

A Mathematical Framework for Innovation: Cognitive Gauge Theory and Empirical Validation

Jian Gao

China Development Bank, No. 2, Fuxingmen Nei Street, Beijing 100031, China.

Email: gaojian810@163.com

Manuscript submitted January 10, 2026; accepted August 6, 2026; published October 23, 2026.

doi: 10.17706/ijapm.2026.16.2.43-51

Abstract: How do transformative ideas emerge from the bounded mind? This paper develops and empirically validates a geometric framework for the cognitive architecture of innovation. We first present the Pyramid Knowledge Gauge Theory (PKGT), which models cognition as a seven-layer hierarchical stack where gauge-theoretic principles govern information processing. From PKGT, we derive Cognitive Gauge Theory (CGT) for innovation, which formalizes idea generation as navigation of a cognitive state space, constrained by a gauge field (A_μ) of attentional costs, curvature ($F_{\mu\nu}$) producing path-dependence, and a finite discovery budget (κ). Across three experiments ($N=300$), we confirm CGT's predictions: exploratory search is path-dependent ($\eta^2 = 0.42$), individuals possess stable discovery budgets ($\bar{\kappa} = 3.15$), and conceptual recombinations follow predictable geometric patterns ($\Gamma = 0.45$). CGT provides cognitive microfoundations for economic theories and formalizes key philosophical ideas concerning knowledge structure and growth.

Keywords: bounded rationality, cognitive architecture, gauge theory, idea generation, innovation, mathematical modeling

1. Introduction

How do new ideas emerge from the finite human mind? Economic theory has modeled the effects of innovation, in particular by Schumpeter [1], Aghion and Howitt [2], Romer [3], yet we lack a formal, testable theory of its cognitive origins. This paper bridges this gap by introducing, formalizing, and empirically validating a geometric theory of cognitive innovation.

First, we present the Pyramid Knowledge Gauge Theory (PKGT), a foundational model of the mind as a hierarchical computational stack. PKGT posits that cognition operates across seven layers, from basic sensation to strategic reasoning, governed by gauge-theoretic principles that manage cognitive costs. Second, we derive Cognitive Gauge Theory (CGT) of Innovation, modeling idea generation as navigation in a cognitive state space $\mathcal{M}^*$ with three core constructs: a gauge field (A) quantifying activation costs, curvature (F) producing path dependence, and a discovery budget (κ) limiting exploratory effort. Theorem 1 establishes conditions for an innovative leap: $S[\gamma] \leq \kappa$ and $\oint_C A \neq 0$.

We then measure these constructs across three experiments ($N = 300$), demonstrating path dependence ($\eta^2 = 0.42$), stable discovery budgets ($\bar{\kappa} = 3.15$, $\sigma_\kappa = 0.89$), and predictable geometric distortions under load ($\Gamma = 0.45$). A hierarchical Bayesian model confirms CGT's superior predictive power over benchmark models.

2. Theoretical Foundations

2.1. Pyramid Knowledge Gauge Theory

PKGT proposes a seven-layer cognitive stack (Table 1), where gauge-theoretic principles govern information processing. A “gauge” represents freedom in how knowledge is represented locally; the gauge field quantifies cognitive cost of shifting perspective.

2.2. Architectural Configuration: Infrastructure and Superstructure

The PKGT hierarchy partitions into two functionally distinct strata: the knowledge infrastructure (Layers 1–4)—automatic, stable, path-dependent processing—and the knowledge superstructure (Layers 5–7)—deliberative, resource-intensive, context-sensitive reconfiguration. CGT emerges as the formal instantiation of superstructure dynamics at Layers 6–7, where infrastructure gauge fields $A^{(1-4)}$ constrain feasible trajectories in $\mathcal{M}^*$, while successful innovations recursively reshape infrastructure through learning, flattening curvature F along discovered paths.

2.3. Advantages of the PKGT Configuration

This hierarchical architecture confers three critical advantages for understanding innovation. First, modular decomposition separates automatic knowledge infrastructure from effortful superstructure, explaining why expertise can both enable and constrain creativity—automatized lower layers reduce baseline costs, while flexible upper layers enable novel recombination. Second, cost propagation ensures that inefficiencies at lower layers cascade upward, formally explaining why cognitive biases and fixedness impede innovation. Third, layer-specific gauge fields $A^{(l)}$ allow distinct measurement of different cognitive costs—attentional switching (Layer 1), belief updating (Layer 4), or conceptual integration (Layer 7)—enabling targeted diagnosis of innovative bottlenecks.

