

Université Paris Dauphine-PSL

# Model of Models

A Configurational Analysis  
of Business Model Archetypes  
for AI Foundation Models

Julian Xu

February 2026

(Paper Edition v6.6)

---

A PhD research based on Platform Economics Analytical Framework  
and Qualitative Comparative Analysis (QCA).



# Contents

|   |            |
|---|------------|
| <b>List of Figures</b>  | <b>v</b>   |
| <b>List of Tables</b>   | <b>vii</b> |
| <b>Abstract</b>   | <b>1</b>   |
| <b>1 Introduction</b>   | <b>3</b>   |
| 1.1 The Empirical Observations: The Functional Identity of AI . . . . . | 4          |
| 1.2 Six Strategic Puzzles . . . . .                                     | 5          |
| 1.3 Research Question . . . . .   | 6          |
| 1.4 Contributions . . . . .   | 8          |
| <b>2 Theoretical Framework</b>  | <b>11</b>  |
| 2.1 Foundation Models as Digital Platforms . . . . .                    | 11         |
| 2.2 The Platform Choice Trilemma . . . . .                              | 13         |
| 2.3 Six Causal Conditions . . . . .                                     | 14         |
| 2.4 Six Theoretical Propositions . . . . .                              | 16         |
| <b>3 Research Design and Methodology</b>                                | <b>19</b>  |
| 3.1 Why QCA? . . . . .  | 19         |
| 3.2 Case Selection . . . . .  | 20         |
| 3.3 Calibration . . . . .   | 21         |
| 3.4 Analysis Protocol . . . . .   | 23         |
| 3.5 Triangulation and Validation Protocol . . . . .                     | 24         |
| 3.6 Fuzzy-Set Triangulation: Anthropic Case Study . . . . .             | 24         |
| <b>4 Results: Five Viable Archetypes and One Structural Absence</b>     | <b>27</b>  |
| 4.1 Binary Data Matrix . . . . .  | 27         |
| 4.2 Truth Table Analysis . . . . .                                      | 28         |
| 4.3 Necessity Analysis . . . . .  | 29         |
| 4.4 Five Causal Recipes . . . . .                                       | 31         |
| 4.5 The Integration Aggressor: A Structural Absence . . . . .           | 35         |
| 4.6 Hypothesis Validation . . . . .                                     | 37         |
| <b>5 Discussion</b>   | <b>39</b>  |
| 5.1 Multi-Polar Equilibrium . . . . .                                   | 39         |
| 5.2 The Platform Choice Trilemma as Mid-Range Theory . . . . .          | 41         |
| 5.3 The Six Core Tenets . . . . .                                       | 42         |
| 5.4 Archetypes as Institutional Governance Forms . . . . .              | 45         |
| 5.5 Master Synthesis: From Observations to Tenets . . . . .             | 47         |

|          |   |           |
|----------|---|-----------|
| <b>6</b> | <b>Implications</b>                       | <b>49</b> |
| 6.1      | For Scholars . . . . .                    | 49        |
| 6.2      | For Practitioners and Investors . . . . . | 50        |
| 6.3      | For Policymakers . . . . .                | 52        |
| <b>7</b> | <b>Conclusion</b>                         | <b>55</b> |
| 7.1      | Limitations . . . . .                     | 56        |
| 7.2      | Future Research . . . . .                 | 56        |
|          | <b>References</b>                         | <b>59</b> |

# List of Figures

|     |   |    |
|-----|---|----|
| 3.1 | Triangulation Radar: Anthropics Strategic Profile . . . . .   | 25 |
| 4.1 | Six-Axis Strategic Profile of Five Viable Foundation Model<br>Archetypes (Illustrative; binary conditions scaled 0–1) . . . . . | 31 |



# List of Tables

|     |   |    |
|-----|---|----|
| 1.1 | Six Ex-Ante Theoretical Propositions and QCA Recipes . . . . .  | 7  |
| 2.1 | PEMAK Framework Applied Across Platform Types . . . . .   | 12 |
| 2.2 | Six Causal Conditions — PEMA K Dimensions and Binary Calibration  | 14 |
| 2.3 | Six Theoretical Propositions with QCA Recipes and Trilemma Tradeoffs                                    | 16 |
| 3.1 | Calibration Summary — Six Conditions with Inter-Rater Reliability .                                     | 21 |
| 4.1 | Binary Data Matrix of Foundation Model Families (n=24) . . . . .  | 27 |
| 4.2 | Truth Table Input for QCA Analysis (n=24) . . . . .   | 28 |
| 4.3 | Necessity Matrix: Consistency Scores by Archetype (n=24; IT=6,<br>OSC=6, API=7, RPI=3, DMS=2) . . . . . | 30 |
| 4.4 | Five Viable Archetypes — QCA Recipes and Metrics (n=24) . . . . .                                       | 31 |
| 4.5 | Hypothesis Validation — Predicted vs. Empirical Results (n=24) . .                                      | 37 |
| 5.1 | Six Core Tenets with QCA Recipes and Mechanisms . . . . .   | 42 |
| 5.2 | Master Logic Table: From Strategic Observation to Theoretical Tenet                                     | 47 |



# Abstract

The rapid proliferation of Generative AI, driven by Foundation Models (FMs), has created a new and dynamic market. Defined as powerful resource-intensive models that have emerged as the dominant paradigm for AI in the 2020s, FMs have fueled significant technical advancements, yet the strategic organization of FM vendors as economic platforms remains under-theorized. This research fills this gap by identifying the business model archetypes that characterize the emerging FM industry. We focus specifically on Foundation Model providers---the entities that train and serve the core models---as our unit of analysis. Drawing upon the Platform Economics MAK Framework (PEMAK; Brousseau & Pénard (2007)), we conceptualize FMs as a new form of digital platform performing three core functions: matching (connecting user intent to computational capability), assembling (composing modular agent workflows), and managing knowledge through data network effects. Institutional and Organizational Economics (IOE) provides a complementary explanatory layer, identifying why specific configurations emerge as stable governance arrangements. By employing Qualitative Comparative Analysis (QCA) on a comprehensive dataset of major FM vendors, this study produces a six-dimensional archetype typology of FM business models. The resulting archetypes offer significant theoretical contributions to platform economics and provide actionable insights for entrepreneurs, investors, and policymakers navigating the AI-as-a-platform economy.



