

On the Returns to Scale of Organisational Coordination

A Mathematical Proof for Hub-Mediated Intelligence Architectures

As implemented in the Altimist Intelligent Operating System™

Arthur L. T. Davis

Altimist® Ltd | 4 July 2026 | Version 1.1 (Public Edition)

Working Paper — Public Edition

Keywords: coordination economics, organisational theory, sovereign AI, hub-mediated architecture, returns to scale, emotional intelligence, Coase, Brooks, Dunbar

Abstract. We prove that bilateral coordination architectures — the organisational model used by virtually every firm in history — exhibit strictly negative returns to scale on coordination ($d^2P_B/dn^2 < 0$), with marginal coordination cost growing linearly in team size while the coordination burden grows quadratically. We then prove that hub-mediated intelligence architectures, in which a shared knowledge layer absorbs coordination as a side effect of productive work, exhibit strictly positive returns to scale ($d^2P_H/dn^2 > 0$) under defined conditions. We identify a fourth necessary condition for the proof to hold in practice: emotionally intelligent knowledge governance — the system must preserve the psychological safety required for participants to share genuine knowledge. Without this condition, the knowledge utility function degrades and the positive-returns property is lost. We provide three theorems with formal proofs, a corollary, and specific points of contradiction with Coase (1937), Brooks (1975), Dunbar (1992), Gulick and Urwick (1937), Chandler (1962), Mintzberg (1979), and Porter (1985).

1. Introduction

The question of how organisations coordinate is among the oldest in economics. Coase (1937) established that firms exist because internal coordination is cheaper than market transaction costs, but that coordination costs grow with firm size.¹ Brooks (1975) formalised the constraint: as teams grow, coordination cost grows quadratically.² The number of bilateral communication channels in a team of n people is $n(n-1)/2$. Dunbar (1992) identified a cognitive ceiling of approximately 150 stable relationships, constrained by neocortex size.³ Empirical evidence consistently supports these constraints: the Standish Group found project success rates decline sharply with team size,⁴ and ISBSG analyses identify 3–5 members as optimal.⁵

These constraints apply universally but are particularly visible in knowledge-intensive organisations where work is interdisciplinary and the risk of redundant effort is high. López de Prado (2018) identifies this as the ‘Sisyphus paradigm’ — individual researchers working in isolation repeat known mistakes because institutional knowledge fails to accumulate.⁶ This paper asks whether a new class of technology — sovereign AI agent infrastructure — can structurally change this coordination economics.

Recent scholarship has observed that AI is reshaping organisational structure. Brynjolfsson and McAfee (2014) identified the potential for machine intelligence to restructure work; Iansiti and Lakhani (2020) argued that AI-centric

¹Coase, R.H. (1937). ‘The Nature of the Firm’. *Economica*, New Series, Vol. 4, No. 16, pp. 386–405. Nobel Memorial Prize, 1991.

²Brooks, F.P. Jr. (1975). *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley. The $n(n-1)/2$ formula for bilateral communication channels appears in Chapter 2.

³Dunbar, R.I.M. (1992). ‘Neocortex size as a constraint on group size in primates’. *Journal of Human Evolution*, 22(6), 469–493.

⁴Standish Group (2015). *CHAOS Report*. Small projects (≤ 5 people): 61% success. Large: 11%. Failure attributed to coordination overhead.

⁵ISBSG Repository Analysis (2012). Study of 1,000+ projects: 3–5 members optimal; teams of 9+ exhibited significantly lower productivity.

⁶López de Prado, M. (2018). *Advances in Financial Machine Learning*. John Wiley & Sons. Chapter 1 identifies the ‘Sisyphus paradigm’ — individual researchers working in isolation repeat known mistakes because institutional knowledge does not accumulate.

operating models are replacing traditional firms; Agrawal, Gans and Goldfarb (2018) formalised AI as a prediction technology that lowers the cost of decision-making.⁷ These contributions describe the phenomenon qualitatively. This paper provides a formal mathematical proof.

We present evidence that it can, and that the change is not incremental but structural: the sign of the second derivative of net productive output with respect to team size flips from negative to positive. This is not a refinement of existing theory. It is an inversion.

