Research Square

Preprints are preliminary reports that have not undergone peer review. They should not be considered conclusive, used to inform clinical practice, or referenced by the media as validated information.

From Practice to Performance:A Formally Verified Theory of OrganizationalCapability and Autonomy

Franny Philos Sophia

franny.philos.sophia@elanare.jp

Elanare Institute

Research Article

Keywords:

Posted Date: October 22nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11268483/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: No competing interests reported.

From Practice to Performance: A Formally Verified Theory of Organizational Capability and Autonomy

Abstract

This paper presents RCAS (Reference–Choice Adaptive System), a formal theory that decomposes the arrow from organizational practice to performance into explicit intermediate mechanisms. The theory makes three contributions. First, it provides a mechanism decomposition: each practice maps to state variables that connect to specific performance components. Second, it decomposes productive performance into a per-domain profile—capability ceiling, realized frontier, throughput volume, and resilience count—and proves that knowledge diffusion raises volume and resilience but cannot raise the ceiling, which expands only through capability creation. Third, it separates formal consequences from empirical claims: all theorems and countermodels are machine-checked in Lean 4 (840 declarations, zero `sorry`, standard axioms only), and a countermodel shows that structural self-management does not entail adequate performance.

Keywords: organizational theory · formal verification · team autonomy · capability diffusion · performance decomposition · Lean 4

1 Introduction

Organizational interventions are routinely evaluated by drawing a direct arrow from practice to performance. A firm adopts self-managing teams, agile methods, or knowledge-sharing routines, and the subsequent change in a performance metric is attributed to the intervention. The difficulty is that this arrow hides every mechanism through which the intervention could have acted. When performance improves, the organization cannot identify which intermediate condition was responsible; when it does not, the organization cannot identify which condition was absent.

The empirical literature has converged on a version of this problem. Ryu et al. (2022), pooling 415 effect sizes from 69 studies conducted on approximately 6,035 teams, find significant indirect effects of team autonomy on performance and attitudes via task-focused and relationship-focused team functioning. The mediating constructs, however, are survey-measured team processes and emergent states rather than formally derived mechanisms. A similar gap appears in the shared-leadership literature, where meta-analytic evidence on task

complexity as a moderator is mixed: Wang et al. (2014) find stronger shared-leadership effects on team effectiveness for more complex work, whereas D'Innocenzo et al. (2016), coding task complexity on a scale adapted from Wood (1986), find weaker effects on team performance for more complex tasks. The two studies differ in outcome definition and complexity coding, yet when performance is a single scalar, the source of the disagreement cannot be located.

This paper presents a formal theory called RCAS (Reference–Choice Adaptive System) that decomposes the practice-to-performance arrow into explicit intermediate mechanisms. The theory makes three contributions.

First, it provides a *mechanism decomposition*. Each organizational practice maps to a set of state variables—capability distribution, manager dependence, constraint structure, knowledge-transfer rate, engagement—and each state variable connects to a specific component of performance through a chain of definitions and theorems.

Second, it *decomposes performance* itself. Rather than treating organizational performance as a single number, RCAS defines three components per domain: a realized frontier F (the highest quality or complexity the team can currently achieve), a throughput volume V (how much the team can produce in parallel), and a resilience count N (how many members can be lost before the domain's capability drops below a required threshold). These three components have different causal drivers: capability diffusion raises V and N but cannot raise the potential ceiling; only capability creation—someone exceeding the current maximum—can expand the frontier. The model formalizes the connection from capability to output; the further connection from output to organizational outcome depends on whether the right goals were set, which lies outside the formal boundary.

Third, it *separates what follows from assumptions from what does not*. All formal results and countermodels reported in this article were machine-checked in Lean 4. Every theorem is accompanied by the assumptions from which it follows. Every countermodel is a constructed example showing that removing an assumption breaks the result. Every claim that falls outside the formal boundary is explicitly marked as an empirical hypothesis. To our knowledge, this is among the first organization-theory models to be machine-checked in a proof assistant; prior mechanizations in economics have addressed social choice theory (Nipkow 2009; Holliday et al. 2021) and individual decision theory (Li 2025; Li et al. 2025).