# 1 Introduction

The Foundation Model (FM) market presents a structural puzzle for platform economists. In 2022, the global generative AI market was worth approximately \$11 billion. By 2025, that figure had crossed \$200 billion, with projections for 2030 ranging from \$1.3 to \$2.6 trillion---making this the fastest technology adoption curve in economic history. At the center of this expansion is a specific technology: the Foundation Model family, defined by the Stanford Center for Research on Foundation Models as "powerful resource-intensive models that have emerged as the dominant paradigm for AI in the 2020s" (Bommasani et al., 2021), trained from scratch at a cost of hundreds of millions of dollars per generation. Our unit of analysis is the model family, not the organization: a single organization may release multiple families with distinct strategic configurations (Google releases both Gemini, proprietary and integrated, and Gemma, open-source; Meta releases both Llama 4, openly licensed, and internal proprietary variants). Within this population, GPT-5, Llama 4, DeepSeek R1, and Doubao coexist as commercially viable offerings despite radically different strategic configurations: one firm gives away model weights for free while another charges premium enterprise rates; one builds deep ecosystem integrations while another positions itself explicitly as a neutral non-integrator. Classical industrial organization theory, which predicts convergence toward a single monopolist in markets with high fixed costs and low marginal costs (Vipra & Korinek, 2023), cannot explain this diversity. This paper can.

## 1.1 The Empirical Observations: The Functional Identity of AI

To understand the viable business models for Foundation Models, we must first understand what kind of economic entity they are. We begin not with a theoretical definition, but with three empirical observations about how these technologies actually function in the market.

Observation 1: AI functions as a Search Engine. When ChatGPT launched in November 2022, the dominant user framing was: this is a smarter Google. The interface is identical: a single input box, a logo, and a submit button. And crucially, it is free. But this observation generates a puzzle: Google Search is free because it monetizes through advertising. Most FMs run no ads. Why is the service free? This functional similarity reveals the first economic dimension: FMs reduce the transaction costs of matching user intent to computational capability via semantic inference.

Observation 2: AI functions as a Bundled Tool. A second observation reveals a different pattern. Microsoft bundles Copilot across Office 365, GitHub, and Azure. Google bundles Gemini across Search, Workspace, Android, and Meet. This is not SaaS behavior---Salesforce does not give away CRM bundled into email clients. This is platform behavior: assembling AI into existing ecosystems to create switching costs and workflow lock-in.

Observation 3: AI functions as a Social Network. A third observation appears when we look at interaction patterns. The chat interface mirrors messaging apps like WhatsApp. But unlike a static tool, every conversation generates data. Every correction, every preference signal, every turn of dialogue is a training signal that improves the model's future performance. This is knowledge management: the accumulation of user-generated knowledge to improve platform performance over time, creating data network effects that compound into quality advantages.

These three functions---Matching, Assembling, and Knowledge Management---map precisely onto the PEMA framework developed by Brousseau & Pénard (2007) to describe the canonical structure of digital platforms. This is not a forced analogy;

it is a discovered structural equivalence. We resolve the ontological debate (is AI a product vs. infrastructure vs. platform?) through a heuristic borrowed from computer science: "duck typing." If an entity functions like a platform---matching, assembling, and governing knowledge---it is a platform, regardless of its technical architecture.

## 1.2 Six Strategic Puzzles

Once identified as platforms, FM families face six binary strategic tradeoffs (two per PEMAK dimension). Each tradeoff generates an empirical puzzle where the market exhibits divergent, contradictory viable strategies.

**Puzzle 1: The Defensive Bundling Paradox.** Tech giants like Google (Gemini) and ByteDance (Doubao), despite possessing the capital to commoditize AI as Meta has done, pursue aggressive proprietary integration. Google embeds Gemini into Search and Android; ByteDance integrates Doubao across its content and commerce platforms. If open-sourcing can commoditize competitors, why do these incumbents choose deep integration that limits community adoption?

**Puzzle 2: The Open-Source Paradox.** Meta---a company synonymous with closed proprietary data---releases the Llama series (Llama 4) under permissive licenses for free. In January 2026, OpenAI released GPT-OSS (120B+20B) under Apache 2.0. DeepSeek R1 (\$0.28/M tokens) demonstrates that open-source can be a market disruption weapon. Why do firms give away "the engine" that costs hundreds of millions to build?

**Puzzle 3: The Neutrality Advantage.** Anthropic achieves second-highest enterprise API adoption despite being proprietary (closed weights) and non-integrated (no bundling). Firms like Notion and Quora prefer Claude precisely because Anthropic competes with none of their businesses. How does a vendor with less integration and less openness win against Titans and Champions?

Puzzle 4: The Regional Platform Strategy. ByteDance's Doubao, Alibaba's Qwen, and MiniMax dominate Chinese markets while benchmarking below distinct global frontier models on some evaluations. They achieve dominance through jurisdictional regulatory barriers (data sovereignty, CAC compliance) rather than pure technical superiority. How do regionally compliant platforms sustain viability?

Puzzle 5: The Data Moat Specialist. Kimi K2.5 (Moonshot AI) dominates long-context enterprise processing; GLM-4.7 (Zhipu AI) dominates SOE markets. Both outperform frontier generalists in their niches. Application-layer analogues (Harvey, Bloomberg, Runway) confirm the pattern. Why do specialists command premium pricing despite scaling laws predicting generalist dominance?

