit allocation of responsibility for producing, verifying, contesting and authorising strategic judgements. Contribution. The article separates substitutability from commitment, identifies three observable failure modes, maps thirteen widely used instruments to the framework, and advances seven refutable propositions. Epistemic custody is distinguished from accountability, human oversight and decision rights. Implications. Boards require custody arrangements rather than technical assurance alone. The withdrawal of South Africa’s Draft National Artificial Intelligence Policy in April 2026, following the discovery of fabricated references, is used solely to illustrate the framework’s plausibility. It is not evidence validating the framework, which remains untested.

Keywords: Artificial Intelligence; Strategy Tools; Strategic Decision-Making; Epistemic Custody; Corporate Governance

INTRODUCTION

Strategic management has long relied on analytical instruments to support environmental scanning, competitive analysis and organisational decision-making. Frameworks such as PESTEL, Porter's Five Forces, SWOT analysis, portfolio matrices, the Balanced Scorecard and the Business Model Canvas continue to shape strategic practice despite many having been developed several decades ago (Qehaja & Kutllovci, 2020; George, 2025; Seidl et al., 2024). Their longevity reflects more than familiarity. These instruments provide a structured way of organising information, facilitating strategic dialogue and supporting collective judgement. Research in the strategy-as-practice tradition argues that strategy is not simply something organisations possess but something managers actively construct through shared practices and analytical tools (Jarzabkowski et al., 2022; Kohtamäki et al., 2022; Seidl et al., 2024). Consequently, the value of a strategy instrument extends beyond its analytical output to the organisational process through which that output is produced.

The emergence of artificial intelligence (AI) has fundamentally altered the environment in which these instruments operate. Most established strategy frameworks were designed when analytical capacity was limited, information was costly to obtain and strategic analysis depended almost entirely on human expertise (George, 2025; Kemp, 2024; Krakowski et al., 2023). Their purpose was therefore to simplify complexity by reducing large volumes of information into manageable analytical structures. AI has largely removed that constraint. Large language models and other forms of generative AI can now synthesise evidence, generate strategic options and evaluate business scenarios in seconds. Recent studies demonstrate that AI can perform many cognitive tasks associated with strategic analysis at levels comparable with experienced managers under particular conditions (Csaszar et al., 2024). At the same time, evidence suggests that AI complements rather than replaces managerial expertise, increasing the importance of coordination, interpretation and governance (Aleksieva et al., 2026; Krakowski et al., 2023).

Although important, this question addresses only part of the challenge. The authority of a strategic instrument depends not only on the conclusions it generates but also on how those conclusions are produced and who is accountable for them. Much of the emerging literature therefore asks which established strategy tools can be performed by AI (Keding, 2021; Burke & Wolf, 2021; Shrestha et al., 2019; Simkute et al., 2024). A Five Forces analysis developed through executive deliberation represents both an assessment of industry structure and a record of collective organisational judgement. The same analysis generated by AI may reach similar conclusions while lacking the deliberative process that gives those conclusions organisational legitimacy. The critical question is therefore not whether AI can perform strategic analysis, but whether the strategic value of an instrument changes when its underlying cognitive work is transferred from people to machines. Some strategy instruments lose little through automation because their value lies primarily in information processing. Others retain their familiar appearance while losing much of the epistemic authority that originally justified their use (Burke & Wolf, 2021; Jarzabkowski et al., 2022; Kohtamäki et al., 2022).

This challenge also exposes an important tension within strategic management theory. The resource-based view argues that sustained competitive advantage depends on valuable, rare, imperfectly imitable and non-substitutable resources (Kemp, 2024; Krakowski et al., 2023). If organisations increasingly rely on the same foundation models, trained on similar datasets and prompted in comparable ways, strategic analysis itself may become increasingly homogeneous. Rather than creating differentiation, AI may encourage convergence in strategic thinking. This concern aligns with the concept of algorithmic monoculture, which demonstrates

that widespread reliance on common algorithmic systems can reduce diversity in decision-making even when those systems outperform individual alternatives (Kleinberg & Raghavan, 2021). Similarly, Brynjolfsson argues that the greatest value of AI lies not in replicating human judgement but in extending it through complementary capabilities (Brynjolfsson, 2022).

The governance response has developed rapidly but remains focused primarily on AI systems rather than strategic work itself. The European Union Artificial Intelligence Act introduced a risk-based regulatory framework in 2024 (European Union, 2024). The United States National Institute of Standards and Technology Artificial Intelligence Risk Management Framework provides guidance through the functions of govern, map, measure and manage (National Institute of Standards and Technology, 2023). ISO/IEC 42001 established the first certifiable management system standard for AI (International Organization for Standardization, 2023), while the revised OECD AI Principles provide a common governance vocabulary across jurisdictions (Organisation for Economic Co-operation and Development, 2024). These frameworks establish important governance principles but offer limited guidance on which strategic activities should be delegated to AI, which require continued human deliberation and where accountability should remain.

South Africa provides a timely illustration of this governance challenge. In April 2026, the Department of Communications and Digital Technologies withdrew the Draft National Artificial Intelligence Policy after fabricated academic references were identified in a document approved by Cabinet on 25 March 2026 and published on 10 April 2026; the withdrawal was announced on 26 April and gazetted on 12 June 2026 (Republic of South Africa, 2026; South African Government News Agency, 2026). Rather, it reflected a failure to establish clear responsibility for verifying AI-generated content before it entered a formal policy process. The episode demonstrates that the principal governance challenge may not be the capability of AI itself, but uncertainty regarding who remains accountable for knowledge generated through AI-assisted analysis. Fabricated citation is a measured property of these systems rather than an isolated malfunction: cross-disciplinary testing finds substantial proportions of model-generated references unverifiable (Mugaanyi et al., 2024), and fabrication rates vary with topic familiarity and prompt specificity (Linardon et al., 2025). The incident did not arise because AI was incapable of generating text.