RCAS builds on several traditions without depending on any single one. In assuming that incentives are aligned and that the problem is informational, it belongs to the team-theory tradition of Marschak and Radner (1972); it adds capability distribution and diffusion dynamics. Its distinction between formal authority (who holds the decision right) and real authority (who has the information to exercise it) follows Aghion and Tirole (1997); RCAS contributes a dynamic model of how real authority shifts as front-line knowledge expands, drawing on Garicano's (2000) knowledge-hierarchy framework. Its treatment of constraints as objects with a burden of proof relates to the simple-rules literature (Davis et al. 2009). Its concept of leadership as a distributed function rather than a personal attribute is shared with complexity leadership theory (Uhl-Bien et al. 2007; Hazy 2006), though the formal model was constructed independently. RCAS does not model the hidden-effort regime studied by Holmström (1982); it analyzes a high-observability regime in which action-relevant signals are repeatedly shared, as in planning-and-daily-review environments. Rule updating is modeled as intentional improvement under a stay condition, rather than as the search-and-selection

process of evolutionary economics (Nelson and Winter 1982). The choice set available to each member, $A_i(\rho, R)$, is a related construct to the employee's "area of acceptance" in Simon's (1951) model of the employment relationship, though RCAS defines the set by outcome criteria rather than by the employer's authority.

What RCAS does *not* claim is important. The theory does not assert that autonomy improves performance. It identifies the intermediate mechanisms through which improvement could occur, derives the consequences that follow from explicit assumptions, proves where those consequences fail without the assumptions, and leaves the remaining links to empirical test.

2 Formal Approach and Epistemic Status

The formal content of this paper falls into five categories, which the reader should distinguish throughout.

A **definition** introduces a concept—a reference point, a choice set, a capability matrix, a performance component—and carries no empirical claim. It fixes the vocabulary of the model.

An **assumption** is a condition placed on the model's objects. For example, the knowledge-sharing dynamics assume that the transfer rate satisfies $0 \leq \eta \leq 1$ and that the top member's capability does not decrease. Assumptions are stated as fields of a structure or as hypotheses of a theorem. Whether a given assumption holds in a particular real organization is not determined by the model.

A **theorem** is a statement that follows from assumptions by a chain of deductive steps. Every theorem in this paper has been machine-checked in Lean 4 (v4.34.1, Mathlib v4.34.1): the proof assistant verifies that the conclusion follows from the stated premises and that no unstated premise has been used. The full source code is available as supplementary material.

A **countermodel** is a concrete construction—specific parameter values, specific initial conditions—under which a plausible-sounding claim fails. Countermodels serve two purposes. They show that a theorem's assumptions are not merely technical but load-bearing: remove one and the conclusion breaks. They also show that the theory does not embed success into its definitions: the formal model admits states in which self-management holds but performance does not.

An **empirical hypothesis** is a claim that the model's assumptions hold in some class of real organizations. The model proves that if those assumptions hold, certain consequences follow; it does not prove that the assumptions hold. Each empirical hypothesis is listed with the Lean definitions and theorems that would serve as its formal precondition.

This separation is the core methodological commitment of the paper. When we write that knowledge sharing cannot raise the capability ceiling, we mean that this follows as a theorem from the pure-sharing assumptions. When we write that RCAS practices increase the knowledge-transfer rate η, we mean that this is an empirical hypothesis that the formal model does not address.

The following notation is used throughout:

Symbol	Meaning
$\rho = (Y, B)$	Reference point: acceptability criterion Y and boundary constraints B
R	Team rule set
$A_i(\rho, R)$	Feasible action set for member i
$K,\ k_{id}$	Capability matrix; member i's proficiency in domain d
K_d^*	Capability ceiling: $\max_i k_{id}$
e_{id}	Engagement of member i in domain d, $\in [0, 1]$
F_d	Realized frontier: $\max_i g_d(k_{id}, e_{id})$
V_d	Throughput volume: $\sum_i e_{id} \cdot \varphi_d(k_{id})$
N_d	Resilience count: $\lvert\{i : k_{id} \cdot e_{id} \geq \theta_d\}\rvert$
θ_d	Domain threshold for capability
η	Knowledge-transfer rate, $\in [0, 1]$
$D(t)$	Manager dependence at time t
p	Dependence-absorption rate

3 Reference–Choice Organizational Architecture

3.1 Reference and genuine choice

The core construct of RCAS is a *reference point* $\rho = (Y(\rho), B)$, where $Y(\rho)$ defines an acceptability criterion on outcomes and B is a set of boundary constraints. Given a reference point ρ and a rule set R, each member i has a feasible action set $A_i(\rho, R)$ consisting of actions that satisfy both the acceptability criterion and the constraints.

A member has *genuine choice* when $|A_i| \geq 2$: at least two distinct actions are available. This is the minimum condition for autonomy to be meaningful.

Three organizational regimes are defined by the relationship between the reference point and the action set. Under *prescription*, the action set is a singleton: the member is told what to do and how to do it. Under *laissez-faire*, no reference point is set: there is no acceptability criterion against which actions are evaluated. Under *reference-choice*, a reference point is set (Y and B are specified) and the member has genuine choice.