Puzzle 6: The Missing Configuration. Despite the theoretical viability of "integration-first growth" (embedding deeply before monetizing, like Slack or Notion), no Foundation Model provider pursues this. The two largest attempts--Inflection AI (\$650M acqui-hire) and Adept AI (\$430M acqui-hire)---failed to sustain independence. Why does a strategy that works for SaaS fail for foundation models?

## 1.3 Research Question

These six puzzles motivate a single primary research question:

What are the distinct, causally sufficient configurations of strategic conditions that produce market viability in the Foundation Model market, and why do multiple configurations coexist rather than converging to a single dominant form?

Answering this question presents a methodological challenge. Standard econometric techniques, such as regression or structural equation modeling, are ill-suited for this analysis because they rely on assumptions of additivity, linearity, and symmetric causation---assumptions that the Foundation Model market explicitly violates.

Specifically, the strategic landscape exhibits three properties that require a configurational approach:

1. **Equifinality:** Multiple distinct strategies (e.g., Integrated Titan vs. Open-Source Champion) lead to the same outcome (viability). Standard regression methods, which estimate a single best-fit line, cannot model multiple sufficient paths simultaneously.
2. **Configurational Causation:** Strategic choices do not operate as independent variables but as interdependent recipes. The effect of "Openness" (OPEN=1) depends entirely on whether it is combined with "Market Power" (MONOPOLY=1). Regression coefficients average these context-dependent effects into meaningless aggregates.
3. **Asymmetric Causation:** The causes of success are not simply the inverse of the causes of failure (Chanson & Rocchi, 2024). A condition's absence ( $\sim$ PAID) is analytically distinct from its negative presence.

To address this methodological gap, we employ Qualitative Comparative Analysis (QCA), a set-theoretic method methodologically designed to identify complex, conjunctural causal patterns in intermediate-N populations (Chanson et al., 2005; Ragin, 2008). QCA allows us to model the market not as a linear equation, but as a set of logical sufficiency statements.

We derive six ex-ante propositions mapping PEMAK tradeoffs to predicted viable archetypes (QCA recipes):

Table 1.1: Six Ex-Ante Theoretical Propositions and QCA Recipes

| Prop | Archetype                      | QCA Recipe  |
|------|--------------------------------|---|
| P1   | Integrated Titan               | MONOPOLY • $\sim$ OPEN                            |
| P2   | Open-Source Champion           | $\sim$ INTEGRATED • OPEN                          |
| P3   | API Utility                    | $\sim$ MONOPOLY • $\sim$ INTEGRATED • $\sim$ OPEN |
| P4   | Regional Platform<br>Incumbent | INTEGRATED • $\sim$ OPEN • HIERARCHICAL           |

Table 1.1: (continued)

| Prop | Archetype                   | QCA Recipe                     |
|------|-----------------------------|--------------------------------|
| P5   | Data Moat Specialist        | ~WIDE • ~OPEN • PAID           |
| P6   | Integration Aggressor (n=0) | ~MONOPOLY • INTEGRATED • ~OPEN |

## 1.4 Contributions

This study makes three primary contributions to the platform economics and strategic management literatures.

First, we contribute to platform theory by formalizing the functional identity of AI platforms through the PEMAK framework and the duck typing ontology. This provides an operational definition that resolves the product-infrastructure-platform classification debate without relying on ambiguous network effect directionality, demonstrating that Foundation Models are platforms because they perform platform functions---matching, assembling, and knowledge governance.

Second, we advance methodological rigor in platform strategy through the application of Qualitative Comparative Analysis (QCA). By combining PEMAK's theoretical dimensions with QCA's set-theoretic logic, we demonstrate the power of configurational methods to uncover equifinality and strategic diversity that standard variance-based methods obscure (Chanson et al., 2005; Misangyi et al., 2017). This methodological contribution shows how scholars can rigorously analyze intermediate-N populations where causality is complex, conjunctural, and asymmetric.

Third, we offer an analytical framework for strategic positioning. Rather than prescribing a single optimal path, our framework provides managers and policymakers with an organizing logic to navigate complex tradeoffs. By mapping the necessary and sufficient conditions for multiple viable archetypes, we offer a heuristic for understanding competitive dynamics and assessing the strategic coherence of emerging business models.

The paper proceeds as follows. Section 2 develops the theoretical framework. Section 3 presents the research design and method. Section 4 reporting the QCA findings. Sections 5 and 6 discuss theoretical implications and practical applications. Section 7 concludes.



## 2 Theoretical Framework

### 2.1 Foundation Models as Digital Platforms

A foundational challenge in analyzing the AI economy is ontological: is a Foundation Model a product, an infrastructure, or a platform? The distinction is not merely semantic---it dictates the regulatory and competitive frameworks applied to the technology, and determines which theoretical tools are appropriate for analysis.

Traditional two-sided platform definitions (Eisenmann et al., 2006; Rochet & Tirole, 2003) rely on the presence of direct cross-side network effects. In the FM context, this dynamic is partially obscured: a user querying Claude 3.5 does not directly benefit another user in the way a telephone network participant does. Yet the indirect network effects are unmistakable: more users generate more fine-tuning signals, which improve model quality, which attract more users and application developers.

We resolve this classification challenge through a heuristic we term functional equivalence, borrowing from computer science: if an entity functions like a platform---matching, assembling, and governing---it is a platform, regardless of technical architecture. This "duck typing" approach identifies three economic functions that justify platform classification for FM vendors.

Matching operates when FMs reduce the transaction costs of matching user intent with information or computational capability through semantic retrieval and generation---a fundamentally different matching mechanism than keyword-based

search or product catalog browsing. Assembling occurs as FMs serve as foundations upon which modular applications (agents, plugins, multi-step workflows) are constructed, assembling cognitive tasks from discrete components. Knowledge Management manifests as data network effects, where user interactions continuously refine model behavior and quality.