This article addresses that gap by developing the Substitutability-Commitment Framework, which classifies strategic instruments according to two independent dimensions: the extent to which their underlying cognitive work can be substituted by AI and the degree to which their outputs commit organisational resources. The framework is complemented by the concept of epistemic custody, defined as the explicit allocation. Together, these concepts provide a practical basis for determining where AI can appropriately augment strategic management and where human accountability remains indispensable. responsibility for producing, verifying, contesting and authorising AI-assisted strategic judgements

Three research questions follow. RQ1 asks under what conditions the organisational authority of a strategy tool survives the automation of its analytical core. RQ2 asks how responsibility for producing, verifying, contesting and authorising machine-assisted strategic analysis should be allocated. RQ3 asks which observable failure modes arise when that responsibility is left unallocated.

The article makes four contributions. First, it develops a conceptual framework for classifying strategic instruments according to their substitutability and governance requirements. Second, it identifies three observable failure modes associated with AI-assisted strategy development and explains how they may be recognised in practice. Third, it proposes a board-level diagnostic together with a crosswalk that maps widely used strategy instruments

within the framework. Finally, drawing on South African governance arrangements and emerging international standards, it argues that many of the responsibilities needed to govern AI-assisted strategy already exist within established legal and corporate governance frameworks (Institute of Directors in South Africa, 2016; Republic of South Africa, 2008; Republic of South Africa, 2013). What has been missing is a practical framework that identifies where those responsibilities arise and how they should be exercised as AI becomes embedded in strategic management.

Terminology is used as follows throughout. AI-assisted describes a process in which a machine contributes to analytical work for which a person remains answerable. Machine-generated describes an artefact produced by a machine. Automated describes a task performed without human contribution at the point of execution. Substitution denotes the replacement of human cognitive work by machine execution. Augmentation denotes the combination of human and machine contribution to one task; substitution and augmentation are treated as distinct outcomes rather than poles of a continuum (Raisch & Krakowski, 2021). Epistemic custody, algorithmic monoculture and the substitutability–commitment distinctions are constructs developed in this article and are defined at first use.

LITERATURE REVIEW

What the Evidence on Machine-Generated Strategy Actually Shows

The empirical literature has moved quickly and now supports several firm conclusions. Csaszar et al. (2024), drawing on an accelerator programme and a start-up competition, find that large language models produce and assess strategies at a standard, and they map the consequences for search, representation and aggregation. Comparable to experienced practitioners, a finding extended by evidence that generative tools also shift how strategic decisions are evaluated (Doshi et al., 2025) and raise measured productivity in professional work (Brynjolfsson et al., 2025). Krakowski et al. (2023) establish that substitution and complementation occur together rather than in sequence, so AI destroys some sources of advantage while creating others simultaneously. Shrestha et al. (2019) offer the most directly prescriptive contribution, specifying when decision authority should be fully delegated to algorithms, hybridised in sequence, or aggregated across human and machine judgement.

Two qualifications recur, and both point the same way. On the direction of the effect, Alekseeva et al. (2026) find AI adoption associated with rising demand for managerial skill, with the emphasis moving from budgeting and planning towards coordination and stakeholder work; Raisch and Krakowski (2021) theorise this as an automation–augmentation paradox that organisations must manage rather than resolve. On the object of substitution, Brynjolfsson (2022) warns that pursuing machines which imitate people crowds out machines which extend them, with consequences for bargaining power. Applied to strategy tools, both suggest interesting design question is not how much of the task a machine can assume, but which division of labour leaves the resulting decision better warranted (Csaszar et al., 2024; Raisch & Krakowski, 2021; Shrestha et al., 2019).

Three Gaps

A third literature, largely disconnected from the first two, concerns what happens when machine output enters consequential processes unverified. Kleinberg and Raghavan (2021) show formally that reliance on a single algorithmic assessor can worsen aggregate outcomes even where that assessor is individually superior, because correlated errors no longer cancel. Vanneste and Puranam (2025) examine how trust in artificial agents forms and how perceptions

of agency shape willingness to rely on their output, which bears directly on whether verification happens at all. Simkute et al. (2024) documents how productivity gains from generative systems are offset where review capacity does not scale with output volume. This work has not been connected to the strategy tools tradition and making that connection is the task here.

Three unresolved problems follow. First, the field lacks a principled basis for deciding which instruments to delegate: existing prescriptive work specifies conditions for delegating decision authority (Shrestha et al., 2019) but does not distinguish tools whose value survives automation from tools whose value it destroys (Burke & Wolf, 2021; Seidl et al., 2024). Second, the commoditisation of analysis has not been reconciled with a theory of advantage predicated on heterogeneity, leaving the resource-based view and the evidence on machine-generated strategy in unresolved tension (Csaszar et al., 2024; Kemp, 2024; Kleinberg & Raghavan, 2021). Third, verification is treated as a technical property of systems rather than an allocation of accountability inside an organisation, which is why episodes such as the South African withdrawal get read as quality-control failures rather than governance-design failures (Institute of Directors in South Africa, 2016; Republic of South Africa, 2013; Republic of South Africa, 2026). Section 4 addresses all three.

Epistemic custody must be distinguished from adjacent constructs. Accountability concerns who answer for an outcome after the fact; custody concerns who is answerable for the epistemic status of an input before it is acted upon. Human oversight requires a person in the loop without specifying what that person must be able to establish. Responsibility allocation and decision rights assign authority over a choice, whereas custody assigns authority over the warrant for the evidence on which the choice rests. Assurance and auditability are retrospective and evidentiary, and the audit literature accordingly places its instruments after deployment (Kazim et al., 2021), contestability concerns standing to challenge a decision already taken. Custody is prospective and artefact specific.