Proposition 1 (P1: Non-prescription and non-laissez-faire)**.** *Reference-choice is neither prescription nor laissez-faire.*

Proposition 2 (P2: Distinct methods, same criterion)**.** *Under genuine choice, there exist at least two distinct actions that satisfy the same acceptability criterion. Outcome standardization does not require method standardization.*

Proposition 3 (P3: Standardization with autonomy)**.** *There exists an assignment of actions to members such that every member's action satisfies the acceptability criterion (outcome standardization) and at least two members use different methods (method diversity). Outcome standardization and method diversity are jointly satisfiable.*

3.2 Nested control architecture

Organizational structure is modeled as a three-layer nesting:

$$B_{\text{org}} \;\rightarrow\; R_{\text{team}} \;\rightarrow\; A_{\text{individual}}.$$

The organizational boundary constrains the team rule set, which in turn constrains the individual action set. A *soundness* condition requires that the nesting is coherent: every individual action permitted by the team rules falls within the organizational boundary. This condition is an explicit assumption, not a consequence of the definitions.

Countermodel 1 (`sound_not_automatic`). A three-layer structure can be defined in which the nesting is not sound: a team rule permits an action that falls outside the organizational boundary. Soundness must be actively maintained—it does not arise automatically from having layers.

3.3 Minimal sufficient constraints

A constraint set B is *admissible* if it satisfies the organization's requirements. A constraint set is *minimal sufficient* if it is admissible and no admissible constraint set with strictly lower burden exists.

Proposition 4 (C1: Existence)**.** *When the set of admissible candidates is finite, a minimal sufficient constraint exists.*

Proposition 5 (C2: Dominance)**.** *If two constraint sets are both admissible and one has strictly lower burden, the higher-burden set is not minimal sufficient.*

Proposition 6 (C3: Burden of proof)**.** *If a constraint set contains an element whose removal leaves the set admissible, then the original set is not minimal sufficient. Every element of a minimal sufficient constraint set is individually necessary.*

This provides a formal version of a principle familiar in lean management: a rule that can be removed without consequence should be removed. The formal result is conditional—it holds given an admissibility criterion and a cost function—and is related to, though not identical with, the simulation finding of Davis et al. (2009) that the structure–performance relationship is asymmetric, with greater penalty for too little structure than for too much.

3.4 Distributed capability and dependence

The capability of a team is represented by a matrix K where $K_{ij} \in [0, 1]$ indicates member i's proficiency at requirement j. The team collectively *covers* a set of requirements if, for each requirement, at least one member is proficient.

Proposition 7 (P7: Hall condition)**.** *The team can assign each requirement to a distinct proficient member if and only if Hall's marriage condition holds: for every subset S of requirements, $|\{i : \exists j \in S,\ k_{ij} > 0\}| \geq |S|$.*

Proposition 8 (K1: Diffusion monotonicity)**.** *Under a knowledge-sharing dynamic* $k_{ij}(t+1) = \min(1, k_{ij}(t) + \eta \cdot L_{ij}(t))$ *with* $\eta \geq 0$*,* $L \geq 0$*, and initial capabilities in* $[0, 1]$*, each component of* K *is monotonically non-decreasing, and collective coverage, once achieved, is preserved.*

Proposition 9 (K2: Redundancy)**.** *If every requirement remains covered after the removal of any single member—that is, each requirement has at least two independent qualified members—then removing any one member preserves collective coverage.*

Countermodel 2 (`concentrated_cover_breaks`)*.* Without redundancy, removing one member can destroy coverage.

Manager dependence. Manager dependence is modeled as a scalar $D(t)$ representing the fraction of decisions that require escalation. Under the dynamic $D(t+1) = (1-p)\,D(t)$ with $p > 0$, the parameter p is interpreted as reflecting the rate at which the team absorbs decision-making capacity. The exponential decay is an assumption, not a derived consequence.

Proposition 10 (D1–D5)**.** $D(t) = (1-p)^t D(0)$ *(closed form);* D *is strictly decreasing when* $D(0) > 0$ *and* $0 < p < 1$*;* $D(t) \to 0$ *as* $t \to \infty$*; for any* $\varepsilon > 0$ *there exists finite* T *with* $D(t) \leq \varepsilon$ *for all* $t \geq T$*; these results extend to time-varying* $p(t) \geq p_{\min} > 0$*.*

The connection between capability diffusion and manager dependence is formalized through an escalation model (the diffusion dynamics themselves are developed in §5). If problems have difficulties drawn from a distribution $\mathcal{F}$, and a member with knowledge level z can solve problems of difficulty at most z, then the escalation rate is $1 - \mathcal{F}(z)$. As knowledge sharing raises z across the team (K1), escalation decreases (ESC1). In Garicano's (2000) knowledge hierarchy, cheaper knowledge acquisition expands front-line knowledge and reduces escalation, whereas cheaper communication has the opposite effect. RCAS formalizes the knowledge-acquisition side of this comparative statics as a dynamic process, providing a formal mechanism for the shift from formal to real authority described by Aghion and Tirole (1997).