Table 1. Seven Layers of PKGT

Function	Layer	Name	Function & Gauge Principle
Infrastructure	1	Sensation	Raw perception; attention-switching cost
	2	Pattern	Feature extraction; reconfiguration cost
	3	Skill	Procedural execution; routine-switching cost
	4	Belief	World modeling; belief-updating cost
Superstructure	5	Preference	Option evaluation; trade-off deliberation
	6	Planning	Multi-step strategy; horizon cost
	7	Strategic	Conceptual recombination; integration cost

2.4. Deriving Cognitive Gauge Theory (CGT) for Innovation

Cognitive Gauge Theory is the formal instantiation of PKGT’s principles at Layers 6 and 7, specifically applied to the problem of innovative search. Where PKGT describes the full cognitive architecture, CGT provides the specific mathematics for modeling navigation within the high-level conceptual spaces where new ideas are formed.

We define the cognitive innovation space $\mathcal{M}^*$ as the sub-manifold of the full PKGT state space corresponding to activated concepts and their potential relationships. A point $z \in \mathcal{M}^*$ is a specific configuration of thoughts relevant to a creative problem.

Axiom 1 (Cognitive Innovation Space (from PKGT L7)). The space $\mathcal{M}^*$ of conceptual configurations relevant to a creative problem is a smooth, finite-dimensional manifold. It inherits its geometric structure from the gauge fields of the underlying PKGT layers.

Axiom 2 (Cognitive Gauge Field of Exploration (The Cost of Thought)). Attached to $\mathcal{M}^*$ is a gauge field one-form $A = A_\mu(z)dz^\mu$. The component $A_\mu(z)$ represents the marginal cognitive cost, in terms of

attentional effort, of shifting cognitive focus to, or strengthening the mental activation of, conceptual element μ when at cognitive state z . This field A is the direct manifestation of the PKGT cost principle at the recombinant reasoning layer.

Definition 1 (Cognitive Curvature and Recombinant Path-Dependence). The cognitive curvature is the two-form $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$. Non-zero curvature ($F_{\mu\nu} \neq 0$) means the cognitive cost of a sequence of thought depends on its order: exploring concept ν after μ has a different total cost than the reverse sequence. This formalizes cognitive fixedness and makes innovation path-dependent.

Axiom 3 (Cognitive Discovery Budget (Resource Constraint)). An agent possesses a finite cognitive discovery budget > 0 , a personal resource limit for deliberate, effortful exploration in $\mathcal{M}^*$. For any exploratory path γ undertaken, the incurred cost is constrained:

$$\int_\gamma A_\mu(z) dz^\mu \leq \kappa$$

This budget is a higher-level, aggregate constraint emerging from the resource limitations present across all PKGT layers.

Proposition 1 (The Innovative Cognitive Profile). Within the PKGT framework, individuals prone to innovation can be characterized by a specific cognitive profile: (i) a high discovery budget (κ) permitting extended search; (ii) a gauge field (A) shaped by diverse knowledge, creating low-cost “bridges” between distant conceptual domains; and (iii) a high tolerance for navigating curvature (F), allowing them to execute and trust non-commutative thought sequences that others abandon.

2.5. Core Theorems: The CGT Model of Idea Generation

From these axioms, which are grounded in PKGT, we derive the core mechanics of innovative thinking.

Theorem 1 (The Condition for an Innovative Leap). An innovative idea, represented by a novel path γ_{new} in $\mathcal{M}^*$, becomes cognitively feasible and is likely to be discovered when:

$$S[\gamma_{\text{new}}] = \int_{\gamma_{\text{new}}} A_\mu(z) dz^\mu \leq \kappa \text{ (Feasibility Condition)}$$

and γ_{new} forms a non-contractible loop with significant curvature flux, indicating a non-trivial recombination that could not have been found by incremental search.

Proof. We prove the two conditions for an innovative leap sequentially, grounding each step in the axioms and definitions of CGT and its interpretation within the PKGT framework.

Part 1: Feasibility Condition.

Let $\gamma_{\text{new}} : [0,1] \rightarrow \mathcal{M}^*$ be a piecewise-smooth path representing a specific sequence of conceptual explorations, with $\gamma_{\text{new}}(0) = z_i$ (initial cognitive state) and $\gamma_{\text{new}}(1) = z_f$ (final state containing the new idea).