Four adjacent literatures bear on the problem. Research on algorithm aversion establishes that reluctance to rely on machine output is systematic and varies with algorithm, individual, task and setting (Mahmud et al., 2022), which matters because a governance arrangement assuming uniform willingness to challenge machine output will misallocate the contesting role. Work on explainable AI shows that explanation for end users pursues understandability, trustworthiness, transparency, controllability and fairness, and that no single form satisfies all of them (Laato et al., 2022), a custodian who cannot interrogate an output cannot answer for it. Research on organisational AI governance defines the field as the translation of ethical principles into firm-level process (Mäntymäki et al., 2022) and documents the practices and barriers involved (Papagiannidis et al., 2023), while a recent synthesis maps its unresolved gaps (Birkstedt et al., 2023). Evidence on human–AI teaming qualifies the optimistic reading: combinations improve performance on content creation but lose performance on decision tasks (Vaccaro et al., 2024). Across all four, the object of governance is the AI system and its lifecycle. None addresses the analytical artefact once it enters a strategic process, which is the gap this construct occupies. the value of assistance varies with the experience of the recipient (Wang et al., 2023), and AI is better studied as part of a system of work than as a discrete instrument (Anthony et al., 2023).

RESEARCH METHOD

This is a conceptual article, and its method is theory construction rather than hypothesis testing. It begins from an inconsistency between two established bodies of work: the strategy tools tradition, which holds that a tool's authority depends on its process of production (Burke

& Wolf, 2021; George, 2025; Kohtamäki et al., 2022), and the evidence that machines now produce those outputs to professional standard (Alekseeva et al., 2026; Csaszar et al., 2024; Krakowski et al., 2023). Where two literatures generate incompatible implications, conceptual work is the appropriate response, and the resulting framework is assessed against the usual criteria: whether it explains what existing apparatus cannot, whether its constructs are distinguishable and observable, and whether it yields propositions capable of refutation (Doshi et al., 2025; Raisch & Krakowski, 2021; Shrestha et al., 2019).

The framework is illustrated rather than tested. Illustration draws on documented cases, chiefly the withdrawal of South Africa's draft AI policy in April 2026, reported consistently across the Government Gazette record and multiple independent outlets (Department of Communications and Digital Technologies, 2024; Republic of South Africa, 2026; South African Government News Agency, 2026), and on regularities established in the studies reviewed above (Csaszar et al., 2024; Kleinberg & Raghavan, 2021; Krakowski et al., 2023). No claim of statistical generalisation is made. The South African material serves two purposes: it grounds an abstract argument in a setting where consequences were observable and severe, and it allows the governance implications to be connected to specific statutory duties rather than to regulation in prospect (Institute of Directors in South Africa, 2016; Republic of South Africa, 2008; Republic of South Africa, 2013). Section 7 states the limits of this approach and specifies the empirical work needed to test the propositions.

This article is conceptual and adopts a theory-construction design in which the framework is derived from an identified inconsistency between literature rather than induced from data. The derivation is set out so that the dimensions can be assessed as theoretically generated rather than author imposed.

Literature was identified through Scopus, Web of Science and Google Scholar, searched between January and March 2026. Search strings combined strategy tool terms with artificial intelligence terms, and separately with governance and accountability terms. Coverage was restricted to peer-reviewed articles in English, supplemented by regulatory instruments where these were the primary source for a normative claim. Sources were included where they addressed the production or use of strategic analysis, the substitution of analytical labour, or responsibility for machine-generated content, and excluded where AI appeared only as a technical topic without bearing on strategic decision-making. Backward and forward citation searching was applied until no further relevant work emerged.

The two dimensions were derived by asking, of each retained source, what property of a strategy tool its argument turned on. Two recurred and were consistently conflated. Analytical substitutability came from the automation literature; commitment durability from the strategy-as-practice and governance literatures. Three alternatives were rejected: task complexity, which correlates with substitutability without adding governance implications; data availability, an antecedent rather than a property of the tool; and organisational level, since instruments at the same level differ on both retained dimensions.

Both dimensions carry observable indicators so that a classification can be checked rather than asserted. Substitutability is judged high where the task is bound, the information it requires is documented, and performance is checkable against an external referent. Commitment is judged binding where the output triggers resource allocation and that allocation either cannot be reversed within one planning cycle or carries reputational or contractual cost on reversal.

The thirteen instruments were selected because they appear in standard strategic management texts and in the studies reviewed, and because together they span both dimensions. Where an instrument satisfied the criteria for more than one quadrant it was recorded as splitting

rather than forced into a single cell, and the split is reported in Table 3. Classification was performed by the author alone; the absence of independent coding is stated as a limitation. The South African case was selected on three criteria: the failure is documented in the public record and therefore verifiable; it occurred where analysis converted into binding commitment; and it arose in a jurisdiction with codified directors' duties. It is used to illustrate plausibility. A single case cannot validate a framework or establish three failure modes, and no such claim is made.

RESULTH AND DISCUSSION

The Substitutability–Commitment Framework Two Dimensions the Literature Has Run Together

The framework rests on two dimensions that existing work has treated as one. The first, analytical substitutability, is how far from the cognitive work a tool performs can be executed by machine at or above the standard a competent analyst would reach. It is high where the task is bounded, the relevant information documented and performance checkable against an external referent, which describes most environmental scanning and much comparative analysis (Csaszar et al., 2024; George, 2025; Shrestha et al., 2019). It is low where the task requires constructing a novel representation of an ambiguous situation, since representation, unlike search, resists delegation (Csaszar et al., 2024; Doshi et al., 2025; Kim et al., 2024). Substitutability belongs to the task rather than the tool's name, and it moves as capability moves, which is why the framework is applied recursively (Brynjolfsson, 2022; Krakowski et al., 2023). The second dimension, commitment durability, is how far the output binds resources in ways costly or impossible to reverse. It is low where a decision can be revisited cheaply, such as a periodic scan of the competitive environment, and high where it commits capital, capability or reputation over a horizon in which reversal is expensive (George, 2025; Kemp, 2024; Krakowski et al., 2023). This captures what the dynamic capabilities literature identifies as path dependence: the weight of a strategic choice lies less in its analytical content than in the options it closes off (George, 2025; Krakowski et al., 2023). The dimensions are independent, and that independence is the framework's structural claim. A task may be highly machine-executable and still bind the firm irreversibly, and it is exactly this combination where governance is hardest (Kleinberg & Raghavan, 2021; Shrestha et al., 2019; Vanneste & Puranam, 2025).