2. Definitions

Definition 2.1 (Bilateral Coordination Architecture).

An organisation of n human participants in which each participant must maintain a direct communication channel with every other participant. The number of channels is:

$$C_B(n) = n(n - 1) / 2$$

Definition 2.2 (Hub-Mediated Intelligence Architecture).

An organisation of n human participants in which each participant communicates exclusively through a shared intelligence hub H . The hub maintains a persistent knowledge layer K that accumulates as a byproduct of productive interaction.⁸ The number of channels is:

$$C_H(n) = n$$

Definition 2.3 (Coordination Cost per Capita).

Let each bilateral channel require $c > 0$ units of effort per unit time. The coordination cost per capita is total coordination effort divided by team size.

Definition 2.4 (Knowledge Utility Function).

In a hub architecture, define $U(n)$ as the utility each participant derives from accumulated knowledge. We assume $U(n) = \delta \cdot \ln(1 + n)$ for some $\delta > 0$, reflecting diminishing marginal returns to additional contributors but strictly increasing total utility.

The choice of logarithmic form requires justification, as it is the load-bearing assumption of the proof. Each additional contributor adds knowledge but with diminishing marginal novelty (hence U is concave in n). However, the total knowledge utility across all participants — $n \cdot U(n) = \delta n \cdot \ln(1+n)$ — must be convex for the positive-returns result to hold. The logarithmic function satisfies both properties. Critically, the result generalises: any utility function where $n \cdot U(n)$ is convex will produce a positive second derivative. The logarithmic form is sufficient but not uniquely required.⁹

Definition 2.5 (Net Productive Output).

$$\text{Bilateral: } P_B(n) = \alpha n - c \cdot n(n - 1) / 2$$

$$\text{Hub: } P_H(n) = \alpha n - hn + \delta n \cdot \ln(1 + n)$$

where $h > 0$ is the per-capita hub connection cost.

Remark 2.5.1. In both models, n counts human participants only. In a hub-mediated architecture with AI agents (such as the AIOS), agents operate within the hub, not alongside it. They are part of the hub's capacity — expanding its knowledge, speed, and domain coverage — not additional participants requiring coordination.

⁷Brynjolfsson, E. & McAfee, A. (2014). *The Second Machine Age*. W.W. Norton. Iansiti, M. & Lakhani, K.R. (2020). *Competing in the Age of AI*. Harvard Business Review Press. Agrawal, A., Gans, J. & Goldfarb, A. (2018). *Prediction Machines*. Harvard Business Review Press.

⁸For the reference implementation of a persistent, shared knowledge layer of this kind, see Altimist Ltd (2026), *Altimist Intelligent Operating System Architecture Brief*, available from the author on request. The proof itself does not depend on any particular implementation.

⁹The logarithmic form follows the standard network economics assumption of diminishing marginal novelty per contributor. Metcalfe's Law (network value $\propto n^2$) and Reed's Law (network value $\propto 2^n$) describe communication and group-forming networks respectively. For knowledge networks, the appropriate form is sublinear in the marginal contribution but superlinear in total utility: $U(n)$ concave, but $n \cdot U(n)$ convex. The logarithmic function satisfies both properties. The result is robust to any utility function where $n \cdot U(n)$ is convex — log is sufficient but not uniquely required.

Deploying a new agent increases the hub's capability without increasing n or adding any coordination cost. This is the mechanism by which capability decouples from headcount.

Definition 2.6 (Returns to Scale on Coordination).


An architecture exhibits positive returns to scale if $d^2P/dn^2 > 0$. It exhibits negative returns to scale if $d^2P/dn^2 < 0$.

Remark 2.6.1. Throughout this paper, n is treated as a continuous variable for analytical convenience, consistent with standard practice in production economics and organisational theory. The results hold under finite differences for discrete n : the second-order finite difference of a strictly convex (or strictly concave) function preserves the sign of the continuous second derivative.