By Axiom 2, the total cognitive cost incurred in traversing this path is given by the cognitive action:

$$S[\gamma_{\text{new}}] = \int_{\gamma_{\text{new}}} A_\mu(z) dz^\mu = \int_0^1 A_\mu(\gamma_{\text{new}}(t)) \dot{\gamma}_{\text{new}}^\mu dt,$$

where $\dot{\gamma}_{\text{new}}$ denotes the derivative with respect to the path parameter. This integral sums the marginal costs A_μ of activating each conceptual element along the specific trajectory of thought.

Axiom 3 posits the existence of a finite cognitive discovery budget κ . For the agent to complete the exploration along γ_{new} without exceeding their cognitive resources, it is necessary and sufficient that:

$$S[\gamma_{\text{new}}] \leq \kappa.$$

This is the feasibility condition. It formalizes the intuitive notion that an idea remains undiscovered if the mental effort required to conceive it exceeds an individual's or a community's available cognitive resources. Paths with cost $S[\gamma_{new}] > \kappa$ are cognitively inaccessible.

Part 2: Condition for Non-Trivial Innovation (Curvature Flux).

An innovative idea must represent a recombination that is not merely a linear extrapolation or trivial variation of existing knowledge. In the geometric language of CGT, a trivial extension of known ideas can be represented by a path γ_{known} that is contractible within a region of low or zero cognitive curvature.

Consider a closed loop $\mathcal{C}$ formed by concatenating the novel path γ_{new} with a segment of a pre-existing, known path γ_{known} that connects the same endpoints: $\mathcal{C} = \gamma_{new} \circ \gamma_{known}^{-1}$. The cognitive cost of traversing this loop is:

$$\oint_{\mathcal{C}} A_{\mu} dz^{\mu} = \int_{\gamma_{new}} A_{\mu} dz^{\mu} - \int_{\gamma_{known}} A_{\mu} dz^{\mu}.$$

By Stokes' theorem, this loop integral is equal to the flux of the cognitive curvature $F_{\mu\nu}$ through any surface Σ bounded by the loop $\mathcal{C}$:

$$\oint_{\mathcal{C}} A_{\mu} dz^{\mu} = \iint_{\Sigma} F_{\mu\nu} d\sigma^{\mu\nu},$$

where $\sigma^{\mu\nu}$ is the area element on Σ .

A non-trivial innovative leap occurs when the novel path γ_{new} is not cognitively equivalent to simply traversing known paths. This is captured by the condition that the loop $\mathcal{C}$ is not cognitively costless:

$$\oint_{\mathcal{C}} A_{\mu} dz^{\mu} = \iint_{\Sigma} F_{\mu\nu} d\sigma^{\mu\nu} \neq 0$$

A non-zero result implies two things mathematically and cognitively:

1. The curvature $F_{\mu\nu}$ is non-zero in the region Σ enclosed by the thought process. This means the cognitive costs are non-integrable (path-dependent), a signature of complex, interacting conceptual dimensions.
2. The novel path γ_{new} and the known path γ_{known} are not functionally equivalent; their difference has a measurable cognitive cost, indicating the new path involves a unique, non-reducible recombination of concepts. The magnitude of the flux quantifies the degree of conceptual novelty or "creative surprise" inherent in the idea.

Therefore, for an idea to be both feasible and innovative, the path γ_{new} that represents it must satisfy the dual criteria:

$$S[\gamma_{new}] \leq \kappa \text{ and } \oint_{\gamma_{new} \circ \gamma_{known}^{-1}} A \neq 0$$

This completes the proof.

The proof above establishes that local cognitive costs (encoded in A) integrate to a global action $S[\gamma]$ that must stay within the discovery budget κ ; genuine novelty requires non-zero curvature F , i.e., path-dependence. This link between microeconomic effort and macroeconomic innovative outcomes, mediated by Stokes' theorem, provides the blueprint for our empirical tests: we operationalize A via effort allocation, κ via search stopping points, and F via path-dependence effects. This geometric relationship is visualized in Fig. 1.




Fig. 1. Geometric foundations of CGT. Innovation as navigation on $\mathcal{M}^*$ with gauge field A_μ , curvature $F_{\mu\nu}$, and budget κ . Novel paths γ_{new} enclose non-zero curvature flux.

3. Empirical Validation

We now test whether CGT's constructs correspond to measurable phenomena through three experiments.

Experiment 1 ($N = 120$). used an alien plant classification task with interacting traits. The interaction between path complexity and environmental curvature was significant ($\eta^2 = 0.42$, $p < 0.001$). High-curvature landscapes produced 65% failure rate vs. 22% in low-curvature landscapes.