Four Classes of Instrument

Crossing the dimensions produces four classes, shown in Figure 1 and specified in Table 1. The classification is not ranking. Each class describes a different relationship between machine capability and human accountability, and each requires a different design response (Burke & Wolf, 2021; Raisch & Krakowski, 2021; Shrestha et al., 2019).




Figure 1. The Substitutability–Commitment Framework

Instruments are located by how far their cognitive work is machine-executable and how irreversibly their output binds resources. Epistemic custody applies in all four quadrants; the custody deficit is most damaging in Quadrant II, where machine-generated analysis supports binding commitment. Source: Author’s construction.

Quadrant II, augmented deliberation, is the contested zone and the framework’s main concern. Scenario generation, portfolio choice, business model design and valuation are substantially machine-executable yet issue in binding commitment (Qehaja & Kutlllovci, 2020; Csaszar et al., 2024; George, 2025). The appropriate division of labour has the machine generating and evaluating options at a scale no team could match, while humans retain the commitment decision and the obligation to explain it (Raisch & Krakowski, 2021; Shrestha et al., 2019). This is where the evidence of Csaszar et al. (2024) bites hardest, and where custody becomes indispensable, because machine-generated analysis arrives without the deliberative record that ordinarily makes a commitment defensible (Burke & Wolf, 2021; Vanneste & Puranam, 2025).

Table 1. The four classes of strategic instrument and the design response each implies

Quadrant	Substitutability	Commitment	Representative instruments	Design response
I. Continuous instrumentation	High	Weak	PESTEL scanning, competitor tracking, market sizing, benchmarking	Dissolve the periodic exercise into standing monitoring; specify attention triggers
II. Augmented deliberation	High	Binding	Scenario generation, portfolio choice, business model design, valuation	Machine generates options at scale; humans own commitment and

Quadrant	Substitutability	Commitment	Representative instruments	Design response
				must be able to explain it
III. Deliberative ritual	Low	Weak	SWOT and TOWS workshops, alignment sessions, planning away-days	Protect the deliberative process; treat machine output as input, never as product
IV. Irreducible judgement	Low	Binding	Purpose and mandate, ethical trade-offs, capability bets under uncertainty	Reserve explicitly from delegation; record the reservation and its basis

Substitutability is a property of the task rather than the tool's label, so classifications should be revisited as machine capability advances.

Source: Author's construction.

Epistemic Custody

The two dimensions establish where machine capability may be deployed but say nothing about when its output may be relied upon. Epistemic custody fills that gap. It is defined here as the identifiable allocation of three obligations: to verify the inputs and factual claims on which an output rests, to contest that output before it becomes binding, and to answer for the decision that follows. Custody is an organisational allocation rather than a technical property, and it can be absent even where the system performs well (Burke & Wolf, 2021; Shrestha et al., 2019; Vanneste & Puranam, 2025). It runs across all four quadrants, though what it requires differs: data provenance in Quadrant I, a named reviewer and a contestation point in Quadrant II, protection of the deliberative process in Quadrant III, and discharge by the accountable person in Quadrant IV (Institute of Directors in South Africa, 2016; Raisch & Krakowski, 2021; Republic of South Africa, 2008).

Three Failure Modes

Custody deficit occurs when machine-generated material enters a process faster than it can be verified, so verification degrades into formality while accountability remains nominally assigned. Its signature is the appearance, in consequential documents, of claims nobody identifiable checked. South Africa's draft AI policy is a documented instance: the tool worked, the process assigned responsibility on paper, and in a Cabinet-approved document a reported ten per cent or more of the academic references were nonetheless fabricated (Department of Communications and Digital Technologies, 2024; Republic of South Africa, 2026; South African Government News Agency, 2026). The Minister's own diagnosis, that citations were included without proper verification, describes custody deficit rather than model failure (Republic of South Africa, 2026; Rest of World, 2026; Vanneste & Puranam, 2025).

Applying The Framework

Figure 2 sets out the procedure, which is recursive rather than linear because substitutability moves. An instrument correctly placed in Quadrant III today may become partly machine-executable within a product cycle, and a classification treated as permanent will be wrong before anyone reviews it (Brynjolfsson, 2022; Csaszar et al., 2024; Krakowski et al., 2023). The first three steps establish the division of labour and the custody allocation; the last two institutionalise monitoring and reclassification, without which the framework becomes

another static artefact of the kind it exists to prevent (Burke & Wolf, 2021; George, 2025; National Institute of Standards and Technology, 2023).




Figure 2. Applying The Framework: A Five-Step Recursive Procedure

Steps 1 to 3 establish classification, division of labour and custody; steps 4 and 5 institutionalise monitoring for the three failure modes and periodic reclassification as machine capability advances. Source: author's construction.