Leadership in this framework is not a property of a designated leader. It is a distributed function: $D(t) \to 0$ does not mean a leader has been removed; it means the leadership function has been absorbed by the team. This conceptual position is shared with complexity leadership theory (Uhl-Bien et al. 2007; Hazy 2006), though the formal model was constructed independently.

3.5 Rule updating

Rules evolve through an update dynamic. At each step, the team evaluates nearby rule sets and adopts the one with the highest evaluation, subject to a *stay condition*: the current rule set is always a candidate.

Proposition 11 (P8: Monotonicity)**.** *Under the stay condition, the evaluation of the rule set is monotonically non-decreasing over time.*

Proposition 12 (P9: Strict improvement from deletion)**.** *If the rule set contains an element whose removal would improve the evaluation, then the update produces a strict improvement and the rule set changes.*

Countermodel 3 (`stay_needed`). Without the stay condition, the update can decrease the evaluation. Any rule-update mechanism must include the option of keeping the current rules.

This model differs from evolutionary economics (Nelson and Winter 1982) in a specific way. In their framework, variation arises from deliberate but boundedly rational search, and market selection filters the results. In RCAS, the team itself performs both search and selection through the evaluation-and-stay dynamic. The stay condition ensures that the process does not regress.

3.6 Premise-based decisions

Proposition 13 (P10a)**.** *If whether each premise holds is decidable, then whether the decision is reopenable is also decidable.*

Proposition 14 (P10b)**.** *If the premise set is empty, the decision is never reopenable through the premise-triggered mechanism. The formal reopening mechanism requires recorded premises to operate on; other means of revision are not excluded.*

4 Performance as a Vector, Not a Scalar

4.1 Three dimensions per domain

Three levels of organizational performance should be distinguished.

Capability is what the team can do: the distribution of proficiency across members and domains (K, K^*), and the redundancy of that distribution (N). Capability is a stock variable that changes through learning, hiring, and knowledge sharing.

Output is what the team actually produces, given its capability and the engagement of its members. The realized frontier F and the throughput volume V are output-level quantities.

Outcome is the value that the organization's environment assigns to the output. Outcome depends on the reference point ρ—whether the right goals and criteria were set—and on external conditions that the team does not control.

RCAS formalizes the first two levels: the connection from capability to output. The third level lies outside the formal boundary. When this paper uses the word *performance*, it refers to the capability-to-output connection unless otherwise stated.

In a common governance arrangement where ρ is set externally, the team controls capability development and the quality of its output relative to ρ; it does not control the quality of ρ itself. The implications of this asymmetry are discussed in §9.

For each domain d, RCAS defines a *productive-performance profile* with four components.

The **potential capability ceiling** is $K_d^* = \max_i k_{id}$.

In general, a member's productive contribution is $q_{id} = g_d(k_{id}, e_{id})$, where g_d is monotonically increasing in both arguments and $g_d(k, 0) = 0$. The baseline model adopts the

multiplicative form $g_d(k, e) = k \cdot e$. The formal arguments requiring only monotonicity generalize to any g_d with these properties; those depending on the multiplicative structure are identified in the supplementary correspondence table.

The **realized frontier** is

$$F_d = \max_i g_d(k_{id},\, e_{id}),$$

the highest productive contribution any member achieves. Under the baseline form, $F_d = \max_i(k_{id} \cdot e_{id})$. Here $e_{id} \in [0, 1]$ is member i's engagement in domain d—the degree to which capability is actually exercised. (A motivation model connecting engagement to the self-determination theory continuum is developed in the supplementary material.)

The **throughput volume** is

$$V_d = \sum_i e_{id} \cdot \varphi_d(k_{id}),$$

where φ_d is monotonically increasing.

The **resilience count** is $N_d = |\{i : k_{id} \cdot e_{id} \geq \theta_d\}|$.