Figure 1: Bilateral vs Hub-Mediated Coordination Architecture

Bilateral: 28 channels

$$n(n-1)/2 = 8(7)/2 = 28$$



VS

Hub-Mediated: 8 channels

$$n = 8 \text{ (+ knowledge layer)}$$




Figure 1. Left: bilateral coordination architecture with 8 participants requiring 28 communication channels. Right: hub-mediated architecture with the same 8 participants requiring only 8 connections to the central hub (H). The gold ring represents the shared knowledge layer that accumulates as a side effect of productive interaction.




Figure 2. Illustrative net productive output for bilateral (red, concave) and hub-mediated (green, convex) architectures. The bilateral curve peaks at n^* then declines. The hub curve is strictly convex, accelerating as the knowledge layer grows. The crossover n_0 marks hub dominance. Curves drawn using the generalised convex utility family (Definition 2.4). Exact shape depends on α , c , h , δ .

3. Theorems and Proofs

3.1 Negative Returns in Bilateral Architectures

Theorem 1 (Negative Returns to Scale — Bilateral Architecture). In a bilateral coordination architecture with $n \geq 2$ participants, coordination cost per capita grows linearly in n , and net productive output exhibits strictly negative returns to scale.

Proof. Total coordination cost: $C_B(n) = c \cdot n(n-1) / 2$. Per capita: $c(n-1) / 2$ — linear in n , strictly increasing.

$$P_B(n) = \alpha n - cn(n-1) / 2$$

$$dP_B/dn = \alpha - cn + c/2$$

$$d^2P_B/dn^2 = -c$$

Since $c > 0$, we have $d^2P_B/dn^2 = -c < 0$ for all n .¹⁰ The function is strictly concave. Maximum at $n^* = \alpha/c + 1/2$.¹¹ ■

3.2 Positive Returns in Hub-Mediated Architectures

Theorem 2 (Positive Returns to Scale — Hub Architecture). In a hub-mediated intelligence architecture with knowledge utility $U(n) = \delta \cdot \ln(1+n)$, the net productive output exhibits strictly positive returns to scale for all $n \geq 0$.

Proof.

$$P_H(n) = \alpha n - hn + \delta n \cdot \ln(1+n)$$

The third term $\delta n \cdot \ln(1+n)$ is the knowledge network effect.¹²

$$dP_H/dn = (\alpha - h) + \delta \cdot \ln(1+n) + \delta n/(1+n)$$

$$d^2P_H/dn^2 = \delta/(1+n) + \delta/(1+n)^2 = \delta(2+n)/(1+n)^2$$

Since $\delta > 0$ and $n \geq 0$: $d^2P_H/dn^2 > 0$ for all n .¹³ The function is strictly convex and admits no interior maximum. ■

Remark 3.2.1. Convexity requires only $\delta > 0$; the condition $\delta > h$ is not needed for Theorem 2. Its force lies in Theorem 3; if $\delta > h$, then $n_0 = e^{h/\delta} - 1 < e - 1 < 2$, so the hub architecture dominates the bilateral architecture for all $n \geq 2$. The condition $\delta > h$ is therefore best read as a dominance-from-the-outset condition, not a convexity condition.

3.3 Crossover Dominance

Theorem 3 (Crossover). There exists a finite crossover point $n_0 \leq e^{h/\delta} - 1$, with equality in the limiting case $c \rightarrow 0$, beyond which the hub architecture strictly dominates the bilateral architecture.

Proof. $\Delta(n) = P_H - P_B = -hn + \delta n \cdot \ln(1+n) + cn(n-1)/2$. As $n \rightarrow \infty$, $\Delta(n) \rightarrow \infty$. Best-case bilateral ($c=0$): crossover at $n_0 = e^{h/\delta} - 1$, finite for all $h, \delta > 0$. Since $\Delta(n)/n = -h + \delta \cdot \ln(1+n) + c(n-1)/2$ is strictly increasing in n , the crossing is unique; for $c > 0$ it occurs strictly earlier. ■

3.4 Zero Marginal Synchronisation Cost

Corollary 3.4.1. In the hub architecture, the marginal cost of synchronisation for participant $n+1$ is bounded by the constant h , not growing with n .¹⁴ In the bilateral architecture, marginal synchronisation cost is cn (linear growth). ■

¹⁰Brooks (1975), op. cit. The quadratic growth of communication overhead is the central mechanism of Brooks's Law: 'Adding manpower to a late software project makes it later.'