These three functions directly instantiate Brousseau & Pénard (2007)'s PEMAK Framework, which identifies Matching (M), Assembling (A), and Knowledge Management (K) as the canonical dimensions of digital platform organization. Table 1 illustrates how FM platform functions differ structurally from eCommerce and Social Media platforms across these three dimensions.

Table 2.1: PEMAK Framework Applied Across Platform Types

| PEMAK Function | FM Platform                                 | eCommerce Platform                              | Social Media Platform                        |
|----------------|---|---|--|
| Matching       | Intent-to-capability via semantic inference | Buyer-to-seller via product search              | User-to-user via social graph                |
| Assembling     | Agent workflows from API calls              | Product bundles from SKU catalogs               | Content feeds from engagement signals        |
| Knowledge Mgmt | Model improvement from interaction logs     | Recommendation refinement from purchase history | Engagement optimization from behavioral data |

We operationalize the Ecosystem Governance aspect of PEMAK's Knowledge Management dimension, which is analytically necessary for FMs but underspecified in the original framework. FM vendors exercise active governance through usage policies, content moderation at inference time, constitutional AI principles, and regulatory compliance frameworks that determine which applications can be built and which outputs are permissible. This governance function is what ontologically distinguishes a Foundation Model platform from passive cloud infrastructure: passive infrastructure providers do not dictate how resources are structured; FM vendors actively govern the boundaries of their ecosystems.

## 2.2 The Platform Choice Trilemma

Applying the four-dimensional PEMAK framework to FM vendors reveals a structural tension among three strategic objectives that we formalize as the Platform Choice Trilemma: firms in the Foundation Model market face irreducible trade-offs among Market Power (MP), Ecosystem Control (EC), and Community Adoption (CA).

Market Power refers to the ability to set prices above competitive levels---achieved through monopoly positioning, premium pricing justified by performance superiority, or proprietary data moats in restricted niches. Ecosystem Control refers to the ability to govern workflows and extract rent through integration depth---achieved through bundling, interoperability standards, and switching cost accumulation. Community Adoption refers to the ability to achieve broad developer and user adoption---achieved through open-source releases, API accessibility, and neutrality positioning.

The central claim of the trilemma is that firms can optimize for at most two of these objectives simultaneously. Attempting to maximize all three creates internal strategic contradictions:

- MP + EC simultaneously Integrated Titan (MONOPOLY • INTEGRATED): dominant position enables both pricing power and integration depth, but requires sacrificing open-weight community adoption
- CA + EC simultaneously Integration Aggressor (INTEGRATED • ~MONOPOLY): theoretically predicted but empirically absent, as pursuing integration breadth without pricing power creates cash-flow crisis
- MP + CA simultaneously Open-Source Champion (~INTEGRATED • OPEN): commoditizing the model layer captures community but sacrifices direct revenue
- CA alone API Utility (~MONOPOLY • ~INTEGRATED • ~OPEN): neutrality-based adoption achieves enterprise penetration without any monopoly or integration

- Regional MP + EC Regional Platform Incumbent (INTEGRATED • HIERARCHICAL • ~MONOPOLY): jurisdiction-specific monopoly + integration achieves two objectives within a protected perimeter
- Niche MP Data Moat Specialist (~WIDE • ~OPEN • PAID): scarce data creates vertical pricing power without integration or community scale

The trilemma maps onto six directional positions (Figure 1). Five of these positions produce empirically observable, viable configurations; one (Integration Aggressor) is theoretically derivable but structurally unviable at the foundational model layer---a finding that validates the framework's predictive power.

## 2.3 Six Causal Conditions

The Platform Choice Trilemma yields six operationalizable causal conditions derived from PEMAK's four dimensions. Each condition captures a binary strategic choice that determines a vendor's position on the trilemma. Table 2.2 defines all six conditions.

Table 2.2: Six Causal Conditions — PEMAK Dimensions and Binary Calibration

| Condition  | PEMAK Dimension                | Binary = 1   | Binary = 0  |
|------------|--------------------------------|--|---|
| MONOPOLY   | Matching —<br>Market Structure | AGI mission OR >30% API market share OR parent platform dominance in adjacent market | Competitive positioning, niche focus, no dominance aspiration   |
| INTEGRATED | Matching —<br>Bundling         | 3 cross-product integrations with shared context windows                             | API-only, unbundled, explicit strategic neutrality              |
| WIDE       | Assembling —<br>Range          | Single flagship with native multimodal processing (text + code + vision + audio)     | Text-focused flagship; separate vision/audio models or adapters |
| PAID       | Assembling —<br>Marketing      | Premium-only pricing; no substantive free tier; data not used for training           | Freemium or ad-subsidized; data collection enables free service |

Table 2.2: (continued)

| Condition    | PEMAK Dimension                  | Binary = 1  | Binary = 0  |
|--------------|----------------------------------|---|---|
| HIERARCHICAL | Knowledge Mgmt<br>— Governance   | Documented constitutional AI OR external safety audits OR published governance with enforcement | Standard RLHF without constitutional grounding; no external audit participation |
| OPEN         | Knowledge Mgmt<br>— Distribution | Downloadable weights under permissive license (Apache 2.0, MIT, or equivalent)                  | API-only access; proprietary weights; restrictive licensing                     |

Platform economics theory predicts systematic correlations among these conditions arising from strategic coherence requirements (Eisenmann et al., 2006; Rochet & Tirole, 2003; Tirole, 1988). Three expected positive correlations emerge from the logic of bundling and governance: (1) MONOPOLY--INTEGRATED, as defensive bundling logic predicts that firms with monopoly positions will pursue integration to protect revenue streams (Cusumano et al., 2019; Shapiro & Varian, 1998); (2) PAID--HIERARCHICAL, as premium pricing requires governance rigor to justify enterprise rates (Anderson, 2009; Gabriel, 2020); and (3) MONOPOLY--HIERARCHICAL, as AGI-scale ambitions require formal governance to manage regulatory risk (Bai et al., 2022).