Experiment 2 ($N = 100$). used remote associates problems where participants purchased hints. Hint-purchasing revealed significant heterogeneity ($\bar{\kappa} = 3.15$, $\sigma = 0.89$), stable across problems (test-retest $r = 0.78$), correlating with Alternative Uses Test ($\rho = 0.41$, $p < 0.01$).

Experiment 3 ($N = 80$). used an economic simulation where participants learned a function then faced a shock under load. High load produced predictable errors, fitting CGT's failing connection model (Bayes Factor > 100), with $\hat{F} = 0.45[0.38, 0.52]$ (See Table 2 and Fig. 2).

Table 2. Synthesis of Experimental Results

Exp	Task & Key Finding	Effect Size	CGT Parameter
1	Alien plant classification: path-dependence amplified by curvature	$\eta^2 = 0.42$	$\hat{F} > 0$
2	Remote associates with hint purchase: stable individual budgets	$\bar{\kappa} = 3.15$; $\sigma = 0.89$	κ trait
3	Economic simulation with load: systematic geometric distortion	$\hat{F} = 0.45[0.38, 0.52]$	Connection fragility




Fig. 2. Empirical validation of CGT constructs. (a) Path-dependence ($\eta^2 = 0.42$), (b) Discovery budget distribution, (c) Connection fragility under load.

4. Toward a Predictive Model: Hierarchical Bayesian Integration

Having established empirical support for CGT's three core constructs—path-dependence (Exp. 1), discovery budgets (Exp. 2), and geometric distortion (Exp. 3)—we now integrate these findings into a unified predictive framework.

4.1. Synthesis Process

Stage 1: Why Hierarchical Bayesian Modeling? Standard NHST tells us whether an effect exists, but we need three deeper answers: (1) Can a single parameter set $(\alpha, \beta, \mu_\kappa, \sigma_\kappa, \Gamma)$ account for all three experiments? (2) Are discovery budgets κ stable individual differences? (3) Does CGT predict out-of-sample better than benchmarks? Hierarchical Bayesian modeling answers all three by pooling information across tasks.

Stage 2: Model Specification. We constructed a three-level hierarchy: (Level 1) likelihood—CGT-derived probability functions for each task; (Level 2) Participant—latent κ_i and $\bar{f}_i$ drawn from population distributions; (Level 3) Population—means $(\bar{\kappa}, \bar{f})$ and variances with weakly informative priors.

Stage 3: Model Comparison. We compared CGT against three benchmarks: Rational Inattention Sims [4]; Caplin and Dean [5], Drift-Diffusion Ratcliff [6], and Heuristic Tallying Gigerenzer and Goldstein [7], each fitted with equivalent hierarchical structures.

4.2. Results

The hierarchical model recovered coherent population-level CGT parameters (Table 3), confirming a

single geometric framework accounts for search, persistence, and recombination. CGT substantially outperformed benchmarks in out-of-sample prediction.

Table 3. Hierarchical Bayesian Estimates and Model Comparison

Parameter	Posterior Mean	95% CI	Interpretation
α	0.34	[0.28,0.40]	Avg. exploration cost
β	0.12	[0.08,0.16]	Avg. non-commutativity
μ_k	3.15	[2.80,3.50]	Avg. discovery budget
σ_k	0.89	[0.77,1.02]	Heterogeneity
Γ	0.45	[0.38,0.52]	Connection fragility
Model Comparison	BIC	Log-Lik	Pred. R^2
CGT	1256.7	-598.3	0.67
Rational Inattention	1342.1	-641.1	0.58
Drift-Diffusion	1398.5	-674.2	0.52
Heuristic Tallying	1423.9	-696.9	0.49

4.3. Key Insights

$\alpha = 0.34$ indicates moderate exploration costs;

$\beta = 0.12$ confirms non-commutative structure;

$\mu_k = 3.15$ suggests 3–4 recombinations before resource exhaustion;

$\sigma_k = 0.89$ implies meaningful individual differences (16% have $k > 4.0$). CGT decisively outperforms benchmarks ($\Delta BIC > 85$), with predictive $R^2 = 0.67$ vs. 0.49–0.58 (See Fig. 3). This advantage stems from capturing three phenomena benchmarks cannot: (a) non-linear error accumulation (Exp. 1), (b) heterogeneity in persistence (Exp. 2), and (c) distortion patterns under load (Exp. 3).