¹¹Schweik, C.M. & English, R. (2007). 'Brooks' Versus Linus' Law: An Empirical Test of Open Source Projects'. NCDG Working Paper No. 07-009.

¹²The assumption that knowledge accumulates as a byproduct of productive interaction — rather than through dedicated synchronisation work — is the defining property of Definition 2.2 and is grounded in the network-economics literature cited at footnote 9. An implementation reference is Altimist Ltd (2026), op. cit.

¹³The logarithmic knowledge utility $U(n) = \delta \cdot \ln(1+n)$ is standard in network economics. The key property: $n \cdot U(n)$ is convex for all $\delta > 0$, producing the positive second derivative.

¹⁴The bound follows directly from Definition 2.5: the only per-participant cost in the hub model is the connection cost h . Synchronisation in such an architecture is emergent — a property of the shared layer — rather than performed by participants. An implementation reference is Altimist Ltd (2026), op. cit.

4. Comparative Summary

Property	Bilateral Architecture	Hub-Mediated Architecture
Channels	$n(n-1)/2$ (quadratic)	n (linear)
Cost per capita	$c(n-1)/2$ (linear in n)	h (constant)
Knowledge utility	0 (no shared layer)	$\delta \cdot \ln(1+n)$ per person
d^2P/dn^2	$-c$ (negative)	$\delta(2+n)/(1+n)^2$ (positive)
Returns to scale	Strictly negative	Strictly positive
Sync cost/new member	cn (linear)	$\leq h$ (constant; Corollary 3.4.1)
Optimal team size	Finite: $n^* = \alpha/c + 1/2$	Unbounded (convex)

5. Discussion

The proof rests on two structural properties. First, the replacement of bilateral channels (quadratic) with hub connections (linear) eliminates the combinatorial explosion that drives Brooks’s Law.¹⁵ Second, the knowledge accumulation property introduces a positive externality that grows with participation.

This has a direct practical consequence identified by López de Prado as the ‘Sisyphus paradigm’: in organisations without institutional knowledge accumulation, individual researchers working in isolation repeat known mistakes.¹⁶ The hub-mediated architecture eliminates the Sisyphus paradigm, because every decision, every rejected hypothesis, and every lesson learned enters the permanent knowledge layer. No participant can unknowingly repeat a failed approach — the system remembers what was tried and why it failed.

An objection should be addressed. The bilateral model (Definition 2.1) assumes pure all-to-all coordination, which no organisation above approximately ten people actually uses. Real organisations deploy hierarchies, reducing the effective channel count from $n(n-1)/2$ to approximately $k \cdot n$, where k is the average number of direct relationships per person. This does not invalidate the proof for three reasons.¹⁷ First, hierarchies are themselves a coordination cost — the management layers required to maintain the hierarchy represent overhead that the hub eliminates. Second, hierarchical relay introduces information degradation at each level; the hub transmits knowledge with materially lower degradation than multi-level relay. Third, even hierarchical organisations exhibit super-linear coordination costs because the management layers must themselves coordinate with each other. The bilateral model is therefore not a straw man; it is the limiting case that all traditional architectures approach as complexity increases.

A conforming implementation must satisfy all four conditions (enumerated formally in Section 8); the AIOS serves as the reference implementation. A shared knowledge layer accessible to all participants eliminates the knowledge silos that would suppress the δ parameter. Event-driven propagation for high-priority decisions ensures that knowledge transfer approaches the real-time assumption of the proof. Federated semantic search and adaptive communication reduce the per-capita hub connection cost h . And emotionally intelligent knowledge governance — knowledge classification, attribution sensitivity, propagation timing, and cognitive privacy — maintains the knowledge quality parameter q close to 1 (q is defined formally in Section 6). The net effect on the crossover point is significant: since $n_0 = e^{h/\delta} - 1$, increasing δ and decreasing h simultaneously reduces n_0 , meaning the hub architecture dominates the bilateral architecture with fewer participants than a naïve implementation would require.