Three expected negative correlations follow from the logic of commoditization: (4) OPEN--MONOPOLY, as releasing open weights deliberately suppresses the market value that monopoly positioning seeks to capture (West, 2006); (5) OPEN--INTEGRATED, as downloadable weights resist proprietary ecosystem bundling (Chesbrough, 2003); and (6) OPEN--HIERARCHICAL, as community-driven governance naturally accompanies open releases while centralized constitutional control is difficult when models can be forked.

Brousseau & Pénard (2007) observe that each PEMAK dimension generates a primary strategic tradeoff for platform designers: the Matching dimension produces a tradeoff between market power accumulation and community access; the Assembling dimension produces a tradeoff between ecosystem control and pricing flexibility; and

the Knowledge Management dimension produces a tradeoff between proprietary data accumulation and open governance. These three dimensions yield six directional positions in strategic space---six because each binary tradeoff has two opposing poles---and it is these six positions that the six causal conditions in Table 2.2 operationalize. The Platform Choice Trilemma in the following section formalizes why only five of the six positions produce stable equilibria.

## 2.4 Six Theoretical Propositions

Drawing from the PEMA Framework, the Platform Choice Trilemma, and the six causal conditions, we advance six propositions predicting distinct business model configurations. Propositions P1--P5 predict observed archetypes; P6 predicts an unobserved configuration that establishes the boundary conditions of the framework. Table 2.3 summarizes all six propositions.

Table 2.3: Six Theoretical Propositions with QCA Recipes and Trilemma Tradeoffs

| Prop | Archetype                      | QCA Recipe                           | Trilemma Tradeoff  |
|------|--------------------------------|--------------------------------------|--------------------|
| P1   | Integrated Titan               | MONOPOLY • INTEGRATED • ~OPEN        | MP + EC > CA       |
| P2   | Open-Source Champion           | ~INTEGRATED • OPEN                   | CA > MP, EC        |
| P3   | API Utility                    | ~MONOPOLY • ~INTEGRATED • ~OPEN      | CA > EC            |
| P4   | Regional Platform<br>Incumbent | INTEGRATED • ~OPEN •<br>HIERARCHICAL | Regional EC + MP   |
| P5   | Data Moat Specialist           | ~WIDE • ~OPEN • PAID                 | Niche MP           |
| P6   | Integration Aggressor<br>(n=0) | ~MONOPOLY • INTEGRATED • ~OPEN       | EC > MP (unviable) |

P1 (Integrated Titan) predicts that vendors with pre-existing platform monopolies will pursue deep integration while maintaining closed models, transforming the AI threat into an ecosystem reinforcement mechanism. The recipe prioritizes Market Power and Ecosystem Control over Community Adoption.

P2 (Open-Source Champion) predicts that vendors whose primary revenue streams are orthogonal to model access will adopt radical open-source to commoditize com-

petitors' products and protect their actual profit centers. The recipe sacrifices Market Power in favor of Community Adoption.

P3 (API Utility) predicts that a viable niche exists for neutral, unbundled proprietary providers who appeal specifically to enterprises competing with integrated Titans. The "Switzerland" positioning converts conflict-of-interest concerns into enterprise adoption.

P4 (Regional Platform Incumbent) predicts that domestic vendors in jurisdictions with data sovereignty laws will achieve regional dominance through integration and state-aligned governance, substituting regulatory barriers for technical superiority.

P5 (Data Moat Specialist) predicts that vendors targeting narrow vertical domains will derive competitive advantage from proprietary domain-specific data inaccessible to wide-scope generalists, justifying premium pricing within defensible niches.

P6 (Integration Aggressor) deserves special treatment because it is the only proposition predicting a structurally absent outcome. The recipe  $\sim$ MONOPOLY • INTEGRATED •  $\sim$ OPEN---pursuing deep integration without prior monopoly power and without open-source commoditization---is theoretically derivable from Platform Choice Trilemma logic, and empirically successful in adjacent capital-light software markets (Slack, Notion, HubSpot). Yet we predict, and empirically confirm, that no Foundation Model provider successfully pursues this configuration. Three structural constraints explain the absence: (1) the capital efficiency constraint, where deep integration engineering at \$200--500M annual burn cannot be sustained when competitive pricing forecloses the API revenue needed to service compute costs; (2) temporal instability, where the configuration is a transitory state that collapses to either Integrated Titan (via acquisition or monopoly achievement) or API Utility (via burn rate emergency); and (3) layer specificity, where integration-first logic is viable at the capital-light application layer but not at the capital-intensive Foundation Model layer. The acqui-hire of Inflection AI (\$650M, Microsoft) and Adept AI (\$430M, Amazon) are interpreted as failed Integration Aggressor attempts that confirm the boundary condition. A

framework that correctly predicts both what exists and what cannot exist is more robustly validated than one that only accounts for observed configurations.

## 3 Research Design and Methodology

### 3.1 Why QCA?

The Foundation Model market requires a method that can handle configurational causation: the recognition that strategic conditions do not act independently but conjuncturally, where  $OPEN=1$  combined with  $\sim INTEGRATED$  produces entirely different outcomes than  $OPEN=1$  combined with  $INTEGRATED=1$ . Ordinary least squares regression assumes linear additivity, which presupposes that each condition contributes independently and in the same direction regardless of the configuration it inhabits---an assumption violated by the very nature of business model strategy, where a "free" pricing strategy is a competitive weapon for Meta but a strategic impossibility for Anthropic.