Table 2. Diagnostic Instrument Derived from The Framework, Designed to Be Answerable by A Board Committee or Strategy Function Without Specialist Technical Knowledge

Diagnostic question	Failure mode indicated	Remedial action
Can we name the person who verified the factual claims in this output?	Custody deficit	Assign verification before the output enters any binding process
Did anyone contest this output before it became binding?	Custody deficit	Establish a contestation point with standing to reject
Would this artefact bind anyone who did not help produce it?	Substitutability illusion	Restore the deliberative process; reclassify to Quadrant III
Could a competitor using the same tools reach this conclusion?	Strategic monoculture	Add proprietary data or judgement; treat output as baseline, not answer
Has verification capacity grown with output volume?	Custody deficit	Constrain generation to the rate at which review can absorb it
Is any Quadrant IV determination resting on unattributable analysis?	Custody deficit at board level	Reserve explicitly; record the basis of the reservation
When was this classification last revisited?	Stale classification	Re-run steps 1 to 3 on a defined cycle

Source: Author's Construction

Propositions

Stating refutable propositions is what separates a framework from a taxonomy (Doshi et al., 2025; Raisch & Krakowski, 2021; Shrestha et al., 2019). On classification, P1 holds that firms automating Quadrant III instruments will show lower subsequent implementation of the resulting strategy than firms retaining the deliberative process, controlling for analytical quality (Burke & Wolf, 2021; Jarzabkowski et al., 2022; Kohtamäki et al., 2022), P2 holds that automating Quadrant I instruments yields efficiency gains with no measurable effect on strategic commitment (Csaszar et al., 2024; George, 2025; Shrestha et al., 2019). On custody, P3 holds that the incidence of unverified claims rises with the ratio of machine-generated

content to available review capacity rather than with volume as such (Republic of South Africa, 2026; Simkute et al., 2024; Vanneste & Puranam, 2025), P4 holds that organisations with a named contestation point will show fewer such incidents than those relying on diffuse review, whatever the sophistication of their tools (National Institute of Standards and Technology, 2023; Republic of South Africa, 2008; Shrestha et al., 2019).

On advantage, P5 holds that, observable in the similarity of disclosed strategic priorities strategic positions within an industry will converge as the proportion of firms using common foundation models rises, advantage depending on how AI is situated within a firm rather than on access to it (Kemp, 2024; Kleinberg & Raghavan, 2021; Krakowski et al., 2023). P6 holds that firms combining machine analysis with proprietary data unavailable to rivals will retain performance heterogeneity while those relying on machine analysis of public information will not, which follows directly from the requirement that advantage-generating resources be rare and imperfectly imitable (Brynjolfsson, 2022; Kemp, 2024; Krakowski et al., 2023). On context, P7 holds that returns to Quadrant II augmentation are conditional on data infrastructure and verification capacity, so that where both are constrained the observed returns will be lower and custody deficit more frequent (Alekseeva et al., 2026; Department of Communications and Digital Technologies, 2024). That proposition is directly testable in South Africa, and Section 6 takes it up.

A CROSSWALK ACROSS THE CANONICAL TOOLKIT

The Scanning Family

The Resource and Portfolio Family

The Deliberative Family

The Measurement and Commitment Family

The crosswalk is presented in Table 3. Each instrument is located by the analytical task it performs rather than by its label, with the quadrant, the expected line of development and the element that must remain human stated for each. The narrative walkthrough of each instrument family has been removed to meet the page limit, since Table 3 carries the same content.

Table 3. Crosswalk Between the Canonical Toolkit and The Substitutability–Commitment Framework

Instrument	Analytical core	Quadrant	How it evolves	What must remain human
PESTEL / environmental scanning (George, 2025)	Structured gathering of macro-environmental signals	I	Dissolves into continuous automated sensing; the periodic report disappears	Deciding what is material and what triggers attention
Five forces (George, 2025)	Assessment of industry structure and bargaining relations	I → II	Populated and updated by machine in near real time	Defining the industry boundary, which governs everything the analysis shows
Value chain (George, 2025)	Mapping of activities and their cost and value contribution	I	Automated from operational and financial data	Judging which activities are

Instrument	Analytical core	Quadrant	How it evolves	What must remain human
				strategically distinctive
Resource audit / VRIO (Kemp, 2024)	Inventory and assessment of resource positions	II	Inventory automated; assessment machine-assisted	Judging imitability against what rivals could actually replicate
Ansoff matrix	Generation of product and market growth options	II	Option sets generated at scale and evaluated against evidence	Committing capital and carrying the consequences
Growth-share matrix (George, 2025)	Allocation of resources across a business portfolio	II	Continuous rather than annual repositioning	Resource commitment; resisting convergence on the industry baseline
SWOT / TOWS (Qehaja & Kutllovci, 2020)	Systematic matching of internal and external factors	III	Automating the output is proposed to erode the tool's function, which lies in the deliberation	The deliberation itself, which is what creates commitment
McKinsey 7S (Qehaja & Kutllovci, 2020)	Diagnosis of organisational alignment, hard and soft	III → IV	Hard elements measurable; soft elements resist external assessment	Participation of those assessed; authority to prompt change
Scenario planning (Qehaja & Kutllovci, 2020; George, 2025)	Construction of futures to test governing assumptions	II and III	Generation migrates to machine; engagement stays human	Confronting the board with the possibility that its assumptions are wrong
Balanced scorecard (George, 2025)	Translation of strategy into a limited set of measures	I and IV	Tracking fully automated; selection is not	Choosing what counts as important, which commits behaviour
Business model canvas (Qehaja & Kutllovci, 2020)	Explicit statement of a firm's value logic	II	Configurations generated and stress-tested by machine	Commitment to a configuration and its foreclosed alternatives
Strategy canvas (Qehaja & Kutllovci, 2020)	Comparison of value curves against industry convention	II	Machine-assisted, but highly exposed to convergence	Departing from the baseline the models themselves produce
Purpose, mandate, ethical trade-offs	Determination of what the organisation is for	IV	Machine may inform; it cannot hold the determination	The accountable person; in South African law, non-delegable

Quadrant assignments describe the dominant analytical task and should be revisited as capability advances. Several instruments split across quadrants, which is itself the diagnostically useful result.

Source: Author's construction.

What the Crosswalk Reveals

THE SOUTH AFRICAN SETTING

Governance Without Waiting for Legislation

Why the framework binds differently in a middle-income economy

South African conditions make that condition binding rather than theoretical, particularly in smaller firms where analysis and decision are rarely separated (Wyk, 2023).