Proposition 15 (PF1: Frontier bounds)**.** *$k_h \cdot e_h \leq F_d \leq k_h$, where h is the most capable member.*

Proposition 16 (PF2: Full engagement)**.** *If $e_h = 1$, then $F_d = K_d^*$.*

Proposition 17 (PR1: Resilience equivalence)**.** *The team tolerates the loss of any m members while retaining at least one capable member if and only if $m < N_d$.*

4.2 From one expert to distributed expertise

Consider a domain with threshold $\theta = 0.5$, one expert (capability 1.0, engagement 1.0) and two novices (capability 0, engagement 1.0). The team state is ($K^* = 1,\ F = 1,\ V = 1,\ N = 1$). After one step of knowledge sharing with $\eta = 0.5$, each novice's capability rises to 0.5. Since $0.5 \cdot 1.0 = 0.5 \geq \theta$, both enter the capable set. The new state is ($K^* = 1,\ F = 1,\ V > 1,\ N = 3$). The capability ceiling is unchanged, but the team is strictly stronger: more output and more resilience.

4.3 Task-type correspondence

The decomposition has a natural correspondence with Steiner's (1972) task aggregation taxonomy.

Proposition 18 (ST1)**.** *Under pure knowledge sharing ($\eta \leq 1$), the disjunctive aggregate defined on raw capability (the ceiling K^*) is unchanged. The realized frontier F may rise even under pure sharing (see §7, `disjunctive_can_rise`).*

Proposition 19 (ST2)**.** *The additive aggregate (volume) is monotonically non-decreasing.*

Proposition 20 (ST3)**.** *The conjunctive aggregate (minimum capability) is monotonically non-decreasing.*

4.4 Team-level aggregation

Proposition 21 (PT1: Team monotonicity)**.** *With* $P_{\text{team}} = \sum_d w_d \cdot \Psi(F_d, V_d, N_d)$ *where* $w_d \geq 0$ *and* Ψ *is monotone increasing, if knowledge sharing operates in every domain and engagement is held fixed,* P_{team} *is monotonically non-decreasing.*

4.5 Why the same low performance can have different causes

The F-V-N decomposition provides a diagnostic structure. If F_d is low, the team may lack capability at the top—the response is acquisition. If F_d is high but V_d is low, capability is concentrated in one person—the response is diffusion. If K_d^* is high but F_d is low, capability exists but is not utilized—the gap is engagement or assignment. If N_d is low, the team is fragile. These diagnoses are mutually exclusive in mechanism and overlapping in symptom.

5 Knowledge Diffusion, Creation, and Operating Modes

The results in this section are stated for a single domain d; the domain subscript is suppressed.

5.1 Pure diffusion dynamics

Knowledge sharing is modeled as a convergence dynamic:

$$k_i(t+1) = k_i(t) + \eta(t,i) \cdot \big(k_h(t) - k_i(t)\big), \quad i \neq h,$$

where h is the member with the highest capability, and $0 \leq \eta \leq 1$.

Countermodel 4 (`rate_le_one_needed`). Relaxing the $[0,1]$ boundedness assumption to permit $\eta = 2$, capability oscillates: $0, 2, 0, 2, \ldots$ The bound separates diffusion from creation.

Under pure sharing ($\eta \leq 1$):

Proposition 22 (PS2: Ceiling invariance)**.** $K^* = \max_i k_i$ *is constant.*

Proposition 23 (PV1–PV2: Volume non-decrease)**.** V *is monotonically non-decreasing; it strictly increases when a non-top member has positive engagement and positive transfer rate.*

Proposition 24 (PR2: Resilience non-decrease)**.** N *is monotonically non-decreasing.*

5.2 Convergence

If $\eta(t,i) \geq \eta_{\min} > 0$ for all non-top members, then all capabilities converge to k_h (`tendsto_top`); volume converges to $(\sum_i e_i) \cdot \varphi(k_h)$ (PV3); every member with a positive engagement lower bound eventually meets the threshold (PR3).

5.3 Capability-ceiling expansion requires creation

Theorem 25 (PM1: Ceiling expansion)**.** *K^* increases from one period to the next if and only if some member's capability exceeds the previous ceiling. Under pure sharing, this cannot happen.*

Two main results follow:

Diffusion expands utilization: sharing raises V and N without raising K^*. The realized frontier F may also rise if engagement patterns shift.

Creation expands the ceiling: K^* rises only when someone creates knowledge beyond the current maximum.

The unsolvable-problem rate is also invariant under sharing. Since sharing does not change k_h, the fraction $1 - \mathcal{F}(k_h)$ is constant (`ceilingEscalation_const`).