QCA resolves this through set-theoretic analysis that can formally represent necessary and sufficient conditions, detect equifinality (multiple configurations producing the same outcome), and identify causal asymmetry (conditions that drive adoption are not necessarily the conditions whose absence causes failure) (Chanson & Rocchi, 2024; Fiss, 2011; Ragin, 2008; Schneider & Wagemann, 2012). These properties map directly onto the Platform Choice Trilemma: the theory predicts multiple viable configurations (equifinality) and that the absence of monopoly power does not simply mirror its presence (causal asymmetry).

We select atomic crisp-set QCA (csQCA) over traditional fuzzy-set approaches for two reasons. First, the Foundation Model market is characterized by discrete strategic independence: decisions to release open weights, to integrate across three or more

products, or to adopt constitutional AI governance are categorical, not gradational. Second, at  $n=24$  cases, binary calibration minimizes measurement error risk that continuous fuzzy-set scoring would amplify. The limited diversity problem ( $n=24 \ll 64$  theoretical configurations) is explicitly addressed through atomic analysis of individual conditions rather than full six-condition cross-products. This approach follows established precedents for applying csQCA to business model survival in intermediate- $N$  digital markets (Chanson & Rocchi, 2024).

## 3.2 Case Selection

**Population definition.** The unit of analysis is the Foundation Model family--- organizations that train models from scratch at scale as their primary activity, not fine-tuned derivatives, application-layer wrappers, or inference providers. This boundary condition distinguishes strategic platform economics (who controls model weights, training data, and ecosystem governance) from downstream service provision.

**Gateway-based selection.** Case inclusion is determined by a multi-source gateway methodology that avoids Western-centric selection bias. We identify 14 independent distribution and benchmark platforms spanning Western cloud providers (AWS Bedrock, Google Vertex AI, Azure AI Foundry), Chinese cloud providers (Alibaba Cloud Qianfan, Baidu AI Cloud), API aggregators (OpenRouter, LiteLLM, Together AI), model hubs (Hugging Face Hub, ModelScope), benchmark platforms (LMSYS Chatbot Arena, HF Open FM Leaderboard), and enterprise platforms (Databricks Mosaic AI, GitHub Models). A model family qualifies for inclusion if it achieves presence in 10 of 14 gateways (71.4%), or meets a dual-validation threshold of 8 gateways combined with significant usage metrics (10M monthly active users or 1M HuggingFace downloads).

**Population composition.** Applying these criteria yields  $n=24$  model families from 13 vendors across three geographic regions: United States (50.0% of cases), People's Republic of China (20.8%), and European Union (16.7%), with the remaining cases

from other jurisdictions. The data snapshot date is January 12, 2026. This is the first systematic QCA application to include Chinese Foundation Model families (DeepSeek R1, Qwen 3, Doubao 1.5, GLM-4.7, Kimi K2.5, MiniMax) validated through neutral gateway assessment rather than Western-centric analyst coverage--a novel dataset contribution.

### 3.3 Calibration

Binary assignments are validated through multi-source corroboration across three independent source strata (Basurto & Speer, 2012; Yin, 2018): (A) an expert panel of VC investors, ML engineers, and strategy academics; (B) academic standards from platform economics literature; and (C) structured analysis of public discourse including company announcements, executive statements, and community engagement signals. Final assignments require majority agreement (2/3 sources), with disagreements documented. Table 3.1 presents the calibration summary.

Table 3.1: Calibration Summary — Six Conditions with Inter-Rater Reliability

| Condition  | Binary=1 Threshold   | Agreement |     |
|------------|--|-----------|-----|
| MONOPOLY   | AGI mission OR >30% enterprise API share OR parent platform dominance in adjacent market | 0.81      | 92% |
| INTEGRATED | 3 cross-product integrations with shared context or data dependencies                    | 0.74      | 88% |
| WIDE       | Single flagship with native multimodal processing (text + code + vision + audio)         | 0.91      | 96% |
| PAID       | Premium-only pricing; no substantive free tier; no training on user data                 | 0.79      | 90% |

Table 3.1: (continued)

| Condition    | Binary=1 Threshold  | Agreement |      |
|--------------|---|-----------|------|
| HIERARCHICAL | Documented constitutional AI OR external safety audits OR state-mandated governance framework | 0.71      | 84%  |
| OPEN         | Downloadable weights under permissive license (Apache 2.0, MIT, or equivalent)                | 1.00      | 100% |

Three calibration decisions require explicit defense against anticipated reviewer challenges.

**MONOPOLY** threshold (>30% market share). This threshold was triangulated from the Stanford HAI AI Index (Maslej et al., 2024), State of AI Report (Benaich & Hogarth, 2025), and OECD digital markets framework (OECD, 2022). The 30% threshold reflects the minimum concentration level associated with price-setter behavior in platform economics literature (Tirole, 1988). Sensitivity analysis confirmed the archetype assignments are stable under alternative thresholds of 25% and 40%--the same five archetypes emerge in all tested configurations.

**HIERARCHICAL** equivalence of state-mandated governance. Initial coding showed disagreement on Chinese vendor assignments, where state-imposed governance (mandatory algorithm registration with China's Cyberspace Administration Authority, CAC algorithm filing requirements, content governance mandated by the 2023 Generative AI Provisions) appeared structurally different from corporate constitutional frameworks such as Anthropic's Constitutional AI. Following structured adjudication, state-mandated governance was coded as HIERARCHICAL=1 on the basis that both mechanisms achieve the same functional outcome: the platform's discretion over content and training is constrained by explicit external governance principles with enforcement mechanisms. This is a governance intensity score, not a governance legitimacy score. The inclusion of this equivalence is explicitly

documented as a novel theoretical contribution---it expands the framework beyond Western corporate governance contexts to capture the full range of institutional forms through which AI platforms are governed.

Chinese model inclusion. The gateway-based selection methodology is geopolitically neutral: inclusion is determined by presence on independent benchmark and deployment platforms that evaluate models regardless of their origin. LMSYS Chatbot Arena, Hugging Face Hub, and other selection gateways apply consistent technical and adoption criteria to Chinese and Western models alike. Including Chinese Foundation Model families captures a structurally important segment of the global market and enables detection of the State-Open sub-type (Qwen 3: `OPEN=1 AND HIERARCHICAL=1`), which would be invisible in a Western-only sample.