Discussion

What this Changes Theoretically

The framework's principal claim is that substitutability and commitment are independent, and that treating them as a single dimension running from manual to automated has hidden the cases that matter. The literature has organised itself around the substitutability axis, asking what machines can do (Csaszar et al., 2024; Krakowski et al., 2023; Shrestha et al., 2019). Adding commitment shows that the hard cases are not those of highest capability but those where high capability meets binding consequence, and that a second class of difficulty arises where low substitutability meets weak commitment, since those instruments are the most likely to be automated precisely because it appears costless (Burke & Wolf, 2021; George, 2025; Seidl et al., 2024). This connects the strategy-as-practice tradition to the AI and strategy literature, which have developed largely in parallel (Csaszar et al., 2024; Jarzabkowski et al., 2022; Kohtamäki et al., 2022).

The monoculture argument carries a sharper implication for the resource-based view than has been acknowledged. If advantage requires resources that are rare and imperfectly imitable (Kemp, 2024), and strategic analysis becomes available to all rivals at near-zero marginal cost and comparable quality (Csaszar et al., 2024), then analysis is disqualified as a source of advantage and the locus shifts to proprietary data, to the quality of the questions posed, and to judgement at the point of commitment (Brynjolfsson, 2022; Kleinberg & Raghavan, 2021; Krakowski et al., 2023). That is consistent with Krakowski et al. (2023), who find new heterogeneity emerging from human-machine interaction rather than from the machine, and with Brynjolfsson (2022) on the returns to augmentation over imitation. Firms competing on the sophistication of their analytical apparatus are competing on a dimension that is actively commoditising (Kemp, 2024; Kleinberg & Raghavan, 2021; Krakowski et al., 2023).

What this Changes in Practice and Policy

For practice, the framework replaces the prevailing question, how much of our strategy work can AI do, with a sequence a board can adjudicate: which quadrant is this task in, what division of labour follows, and who holds custody (Institute of Directors in South Africa, 2016; Republic of South Africa, 2008; Shrestha et al., 2019). The diagnostic in Table 2 is built for that setting and the crosswalk in Table 3 lets a strategy function locate its existing instruments without first running a classification exercise of its own (Csaszar et al., 2024; George, 2025; Seidl et al., 2024). The most consequential implication is counter-intuitive: the framework recommends constraining generation to the rate verification can absorb, so the correct response to a custody deficit is often to produce less rather than review faster (Republic of South Africa, 2026; Simkute et al., 2024; Vanneste & Puranam, 2025).

For policy, three implications follow. Regulatory attention should attach to the custody allocation rather than the technology, since the South African episode shows failure arising from process design rather than model defect (Republic of South Africa, 2013; Republic of South Africa, 2026; South African Government News Agency, 2026). Guidance to public entities should specify verification obligations for machine-assisted documents, which existing governance codes imply but do not make explicit (European Union, 2024; Institute of Directors

in South Africa, 2016; International Organization for Standardization, 2023). And the National AI Policy process, when resumed, would be strengthened by addressing the state's own use of AI as directly as it addresses regulation of AI in the market, since the state is a substantial deployer and the withdrawal illustrates the consequences (Department of Communications and Digital Technologies, 2024; Republic of South Africa, 2026).

The episode is not jurisdictionally peculiar. Within ASEAN, of which Indonesia is a member state, the ASEAN Guide on AI Governance and Ethics (ASEAN, 2024) and its generative-AI expansion (ASEAN, 2025) set out principles of accountability, transparency and human-centricity across the AI lifecycle. Like the South African material, both govern the system rather than the analytical artefact it produces. That shared omission is what the custody construct addresses, and it suggests the framework applies in any developing economy where binding decisions rest on machine-generated analysis.

For boards and executive teams, the framework converts an open question about capability into a sequence that can be minuted. The diagnostic in Table 2 is designed to be worked through by a board committee in a single sitting without specialist technical input. For each instrument in use, the committee establishes the quadrant, records the division of labour between machine and human, and names the individual holding custody. The custody holder is a named person rather than a function, because a function cannot be asked on what basis a claim is believed on. This sits within a wider re-examination of the board monitoring role under AI, which asks what directors must be able to establish rather than merely receive (Hilb, 2020; Kourabas & Tsang, 2025). Where no name can be supplied, the instrument should not proceed to a committing decision until one can.

Three implementation steps follow. First, the strategy function locates existing instruments using the crosswalk in Table 3, avoiding a classification exercise from first principles. Second, verification obligations attach to Quadrant II instruments specifically, since these combine high substitutability with binding commitment. Third, the classification is reviewed at a fixed interval, because substitutability moves. In South African practice these steps extend existing duties under the Companies Act and King IV rather than creating new ones.

Limits and What Should be Tested

Three limitations should be stated plainly. The framework is conceptual and unvalidated; its propositions are refutable but unrefuted, and the illustrative material establishes plausibility rather than support (Csaszar et al., 2024; Doshi et al., 2025; Shrestha et al., 2019). Its dimensions are presented as continual but operationalised only qualitatively, and measurement instruments, particularly for commitment durability, are needed before the propositions can be tested (Kemp, 2024; Krakowski et al., 2023; Raisch & Krakowski, 2021). And it treats substitutability as exogenous, whereas firms shape it through their own investment and data practices, so a fuller treatment would model the co-evolution of capability and classification (Brynjolfsson, 2022; Krakowski et al., 2023; National Institute of Standards and Technology, 2023). The crosswalk carries a further caveat: its assignments reflect current capability and the dominant task within each instrument, and both will shift (Csaszar et al., 2024; George, 2025; Shrestha et al., 2019).