5.4 Operating modes

Proposition 26 (OM1: Preservation)**.** *In a specialization mode without independent learning, volume and resilience are preserved.*

Proposition 27 (OM2: Monotonicity across modes)**.** *Alternating between diffusion and specialization preserves the monotonicity of V and N.*

Proposition 28 (OM3: Ceiling expansion through specialization)**.** *If independent learning raises a member's capability above the current ceiling, K^* increases. Within the present two-mode model, specialization is the mechanism by which new capability at the top is created.*

This connects to the trade-off identified by Dessein and Santos (2006): adaptation and specialization trade off against each other, and better communication complements adaptation. RCAS adds a temporal dimension: the team can sequence the phases rather than choosing one permanently.

5.5 Relation to prior work

Lazer and Friedman (2007) used agent-based simulation on NK fitness landscapes to show that efficient information sharing improves short-run performance but reduces long-run performance on complex problems. RCAS isolates one mechanism compatible with their result: diffusion expands utilization without expanding the frontier of reachable solutions. What RCAS adds is a theorem-level separation: the distinction between diffusion and creation is not a simulation artifact but a formal consequence of the model's structure.

6 Adaptive Self-Management

6.1 Strong self-managing state

A team is in a *structural self-managing state* when three conditions hold: (1) the team has genuine choice in every domain; (2) the team collectively covers all required capabilities;

(3) the team updates its own rules through the evaluation-and-stay dynamic. *Viable* self-management adds: (4) performance meets or exceeds a threshold in every domain. The central countermodel (§6.3) shows that structural self-management does not entail viable self-management.

The *strong* self-managing state adds two further conditions from a fast-slow dynamical model. In the fast dynamics, the team's operational state x evolves under a contraction mapping f_θ plus an external forcing term $u(t)$ representing managerial directives. In the slow dynamics, the operating parameters θ are updated by the team: $\theta(t+1) = \text{paramUpdate}(\theta(t), x(t), \omega(t))$.

The strong state requires (5) $u(t) = 0$ for all sufficiently large t; and (6) the team takes over parameter updating. The contraction mapping assumption restricts the model to teams whose operational dynamics are inherently stabilizing; chaotic or explosively unstable states would require a different framework. Self-management is autonomous adaptation, not stasis.

6.2 Fast-slow dynamics

Proposition 29 (F1: Forcing does not persist)**.** *If the fast dynamics is governed by a contraction mapping and external forcing ceases, the modeled state variable converges to the fixed point regardless of the forcing history. Within this model, past forcing leaves no permanent effect on the fast-state variable once the structural parameters are held fixed. (In real organizations, directives may also change capability, trust, or relational structure—effects modeled as changes to θ, not as persistence of forcing.)*

Proposition 30 (F2: Directives and constraints are distinct)**.** *In the specified contraction model, two trajectories with different constraint parameters converge to different limits, while two with the same constraints but different forcing histories converge to the same limit.*

6.3 The central countermodel

Countermodel 5 (`performance_not_derivable`). A model exists in which all three structural conditions hold—genuine choice, capability coverage, and rule self-updating—but performance is identically zero. Structural self-management does not entail adequate performance.

This countermodel is the formal expression of the paper's central epistemic claim: RCAS does not embed success into its definitions. Whether the structural conditions lead to good performance is an empirical question, not a formal consequence.

7 Machine-Checked Results, Countermodels, and Boundary of Inference

Claim	Status	Key declaration
Distinct methods can satisfy the same criterion	Theorem	P2
Every element of a minimal constraint set is necessary	Theorem (C3)	`not_minSufficient_of_erase`
Rule evaluation never decreases under stay	Theorem (P8)	`monotone_J_iterate`
Coverage preserved under diffusion	Theorem (K1)	`diffusion_cover_preserved`
Manager dependence $\to 0$	Theorem (D3)	`dependence_tendsto_zero`
Sharing cannot raise K^*	Theorem (PS2)	`ceiling_const`
Sharing increases V	Theorem (PV1)	`volume_monotone`
Ceiling expansion requires exceeding max	Theorem (PM1)	`ceiling_lt_iff`
Specialization can raise the ceiling	Theorem (OM3)	`ceiling_lt_iff_of_specializa`
Structural self-management $\Rightarrow$ performance	Not derivable	`performance_not_derivable`
Nested structure $\Rightarrow$ soundness	Not derivable	`sound_not_automatic`
Rule update without stay preserves evaluation	Not derivable	`stay_needed`
Pure sharing with $\eta > 1$ preserves monotonicity	Not derivable	`rate_le_one_needed`
RCAS practices increase η	Empirical hypothesis	—
Real organizations satisfy these assumptions	Empirical hypothesis	—

A fifth result, `disjunctive_can_rise`, is a construct-distinction witness: the realized frontier F can increase under knowledge sharing even though K^* does not. This confirms that the distinction between K^* and F is substantive.