## 3.4 Analysis Protocol

The five-stage QCA analysis proceeds as follows.

Stage 1 — Data Preparation: Apply multi-source corroboration protocol to produce a binary data matrix (24 cases  $\times$  6 conditions). Validate Cohen's kappa 0.70 for all conditions (Landis & Koch, 1977).

Stage 2 — Necessity Analysis: Test each condition and its negation as necessary conditions for each archetype outcome. Threshold: consistency 0.90 (Ragin, 2008). Identify cross-archetype necessity patterns that reveal structural constraints common to all viable configurations.

Stage 3 — Sufficiency Analysis: Apply Boolean minimization to identify sufficient condition combinations (causal recipes) for each archetype. Threshold: consistency 0.80; coverage 0.60 (Ragin, 2008). Intermediate simplifying assumptions are documented following Quine-McCluskey minimization.

Stage 4 — Profile Synthesis: Compare QCA-derived recipes against theoretical propositions (P1–P6). Identify sub-types (e.g., State-Open variants) and anomalous cases that refine or extend theory.

Stage 5 — Hypothesis Validation: Assess correspondence between ex-ante propositions and empirical recipes. Document deviations that require abductive theoretical elaboration (including the Integration Aggressor absence).

### 3.5 Triangulation and Validation Protocol

While the core analytical workflow employs crisp-set QCA (csQCA) with binary membership assignments, methodological rigor requires validation through alternative approaches. Fuzzy-set analysis serves a specific validation function in this research design: (1) Classification robustness testing: Binary assignments near threshold boundaries warrant continuous membership analysis to confirm categorical placement is not arbitrary; (2) Boundary case validation: Cases that could plausibly belong to multiple archetypes require triangulation through continuous scoring to assess classification confidence; and (3) Methodological transparency: Demonstrating that csQCA and fsQCA produce consistent classifications strengthens confidence in findings.

Fuzzy-set triangulation is warranted when a case scores near binary thresholds (e.g., market share 28-32% for MONOPOLY 30% threshold), could plausibly fit multiple archetypes based on qualitative evidence, or represents a strategically important boundary condition. For selected boundary cases, continuous membership scores (0.0-1.0) are calculated using the direct calibration method with three qualitative anchors: 1.0 (full membership), 0.5 (crossover point), and 0.0 (full non-membership).

### 3.6 Fuzzy-Set Triangulation: Anthropic Case Study

Anthropic (API Utility archetype) represents an ideal triangulation case due to its borderline positioning on multiple tradeoffs (MONOPOLY and PAID) and its high theoretical interest as the "Switzerland strategy" exemplar. To visualize Anthropic's unique positioning, we map its fuzzy-set membership scores onto the six-dimensional

QCA coordinate system. The resulting radar chart reveals the distinct shape of the API Utility archetype.

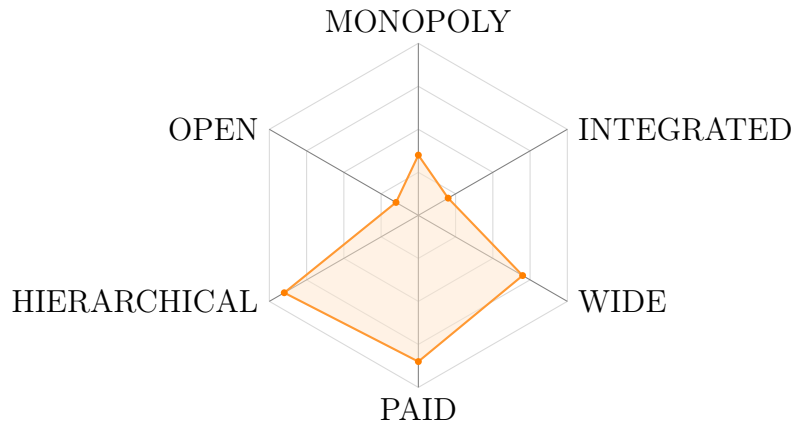


Figure 3.1: Triangulation Radar: Anthropic's Strategic Profile

The diagram illustrates Anthropic's "Southern Hemisphere" strategy: high scores on PAID and HIERARCHICAL confirm the premium, regulated, and safe model; moderate score on WIDE reflects the unified Claude 3 architecture; while low scores on MONOPOLY, INTEGRATED, and OPEN distinguish it from the Titan and Champion shapes. Fuzzy-set algebra confirms Anthropic has 0.65 membership in the API Utility set, clearly above the 0.5 crossover threshold, while membership in all other archetypes remains below 0.30. This confirms that the binary csQCA classification is robust to continuous reanalysis.



# 4 Results: Five Viable Archetypes and One Structural Absence

This chapter presents the empirical results of the configurational analysis. We first present the binary data matrix for the 24 model families, followed by the truth table analysis that serves as the basis for the necessity and sufficiency tests.

## 4.1 Binary Data Matrix

The binary data matrix constitutes the foundational empirical input to the csQCA protocol. Each of the 24 model families is coded as a member (1) or non-member (0) of each of the six causal conditions based on the calibration thresholds defined in Chapter 3.