Four lines of empirical work follow. Field study of strategy processes before and after generative tools were introduced would test P1 and P2 directly, and the strategy-as-practice tradition supplies the method (Jarzabkowski et al., 2022; Kohtamäki et al., 2022; Seidl et al., 2024). Content analysis of disclosed strategic priorities across firms in one industry would test P5 using similarity measures over time (Kemp, 2024; Kleinberg & Raghavan, 2021; Krakowski

et al., 2023). Survey or archival work relating verification capacity to unverified claims would test P3 and P4, with the South African public sector an accessible setting given the documented case (Republic of South Africa, 2013; Republic of South Africa, 2026; South African Government News Agency, 2026). Comparative work across jurisdictions with differing infrastructure and skills endowments would test P7 and establish whether custody deficit is general or concentrated where verification capacity is scarce (Alekseeva et al., 2026; Department of Communications and Digital Technologies, 2024).

Summary of the propositions and how they may be tested. P1 and P2 concern the survival of tool authority under automation and follow from the substitutability dimension; both are testable through comparative field study of strategy processes before and after generative tools were introduced. P3 and P4 concern custody and its failure, testable through survey or archival work relating verification capacity to the incidence of unverified claims in committing documents. P5 concerns analytical convergence, testable through content analysis of disclosed strategic priorities within an industry using similarity measures over time. P6 and P7 concern the governance response, requiring longitudinal work relating custody arrangements to decision quality and to recurrence of the three failure modes. None has been tested here; the framework should be treated as a source of hypotheses rather than a validated instrument.

CONCLUSION

The strategy toolkit was built for a world in which analysis was scarce. That scarcity is over, and the field has responded by asking what machines can now do. This article argued that the question, however natural, leads away from the problem. A strategy tool's output has authority because of how it was produced, so automating production changes the standing of the output even when the content is unchanged. Some instruments come through intact. Others are hollowed out while continuing to look exactly as they always did, which is what makes the failure hard to see.

The Substitutability–Commitment Framework offers a way of telling these apart. It sorts instruments by how machine-executable their cognitive work is and how irreversibly their output binds, giving four classes each with its own design response, and it adds epistemic custody as the condition that must hold across all four. Applied across the canonical toolkit, it shows that several familiar instruments split across quadrants, that the tools most readily automated are those that contributed least to advantage, and that the ones surviving automation intact are those whose value was never informational. Its central practical claim is that the constraint worth managing is verification rather than generation, and that in many organisations the right response to abundant machine analysis is to slow its entry into binding processes rather than accelerate its review.

Three claims should be kept apart. What the literature establishes is that machines now produce strategic analysis to professional standard, that reliance on a common assessor can worsen aggregate outcomes, and that a tool's authority derives from its process of production. What this article proposes, and has not demonstrated, is that substitutability and commitment are independent, that four classes follow, that epistemic custody is the condition holding across them, and that three failure modes arise when custody is unallocated. What must be tested is every one of the seven propositions. The South African withdrawal is consistent with the framework and was the occasion for developing it, but a single documented failure cannot establish a construct.

CONFLICTS OF INTEREST

The author declares no competing interests. The views expressed are his own and do not represent the position of the affiliated entity. No funding was received for this work.

REFERENCES

- Alekseeva, L., Azar, J., Giné, M., & Samila, S. (2026). Artificial intelligence adoption and the demand for managerial expertise. *Strategic Management Journal*, 1(1), 1–27. <https://doi.org/10.1002/smj.70099>
- Anthony, C., Bechky, B. A., & Fayard, A.-L. (2023). "Collaborating" with AI: Taking a system view to explore the future of work. *Organization Science*, 34(5), 1672–1694. <https://doi.org/10.1287/orsc.2022.1651>
- ASEAN. (2024). *ASEAN guide on AI governance and ethics*. https://asean.org/wp-content/uploads/2024/02/ASEAN-Guide-on-AI-Governance-and-Ethics_beautified_201223_v2.pdf
- ASEAN. (2025). *Expanded ASEAN guide on AI governance and ethics – generative AI*. <https://asean.org/wp-content/uploads/2025/01/Expanded-ASEAN-Guide-on-AI-Governance-and-Ethics-Generative-AI.pdf>
- Birkstedt, T., Minkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. *Internet Research*, 33(7), 133–167. <https://doi.org/10.1108/INTR-01-2022-0042>
- Brynjolfsson, E. (2022). The Turing trap: The promise and peril of human-like artificial intelligence. *Daedalus*, 151(2), 272–287. https://doi.org/10.1162/daed_a_01915
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Burke, G. T., & Wolf, C. (2021). The process affordances of strategy toolmaking when addressing wicked problems. *Journal of Management Studies*, 58(2), 359–388. <https://doi.org/10.1111/joms.12572>
- Csaszar, F. A., Ketkar, H., & Kim, H. (2024). Artificial intelligence and strategic decision-making: Evidence from entrepreneurs and investors. *Strategy Science*, 9(4), 322–345. <https://doi.org/10.1287/stsc.2024.0190>
- Department of Communications and Digital Technologies. (2024). *South Africa national artificial intelligence policy framework*. <https://www.dcdt.gov.za/sa-national-ai-policy-framework/file/338-sa-national-ai-policy-framework.html>
- Doshi, A. R., Bell, J. J., Mirzayev, E., & Vanneste, B. S. (2025). Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal*, 46(3), 583–610. <https://doi.org/10.1002/smj.3677>
- European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (artificial intelligence act)*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