¹⁵Brooks, F.P. Jr. (1995). *The Mythical Man-Month, Anniversary Edition*. Addison-Wesley. Brooks notes modular architectures partially mitigate communication overhead but does not identify a mechanism that eliminates it entirely.

¹⁶López de Prado (2018), op. cit. The ‘Sisyphus paradigm’: researchers in isolation unknowingly repeat failed approaches because the organisation has no mechanism for institutional learning. The hub-mediated architecture solves this by making every failed experiment part of the permanent knowledge layer.

¹⁷Williamson, O.E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press. Williamson demonstrated that hierarchies are themselves a governance mechanism with their own transaction costs, including monitoring, information loss at each relay level, and bureaucratic rigidity. The hub-mediated architecture eliminates these hierarchy-specific costs.

A note on the condition $\delta > h$. Early dominance (Theorem 3, via Remark 3.2.1) requires that the knowledge utility parameter exceeds the per-capita hub connection cost. This places a minimum requirement on the quality and capability of the hub: a poorly implemented hub — one with high infrastructure overhead (large h) and weak knowledge accumulation (small δ) — would push the crossover point beyond any realistic team size. The condition is therefore not merely mathematical; it is an engineering constraint that distinguishes a genuine intelligence hub from a simple message-routing system.

6. Emotional Intelligence as a Necessary Condition

The proof assumes that the knowledge entering the hub is genuine. This assumption is non-trivial. If participants self-censor — withholding uncertainties, half-formed ideas, or mistakes — the knowledge layer fills with sanitised, performative content. The utility function $U(n)$ degrades, and the positive-returns property is lost.

Participants will only share genuine knowledge if they feel psychologically safe to do so.¹⁸ This requires the hub to exhibit emotional intelligence (EI)¹⁹ — specifically:

Knowledge governance. The system must distinguish between conclusions (propagate to the knowledge layer) and process (the working-through-it, the wrong turns, the uncertainty). Conclusions are institutional knowledge. Process is cognitive privacy.

Attribution sensitivity. Some knowledge should be attributed ('Person A concluded X'). Other knowledge should become ambient ('The research team identified X'). The system must know when attribution creates accountability and when it creates blame.

Propagation timing. Not all knowledge should propagate immediately. A half-formed concern propagated prematurely can cause organisational panic. The system must distinguish between decisions (propagate now), concerns (hold until confirmed), and reflections (never propagate).

Adaptive communication. Different participants require different register, depth, and framing. The same information delivered in the wrong emotional register fails to land.

We therefore identify a fourth necessary condition for the proof:

Condition (iv): Emotionally Intelligent Knowledge Governance. The hub must preserve the psychological safety necessary for participants to share genuine knowledge. Without this condition, the knowledge utility function $U(n)$ degrades and the positive-returns-to-scale property is lost.

This condition can be given a quantitative anchor. Let $q \in [0,1]$ represent the quality of knowledge entering the hub, where $q = 1$ is fully genuine and $q = 0$ is fully performative. The effective utility function becomes $U(n,q) = q \cdot \delta \cdot \ln(1+n)$. Convexity (Theorem 2) is preserved for any $q \cdot \delta > 0$; the practical force of the result, however, depends on q : the crossover point becomes $n_0 = e^{h/(q \cdot \delta)} - 1$, which grows without bound as $q \rightarrow 0$. High knowledge quality is what keeps hub dominance attainable at realistic team sizes. Condition (iv) is the architectural requirement that maintains $q \approx 1$ in practice. It should be noted that q is not solely a function of the hub's emotional intelligence. Pre-existing organisational culture — a history of blame, mistrust, or performative communication — will suppress q regardless of the hub's design. The hub's EI is necessary for high q but may not be sufficient in hostile cultural environments. Cultural transformation and architectural design must work in parallel.

The AIOS implements all four components architecturally: knowledge classification in the distillation pipeline (tagging outputs as CONCLUSION, PROCESS, or REFLECTION), an attribution policy in the hub's routing logic, three-tier propagation timing (DECISION, CONCERN, REFLECTION), and a cognitive privacy mode allowing participants to flag interactions as exploratory. Engagement metrics provide an early warning system for declining q .