The full Lean formalization comprises 840 declarations with zero uses of `sorry`. The Audit module machine-checks that every declaration depends only on the three standard axioms of Lean's logic (`propext`, `Classical.choice`, `Quot.sound`). The complete source code, correspondence table, deviation record, and empirical hypothesis list are available as supplementary material [repository link omitted for double-anonymous review].

8 Empirical Interface and Diagnostic Implications

8.1 Empirical phenomena compatible with RCAS mechanisms

RCAS mechanism	Prior finding	Relation	What remains untested
PS2: ceiling invariance	Lazer & Friedman (2007): efficient diffusion improves short-run but not long-run performance	Compatible: diffusion raises utilization without expanding the frontier	Whether K^*, V, η map to their search landscape
PF1–PF2: frontier depends on engagement	Ryu et al. (2022): autonomy affects performance indirectly via team functioning	Compatible: autonomy does not directly move the frontier when engagement mediates	Whether e_{id} captures the team-functioning construct
D1–D5, ESC: knowledge diffusion reduces dependence	Garicano (2000): front-line knowledge expansion reduces escalation	Dynamic version of comparative statics	Whether escalation rate is proportional to D
C1–C3: nonminimal constraints dominated	Davis et al. (2009): inverted-U; steeper decline below the optimum	Related: both penalize unnecessary structure	Whether the inverted-U maps to RCAS's cost function
ST1–ST3: task type determines which component moves	Steiner (1972): disjunctive, additive, conjunctive aggregation	Natural correspondence	Whether real tasks decompose cleanly

The word *compatible* is chosen deliberately. These prior findings were not designed to test RCAS, and they do not measure RCAS variables directly.

8.2 An illustrative conjecture: the shared-leadership complexity conflict

Wang et al. (2014) find stronger shared-leadership effects for more complex work. D'Innocenzo et al. (2016), using a different complexity coding and outcome definition, find weaker effects for more complex tasks ($\beta = -0.05$, $p < .05$). RCAS's F-V-N decomposition suggests one possible source: if complexity drives capability concentration in some samples and coordination cost in others, the opposing moderator directions would be expected. This conjecture is not formalized in the present model—it would require decomposing D into capability-dependence and coordination-dependence components—and is offered as a hypothesis for future work.

8.3 Diagnostic representation

The formal theory induces a diagnostic representation. Any organization O, regardless of its management methodology, can be described by a state vector:

$$\mathcal{D}(O) = (F_d,\ V_d,\ N_d,\ \eta_d,\ D,\ \kappa,\ \ldots)$$

The diagnostic question is not whether the organization practices RCAS, but which intermediate variables are binding. The development of this diagnostic framework is the subject of a companion paper.

9 Discussion and Conclusion

Three results stand out. First, performance is not a scalar. The decomposition into frontier (F), volume (V), and resilience (N) reveals that the same symptom—poor performance—can arise from capability gaps, capability concentration, engagement failures, or single points of failure. Second, knowledge diffusion and knowledge creation are formally distinct. Diffusion raises throughput and resilience but cannot raise K^*. The realized frontier F may rise under diffusion if engagement shifts, but this reflects better utilization, not new capability. Third, the theory does not embed success: the countermodel `performance_not_derivable` shows that all structural conditions of self-management can hold while performance fails.

The capability–output–outcome distinction (§4) has a governance implication. In arrangements where ρ is set externally, the team controls capability development and output quality relative to ρ, but not the quality of ρ itself. Output that satisfies a poorly chosen ρ does what was asked; the shortfall is a property of the reference point. This asymmetry suggests that evaluating teams solely on outcome conflates what teams control with what they do not.

The formal model defines the structural principles but does not prescribe how to implement them in daily practice. One concrete instantiation is described in a companion paper, which maps daily practices—including reference-point specification, pair review, and premise-based decision records—onto the framework of complexity leadership theory (Author, submitted). The formal model does not depend on that particular instantiation, but practitioners can consult it as a worked example.

Limitations. The model formalizes the connection from capability to output but not from output to outcome. It analyzes a high-observability regime; low-information-sharing organizations would require different assumptions. The escalation model's prediction that the unsolvable-problem rate is invariant under sharing (`ceilingEscalation_const`) highlights a boundary: sharing makes the team more efficient within its range but does not extend it. The convergence results assume a positive lower bound on η. The formal extensions covering network propagation, threshold dynamics, multi-level attention alignment, and Stuart–Landau amplitude equations are reported in supplementary material and will be developed in companion papers.