- George, B. (2025). Towards purposeful strategic planning: A mixed research synthesis across disciplines. *Long Range Planning*, 58(4), 2–11. <https://doi.org/10.1016/j.lrp.2025.102563>
- Hilb, M. (2020). Toward artificial governance? The role of artificial intelligence in shaping the future of corporate governance. *Journal of Management and Governance*, 24, 851–870. <https://doi.org/10.1007/s10997-020-09519-9>
- Institute of Directors in South Africa. (2016). *King IV report on corporate governance for South Africa 2016*.
- International Organization for Standardization. (2023). *ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system*. <https://www.iso.org/standard/81230.html>
- Jarzabkowski, P., Seidl, D., & Balogun, J. (2022). From germination to propagation: Two decades of strategy-as-practice research and potential future directions. *Human Relations*, 75(8), 1533–1559. <https://doi.org/10.1177/00187267221089473>
- Kazim, E., Denny, D. M. T., & Koshiyama, A. (2021). AI auditing and impact assessment: According to the UK Information Commissioner’s Office. *AI and Ethics*, 1, 301–310. <https://doi.org/10.1007/s43681-021-00039-2>
- Keding, C. (2021). Understanding the interplay of artificial intelligence and strategic management: Four decades of research in review. *Management Review Quarterly*, 71(1), 91–134. <https://doi.org/10.1007/s11301-020-00181-x>
- Kemp, A. (2024). Competitive advantage through artificial intelligence: Toward a theory of situated AI. *Academy of Management Review*, 49(3), 618–635. <https://doi.org/10.5465/amr.2020.0205>
- Kim, H., Glaeser, E. L., Hillis, A., Kominers, S. D., & Luca, M. (2024). Decision authority and the returns to algorithms. *Strategic Management Journal*, 45(4), 619–648. <https://doi.org/10.1002/smj.3569>
- Kleinberg, J., & Raghavan, M. (2021). Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences*, 118(22), 1–7. <https://doi.org/10.1073/pnas.2018340118>
- Kohtamäki, M., Whittington, R., Vaara, E., & Rabetino, R. (2022). Making connections: Harnessing the diversity of strategy-as-practice research. *International Journal of Management Reviews*, 24(2), 210–232. <https://doi.org/10.1111/ijmr.12274>
- Kourabas, S., & Tsang, C.-Y. (2025). The board monitoring function: Artificial intelligence in the era of heightened accountability. *The Chinese Journal of Comparative Law*, 13, 1–29. <https://doi.org/10.1093/cjcl/cxaf013>
- Krakowski, S., Luger, J., & Raisch, S. (2023). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, 44(6), 1425–1452. <https://doi.org/10.1002/smj.3387>
- Laato, S., Tiainen, M., Islam, A. K. M. N., & Mäntymäki, M. (2022). How to explain AI systems to end users: A systematic literature review and research agenda. *Internet Research*, 32(7), 1–31. <https://doi.org/10.1108/INTR-08-2021-0600>

- Linardon, J., Jarman, H. K., McClure, Z., Anderson, C., Liu, C., & Messer, M. (2025). Influence of topic familiarity and prompt specificity on citation fabrication in mental health research using large language models: Experimental study. *Journal of Medical Internet Research*, 12, 1–8. <https://doi.org/10.2196/80371>
- Mahmud, H., Islam, A. K. M. N., Ahmed, S. I., & Smolander, K. (2022). What influences algorithmic decision-making? A systematic literature review on algorithm aversion. *Technological Forecasting and Social Change*, 175, 1–26. <https://doi.org/10.1016/j.techfore.2021.121390>
- Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics*, 2(4), 603–609. <https://doi.org/10.1007/s43681-022-00143-x>
- Mugaanyi, J., Cai, L., Cheng, S., Lu, C., & Huang, J. (2024). Evaluation of large language model performance and reliability for citations and references in scholarly writing: Cross-disciplinary study. *Journal of Medical Internet Research*, 26, 1–7. <https://doi.org/10.2196/52935>
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. <https://www.nist.gov/itl/ai-risk-management-framework>
- Organisation for Economic Co-operation and Development. (2024). *OECD AI principles overview*. <https://oecd.ai/en/ai-principles>
- Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2023). Toward AI governance: Identifying best practices and potential barriers and outcomes. *Information Systems Frontiers*, 25(1), 123–141. <https://doi.org/10.1007/s10796-022-10251-y>
- Qehaja, A. B., & Kutllovci, E. (2020). Strategy tools in use: New empirical insights from the strategy-as-practice perspective. *Management: Journal of Contemporary Management Issues*, 25(1), 145–169. <https://doi.org/10.30924/mjcmi.25.1.9>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- Republic of South Africa. (2008). *Companies Act 71 of 2008*. Government Printer.
- Republic of South Africa. (2013, November 26). *Protection of Personal Information Act*. Government Printer.
- Republic of South Africa. (2026, April 10). *Draft national artificial intelligence policy*. Government Printer.
- Rest of World. (2026, May 7). *Fact-check fail: When AI hallucinations derail governments*. <https://restofworld.org/2026/government-ai-hallucinations-south-africa-deloitte/>
- Seidl, D., Ma, S., & Splitter, V. (2024). What makes activities strategic: Toward a new framework for strategy-as-practice research. *Strategic Management Journal*, 45(12), 2395–2419. <https://doi.org/10.1002/smj.3668>
- Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66–83. <https://doi.org/10.1177/0008125619862257>

- Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2024). Ironies of generative AI: Understanding and mitigating productivity loss in human–AI interaction. *International Journal of Human–Computer Interaction*, 41(5), 1–12. <https://doi.org/10.1080/10447318.2024.2405782>
- South African Government News Agency. (2026, April 26). *Minister announces withdrawal of draft AI policy*. <https://www.sanews.gov.za/south-africa/minister-announces-withdrawal-draft-ai-policy>
- Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8, 2293–2303. <https://doi.org/10.1038/s41562-024-02024-1>
- Vanneste, B. S., & Puranam, P. (2025). Artificial intelligence, trust, and perceptions of agency. *Academy of Management Review*, 50(4), 726–744. <https://doi.org/10.5465/amr.2022.0041>
- Wang, W., Gao, G., & Agarwal, R. (2023). Friend or foe? Teaming between artificial intelligence and workers with variation in experience. *Management Science*, 70(9), 5753–5775. <https://doi.org/10.1287/mnsc.2021.00588>
- Wyk, I. V. (2023). Strategic decision-making in small and medium enterprises in South Africa. *Southern African Journal of Entrepreneurship and Small Business Management*, 15(1), 1–12. https://hdl.handle.net/10520/ejc-sajesbm_v15_n1_a684