Fig. 3. Hierarchical Bayesian integration. Parameters show narrow credible intervals ($\hat{R} \leq 1.02$). CGT outperforms benchmarks ($R^2 = 0.67$ vs. 0.49–0.58).

Convergence diagnostics ($\hat{R} \leq 1.02$, ESS > 10,000) and posterior predictive checks ($p = 0.43$) confirm model adequacy (see Fig. 3). In sum, hierarchical Bayesian analysis validates CGT as a predictive, quantitative framework—not merely a descriptive metaphor.

5. Discussion: CGT as the Cognitive Engine of Innovation

Our experimental validation of CGT's constructs provides the first empirical evidence for a geometric model of innovative thinking derived from the PKGT architecture.

5.1. Bridging Architecture and Process and Indicating the Sources of Cognitive Geometry

PKGT provides the structural account of the mind's hierarchical layers, while CGT provides the dynamical account of how one moves through the high-level conceptual space defined by that structure. The measurable parameters of CGT ($A_\mu, F_{\mu\nu}, \kappa$) can be understood as dynamic signatures of an individual's PKGT configuration.

What shapes an individual's gauge field and budget? The knowledge infrastructure (PKGT Layers 1–4) shapes the baseline cost landscape. The knowledge superstructure (Layers 5–7) directly determines one's ability to manage curvature and budget resources for exploration.

5.2. CGT as a Unifying Framework: Synthesizing Schumpeter, Hayek, and Endogenous Growth Theory

CGT provides a unified mathematical language for three seminal economic traditions. Table 4 summarizes the synthesis.

Table 4. Synthesis of Economic Theories through CGT

Economic Tradition	Core Thesis	CGT Formalization
Schumpeterian	Radical recombination by entrepreneurs	High-curvature path ($F \geq 0$) traversed by high- κ agent
Hayekian	Utilization of dispersed, tacit knowledge	Idiosyncratic gauge fields $A^{(i)}$
Endogenous Growth	Investment in ideas as non-rival good	Collective $\uparrow \mu_\kappa, \downarrow F$

5.3. Philosophical Foundations: Formalizing Kant, Pepper, and Popper

CGT formalizes three philosophical frameworks (Table 5). Kant's categories correspond to gauge fields $A^{(i)}$ with $\mathcal{M}^*$ as phenomenal space. Pepper's root metaphors map to (A, F) configurations; switching worldviews requires high-curvature paths. Popper's Three Worlds interact via PKGT/CGT: World 2 traverses $\mathcal{M}^*$, generating World 3 objects that recursively lower A and flatten F —knowledge growth as curvature reduction. Thus, Kantian categories under resource-constrained exploration generate metaphor-guided evolution that reshapes cognition, closing the loop between mind, culture, and objective knowledge.

Table 5. Formalization of Philosophical Frameworks through CGT/PKGT

Philosophical Framework	Core Thesis	CGT/PKGT Formalization
Kantian	Mind structures phenomena via categories	PKGT layers implement categories as gauge fields $A^{(i)}$; $\mathcal{M}^*$ is phenomenal space
Pepperian	Worldviews built on root metaphors	Root metaphors correspond to specific A, F configurations
Popperian	Three worlds: physical, subjective, objective	World 2 generates World 3 objects that reshape future A, F

5.4. Limitations and Future Directions

This work focuses on individual cognition. Future work should scale to collective cognition, measure natural gauge fields of real-world knowledge domains, and link micro-cognitive geometries to macro-level innovation outcomes.

6. Conclusion

We presented PKGT, a seven-layer hierarchical cognitive architecture, and derived CGT, a geometric formalism modeling idea generation as navigation on $\mathcal{M}^*$ with gauge field A_μ , curvature $F_{\mu\nu}$, and

discovery budget κ . Theorem 1 established conditions for innovative leaps: $S[\gamma] \leq \kappa$ and $\phi_c \mathbf{A} \neq 0$. Experiments ($N = 300$) validated path-dependence ($\eta^2 = 0.42$), stable budgets ($\bar{\kappa} = 3.15$), and geometric distortion under load ($\Gamma = 0.45$). CGT unifies economic theories and formalizes philosophical frameworks, providing cognitive micro-foundations for innovation.

Conflict of Interest

The author declares no conflict of interest.

Acknowledgment

The author thanks the anonymous reviewers for their constructive feedback.

References

- [1] Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.
- [2] Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2), 323–351.
- [3] Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5,2), S71–S102.
- [4] Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690.
- [5] Caplin, A., & Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7), 2183–2203.
- [6] Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
- [7] Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669.

Copyright © 2026 by the authors. This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)).