Future work. (1) Empirical validation: GitHub truck-factor datasets provide a candidate setting; a separate study is being designed to test whether the F-V-N decomposition predicts post-departure performance better than aggregate measures. (2) A companion paper develops $\mathcal{D}(O)$ into a practical organizational assessment framework. (3) The Stuart–Landau dynamical component is being developed as an independent contribution.

We close with the statement that motivated this work. We do not claim that RCAS improves organizational performance. We identify the intermediate mechanisms through which such improvement could occur, derive the consequences that follow from explicit assumptions, prove where those consequences fail without them, and leave the remaining links to empirical test.

References

Aghion P, Tirole J (1997) Formal and real authority in organizations. J Polit Econ 105:1–29. https://doi.org/10.1086/262063

Author (submitted) Connecting practice to complexity leadership theory: How local rules are set and changed in self-managing teams. Manuscript submitted for publication [details omitted for double-anonymous review]

Davis JP, Eisenhardt KM, Bingham CB (2009) Optimal structure, market dynamism, and the strategy of simple rules. Adm Sci Q 54:413–452. https://doi.org/10.2189/asqu.2009.54.3.413

Dessein W, Santos T (2006) Adaptive organizations. J Polit Econ 114:956–995. https://doi.org/10.1086/508031

D'Innocenzo L, Mathieu JE, Kukenberger MR (2016) A meta-analysis of different forms of shared leadership–team performance relations. J Manage 42:1964–1991. https://doi.org/10.1177/0149206314525205

Garicano L (2000) Hierarchies and the organization of knowledge in production. J Polit Econ 108:874–904. https://doi.org/10.1086/317671

Hazy JK (2006) Measuring leadership effectiveness in complex socio-technical systems. Emergence: Complexity & Organization 8:58–77

Holliday WH, Norman C, Pacuit E (2021) Voting theory in the Lean theorem prover. In: Ghosh S, Icard T (eds) Logic, Rationality, and Interaction (LORI 2021), LNCS, vol 13039. Springer, pp 111–127. https://doi.org/10.1007/978-3-030-88708-7_9

Holmström B (1982) Moral hazard in teams. Bell J Econ 13:324–340. https://doi.org/10.2307/3003457

Jordan WC, Graves SC (1995) Principles on the benefits of manufacturing process flexibility. Manage Sci 41:577–594. https://doi.org/10.1287/mnsc.41.4.577

Lazer D, Friedman A (2007) The network structure of exploration and exploitation. Adm Sci Q 52:667–694. https://doi.org/10.2189/asqu.52.4.667

Li J (2025) From axioms to algorithms: Mechanized proofs of the vNM utility theorem. Preprint, arXiv:2506.07066

Li J, Wang J, Zhou L (2025) Quantifying bounded rationality: Formal verification of Simon's satisficing through flexible stochastic dominance. Preprint, arXiv:2507.07052

Marschak J, Radner R (1972) Economic theory of teams (Cowles Foundation Monograph No. 22). Yale University Press, New Haven

Nelson RR, Winter SG (1982) An evolutionary theory of economic change. Belknap Press of Harvard University Press, Cambridge

Nipkow T (2009) Social choice theory in HOL: Arrow and Gibbard–Satterthwaite. J Autom Reasoning 43:289–304. https://doi.org/10.1007/s10817-009-9147-4

Ryu JW, Neubert EM, Gonzalez-Mulé E (2022) Putting the team in the driver's seat: A meta-analysis on the what, why, and when of team autonomy's impact on team effectiveness. Pers Psychol 75:411–439. https://doi.org/10.1111/peps.12468

Simon HA (1951) A formal theory of the employment relationship. Econometrica 19:293–305. https://doi.org/10.2307/1906815

Steiner ID (1972) Group process and productivity. Academic Press, New York

Uhl-Bien M, Marion R, McKelvey B (2007) Complexity leadership theory: Shifting leadership from the industrial age to the knowledge era. Leadersh Q 18:298–318. https://doi.org/10.1016/j.leaqua.2007.04.002

Wang D, Waldman DA, Zhang Z (2014) A meta-analysis of shared leadership and team effectiveness. J Appl Psychol 99:181–198. https://doi.org/10.1037/a0034531

Wood RE (1986) Task complexity: Definition of the construct. Organ Behav Hum Decis Process 37:60–82. https://doi.org/10.1016/0749-5978(86)90044-0

Declarations

Funding: No funding was received for this research.

Competing interests: The author has no relevant financial or non-financial interests to disclose.

Data availability: The complete Lean 4 source code, including all definitions, theorems, countermodels, and the axiom audit module, is publicly available at a permanent repository [link omitted for double-anonymous review]. No other data were generated or analysed during this study.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1leansource.zip