¹⁸Edmondson, A. (2019). *The Fearless Organization: Creating Psychological Safety in the Workplace for Learning, Innovation, and Growth*. Wiley. Established that psychological safety is the critical predictor of team learning and performance.

¹⁹Salovey, P. & Mayer, J.D. (1990). 'Emotional Intelligence'. *Imagination, Cognition and Personality*, 9(3), 185–211. Defined emotional intelligence as the ability to perceive, understand, manage, and use emotions. Popularised by Goleman, D. (1995). *Emotional Intelligence*. Bantam Books.

Emotional intelligence is not a feature. It is a load-bearing component of the architecture.

7. Implications for Organisational Theory

7.1 Coase (1937): The Theory of the Firm

Coase's optimal firm boundary depends on internal coordination costs rising with size.²⁰ Theorem 2 demonstrates that when coordination changes from bilateral to hub-mediated, this assumption no longer holds for the coordination component of Coase's diminishing returns. Other sources of diminishing returns to management — bounded rationality, agency costs, monitoring complexity — may persist. But the primary mechanism that Coase identified as limiting firm size is weakened substantially, and the optimal firm boundary shifts outward.

7.2 Dunbar (1992): The Cognitive Limit

Dunbar's ~150 ceiling is a constraint on the neocortex, not on the hub. Each participant maintains one relationship (with the hub). The cognitive bottleneck moves from brain to infrastructure.²²

7.3 Gulick, Urwick & Graicunas (1937): Span of Control

The 5–9 direct report limit is cognitive.²³ The hub's span is limited by compute and context, not cognition. It does not degrade with more participants.

7.4 Chandler (1962), Mintzberg (1979), Porter (1985)

All three treat coordination overhead as the structural constraint.²⁴ If coordination cost per capita is constant or decreasing, the diseconomy of scale disappears. The long-run average cost curve does not turn upward.

7.5 The Nature of the Contradiction

The proof uses the same $n(n-1)/2$ formula that Brooks used in 1975 and standard calculus. The insight is not mathematical; it is architectural. Previous theory assumed bilateral coordination was the only structure possible, because until now it was. What changed is that for the first time, the hub can understand content, accumulate knowledge, and redistribute it contextually — without any participant performing synchronisation work.

7.6 Sovereign AI and the Firm

Sovereign AI — defined as a nation's capability to produce intelligence using its own infrastructure, data, workforce and business networks²⁵ — has become the dominant policy framework for national AI strategy. The argument is geopolitical: nations that depend on foreign infrastructure for their intelligence production surrender strategic autonomy.

This paper demonstrates that the same structural property holds at the level of the firm. An organisation that operates a hub-mediated intelligence architecture on its own infrastructure achieves sovereignty over its coordination intelligence — with mathematically provable positive returns to scale (Theorem 2). The knowledge layer accumulates domestically, the coordination economics improve with participation, and no external dependency is introduced.

²⁰Coase (1937), p. 395: 'A firm will tend to expand until the costs of organizing an extra transaction within the firm become equal to the costs of carrying out the same transaction by means of an exchange on the open market.'

²¹Coase (1937), p. 394: 'First, as a firm gets larger, there may be decreasing returns to the entrepreneur function.'

²²Dunbar (1992), op. cit. Lindenfors, P., Wartel, A., & Lind, J. (2021). 'Dunbar's number deconstructed'. *Biology Letters*, 17(3), 20200748. 95% confidence intervals of 4–520.

²³Gulick, L. & Urwick, L. (1937). *Papers on the Science of Administration*. Columbia University. Graicunas, V.A. (1933). Supervisory formula $R = n(2^{n/2} + n - 1)$.

²⁴Chandler, A.D. Jr. (1962). *Strategy and Structure*. MIT Press. Mintzberg, H. (1979). *The Structuring of Organizations*. Prentice-Hall. Porter, M.E. (1985). *Competitive Advantage*. Free Press.

²⁵Huang, J. (2024). Keynote address, World Governments Summit, Dubai, February 2024. 'Every country needs to own the production of their own intelligence.' See also NVIDIA Corporation, 'What Is Sovereign AI?', NVIDIA Blog, 2024.

The implication extends the sovereign AI thesis from nations to enterprises. Where Huang (2024) argues that every country should own the production of its own intelligence, the proof presented here establishes that every organisation of sufficient complexity should do the same — not for geopolitical reasons, but because the coordination economics are structurally superior to any alternative architecture. Sovereign AI is not only a national strategy. It is an organisational one.²⁶

It is notable that NVIDIA itself operates a radically flattened organisational structure in which the CEO serves as a de facto human hub with approximately 60 direct reports and no bilateral meetings.²⁷ Information is broadcast to all participants simultaneously; individual one-to-one meetings are eliminated; strategic direction is shared with the entire leadership team at once. This architecture mirrors the hub-mediated model described in this paper. However, as a human hub, the CEO remains bounded by the cognitive constraints identified by Dunbar (1992) and Graicunas (1933) — a limitation that organisational behaviour scholars have noted publicly.²⁸ The hub-mediated intelligence architecture removes this constraint by replacing the human hub with an intelligence layer that does not degrade with participation. The organisational intuition is correct. The proof explains why the AI implementation of that intuition scales without limit.

8. Conclusion

This paper proves that bilateral coordination architectures — the organisational model used by virtually every firm in history — exhibit strictly negative returns to scale on coordination (Theorem 1: $d^2P_B/dn^2 = -c < 0$). It then proves that hub-mediated intelligence architectures exhibit strictly positive returns (Theorem 2: $d^2P_H/dn^2 = \delta(2+n)/(1+n)^2 > 0$). Theorem 3 establishes that the hub architecture dominates the bilateral architecture beyond a finite crossover point.

Four conditions are necessary for the positive-returns property to hold in practice: (i) hub-mediated communication replacing bilateral channels; (ii) knowledge accumulation as a side effect of productive work; (iii) universal access to the knowledge layer without active synchronisation; and (iv) emotionally intelligent knowledge governance preserving the psychological safety required for genuine knowledge to enter the system. The fourth condition is not optional. Without it, the knowledge utility function degrades and the proof's assumptions fail.

The proof contradicts foundational assumptions in organisational theory (Coase, 1937), software engineering (Brooks, 1975), cognitive anthropology (Dunbar, 1992), management science (Gulick & Urwick, 1937; Graicunas, 1933), and strategic management (Chandler, 1962; Mintzberg, 1979; Porter, 1985). It does so using standard calculus and the same $n(n-1)/2$ formula that Brooks used half a century ago. The insight is not mathematical. It is architectural.

Limitations should be noted. The proof holds under specific assumptions — particularly that the knowledge utility function grows at least logarithmically and that knowledge quality q remains close to 1. The architecture also introduces a single point of failure at the hub; resilience, redundancy, and graceful degradation of the hub layer are engineering requirements that the mathematical framework does not address. The proof demonstrates the possibility of positive returns to scale on coordination, not the guarantee.

What is perhaps most significant is that the proof spans two fields that are traditionally treated as distinct. Theorems 1–3 and Section 7 are contributions to organisational theory — the macro-level study of firm boundaries, coordination structures, and why organisations take the shape they do. Condition (iv) crosses into organisational behaviour — the micro-level study of how individuals act within structures, including psychological safety and knowledge quality. The bridge between these fields is where the real contribution sits: a proof about organisational structure that cannot hold without a claim about human behaviour. The engineering must deliver what the mathematics promises — but so must the individual participants of the organisation. The proof holds only when both the infrastructure and the people using it sustain the conditions identified above.

²⁶The commercial implementation described in this paper — the Altimest Intelligent Operating System™ — operates on NVIDIA supercomputing infrastructure with sovereign compute, sovereign knowledge, and zero external AI dependencies. It is, to the author's knowledge, the first enterprise-scale implementation of sovereign AI at the organisational level.

²⁷Huang, J. (2024). Stanford Graduate School of Business, 2024. Huang stated that CEOs should have the largest number of direct reports. See Fortune, '60 direct reports, but no 1-on-1 meetings', November 2024. Confirmed in Lex Fridman Podcast, Episode published March 2026.

²⁸Chaturvedi, S. (2024). Quoted in CNBC, May 2024. Professor of Organisational Behaviour, Imperial College Business School: 'I can tell you for sure that he must be struggling to manage that many direct reports.' The observation illustrates the cognitive constraint that the hub-mediated intelligence architecture is designed to eliminate.

Traditional organisations have negative returns to scale on coordination. Hub-mediated intelligence architectures have positive returns. The sign of the second derivative flips from negative to positive. That is not a refinement. It is an inversion.

Acknowledgements

This paper was prepared with the assistance of Claude (Opus 4.6, Opus 4.8, and Fable 5), a family of large language models developed by Anthropic. Claude was used as a research and editorial collaborator throughout the development of this work, including literature review, mathematical verification, peer review, document drafting, and iterative refinement. All intellectual conclusions, architectural decisions, and claims are the author's own. The author reviewed, directed, and takes full responsibility for the accuracy and integrity of the final work.

References

- [1] Agrawal, A., Gans, J. & Goldfarb, A. (2018). *Prediction Machines*. Harvard Business Review Press.
- [2] Altimist Ltd (2026). *Altimist Intelligent Operating System Architecture Brief*. Internal document, March 2026. Available from the author on request.
- [3] Brooks, F.P. Jr. (1975). *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley.
- [4] Brooks, F.P. Jr. (1995). *The Mythical Man-Month, Anniversary Edition*. Addison-Wesley.
- [5] Brynjolfsson, E. & McAfee, A. (2014). *The Second Machine Age*. W.W. Norton.
- [6] Chandler, A.D. Jr. (1962). *Strategy and Structure*. MIT Press.
- [7] Chaturvedi, S. (2024). Quoted in CNBC, May 2024. Professor of Organisational Behaviour, Imperial College Business School.
- [8] Coase, R.H. (1937). 'The Nature of the Firm'. *Economica*, New Series, Vol. 4, No. 16, pp. 386–405.
- [9] Dunbar, R.I.M. (1992). 'Neocortex size as a constraint on group size in primates'. *Journal of Human Evolution*, 22(6), 469–493.
- [10] Edmondson, A. (2019). *The Fearless Organization*. Wiley.
- [11] Goleman, D. (1995). *Emotional Intelligence*. Bantam Books.
- [12] Graicunas, V.A. (1933). 'Relationship in Organisation'. In Gulick & Urwick (eds.), *Papers on the Science of Administration* (1937).
- [13] Gulick, L. & Urwick, L. (1937). *Papers on the Science of Administration*. Columbia University.
- [14] Huang, J. (2024). Keynote address, World Governments Summit, Dubai, February 2024.
- [15] Iansiti, M. & Lakhani, K.R. (2020). *Competing in the Age of AI*. Harvard Business Review Press.
- [16] ISBSG Repository Analysis (2012). Empirical study of 1,000+ software projects.
- [17] Lindenfors, P., Wartel, A., & Lind, J. (2021). 'Dunbar's number deconstructed'. *Biology Letters*, 17(3), 20200748.
- [18] López de Prado, M. (2018). *Advances in Financial Machine Learning*. John Wiley & Sons.
- [19] Mintzberg, H. (1979). *The Structuring of Organizations*. Prentice-Hall.
- [20] NVIDIA Corporation (2024). 'What Is Sovereign AI?'. NVIDIA Blog.
- [21] Porter, M.E. (1985). *Competitive Advantage*. Free Press.
- [22] Salovey, P. & Mayer, J.D. (1990). 'Emotional Intelligence'. *Imagination, Cognition and Personality*, 9(3), 185–211.
- [23] Schweik, C.M. & English, R. (2007). 'Brooks' Versus Linus' Law'. NCDG Working Paper 07-009.
- [24] Standish Group (2015). *CHAOS Report*.
- [25] Williamson, O.E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press.

Altimist® is a registered trademark of Altimist Ltd. Altimist Intelligent Operating System™ is a trademark of Altimist Ltd.