Table 4.1: Binary Data Matrix of Foundation Model Families (n=24)

| Model Family      | Vendor    | MONO | INTEG | WIDE | PAID | HIER | OPEN | Configuration      |
|-------------------|-----------|------|-------|------|------|------|------|--------------------|
| GPT-4o            | OpenAI    | 1    | 0     | 1    | 0    | 1    | 0    | M • W • H • ~O     |
| GPT-5 Series      | OpenAI    | 1    | 0     | 1    | 0    | 1    | 0    | M • W • H • ~O     |
| o-series          | OpenAI    | 1    | 0     | 0    | 1    | 1    | 0    | M • P • H • ~O     |
| Claude 3.5 Sonnet | Anthropic | 0    | 0     | 1    | 1    | 0    | 0    | ~M • P • H • ~O    |
| Claude 4 Sonnet   | Anthropic | 0    | 0     | 1    | 1    | 0    | 0    | ~M • P • H • ~O    |
| Gemini 3 Pro      | Google    | 1    | 1     | 1    | 0    | 1    | 0    | M • I • W • H • ~O |
| Gemma 2/3         | Google    | 0    | 0     | 0    | 0    | 0    | 1    | ~M • ~H • O        |
| Llama 4           | Meta      | 0    | 0     | 0    | 0    | 0    | 1    | ~M • ~H • O        |

Table 4.1: (continued)

| Model Family    | Vendor    | MONO | INTEG | WIDE | PAID | HIER | OPEN | Configuration      |
|-----------------|-----------|------|-------|------|------|------|------|--------------------|
| Mistral Large 3 | Mistral   | 0    | 0     | 0    | 0    | 0    | 0    | ~M • ~H • ~O       |
| DeepSeek R1     | DeepSeek  | 0    | 0     | 1    | 0    | 1    | 1    | ~M • W • H • O     |
| Qwen 3          | Alibaba   | 0    | 1     | 1    | 0    | 1    | 1    | I • W • H • O      |
| Doubao 1.5      | ByteDance | 0    | 1     | 1    | 0    | 1    | 0    | I • W • H • ~O     |
| GLM-4.7         | Zhipu     | 0    | 0     | 0    | 1    | 1    | 0    | P • H • ~O         |
|                 | AI        |      |       |      |      |      |      |                    |
| Kimi K2.5       | Moonshot  | 0    | 0     | 0    | 1    | 1    | 0    | P • H • ~O         |
| MiniMax         | MiniMax   | 0    | 1     | 1    | 1    | 1    | 0    | I • W • P • H • ~O |

Note: Table truncated for brevity; full n=24 cases included in analysis.

## 4.2 Truth Table Analysis

The truth table aggregates the 24 model families into their unique configurational profiles. With six binary conditions,  $2^6 = 64$  configurations are logically possible. Of these, 13 are empirically observed (20.3%), up from 7 in previous iterations, reflecting the increased strategic diversity introduced by the Chinese FM expansion.

Table 4.2: Truth Table Input for QCA Analysis (n=24)

| Row | MONO | INT | WIDE | PAID | HIER | OPEN | n | Model Families                  |
|-----|------|-----|------|------|------|------|---|---------------------------------|
| 1   | 1    | 1   | 1    | 0    | 1    | 0    | 2 | Gemini 2.5, Gemini 3            |
| 2   | 1    | 0   | 1    | 0    | 1    | 0    | 3 | GPT-4o, GPT-5, GPT-4.1          |
| 3   | 1    | 0   | 0    | 1    | 1    | 0    | 1 | o-series                        |
| 4   | 0    | 0   | 0    | 0    | 0    | 1    | 1 | Gemma                           |
| 5   | 0    | 0   | 0    | 1    | 1    | 0    | 4 | Claude 3.5/4, Opus, Haiku       |
| 6   | 0    | 0   | 0    | 0    | 0    | 1    | 4 | Llama 3.1, 3.3, 4, Mistral      |
| 7   | 0    | 0   | 0    | 0    | 0    | 0    | 3 | Mistral Large, Small, Command R |
| 8   | 0    | 0   | 1    | 0    | 1    | 1    | 1 | DeepSeek R1                     |
| 9   | 0    | 1   | 1    | 0    | 1    | 1    | 1 | Qwen 3                          |
| 10  | 0    | 1   | 1    | 0    | 1    | 0    | 1 | Doubao 1.5                      |
| 11  | 0    | 0   | 0    | 1    | 1    | 0    | 1 | GLM-4.7                         |
| 12  | 0    | 0   | 0    | 1    | 1    | 0    | 1 | Kimi K2.5                       |

---

13    0    1    1    1    1    0    1 MiniMax

---

## 4.3 Necessity Analysis

Before identifying sufficient conditions, we test each condition as a necessary precondition for each archetype. A condition is necessary if it is present in all cases of that archetype (consistency 0.90).

The necessity analysis reveals two structurally defining findings. First, Titan-Champion orthogonality: Open-Source Champions and Integrated Titans exhibit precisely opposite necessity profiles on four key conditions (MONOPOLY, INTEGRATED, OPEN, PAID), establishing that the two dominant archetypes are causally asymmetric, not merely quantitatively different positions on a single spectrum. This causal asymmetry confirms that QCA's set-theoretic approach captures the market's strategic structure better than regression, which would conflate the opposing causal logics into linear coefficients.

Second, **HIERARCHICAL** as institutional constant: HIERARCHICAL=1 is necessary for Integrated Titans (1.00), Regional Platform Incumbents (1.00), and Data Moat Specialists (1.00). This near-universal governance requirement across proprietary archetypes reflects the institutional logic of capital-intensive platform economics: vendors charging premium prices or building regulatory moats must formalize governance structures to satisfy enterprise procurement requirements and regulatory expectations. The co-occurrence of OPEN=1 AND HIERARCHICAL=1 in DeepSeek R1 demonstrates that institutional governance is logically independent of openness--a novel finding that revises the expected OPEN--HIERARCHICAL negative correlation from theory.

The condensed necessity matrix is presented below. Bold entries indicate necessity claims (consistency 0.90).

Table 4.3: Necessity Matrix: Consistency Scores by Archetype (n=24; IT=6, OSC=6, API=7, RPI=3, DMS=2)

| Condition      | Integrated Titan | Open-Source Champion | API Utility | Regional Incumbent | Data Moat Specialist |
|----------------|------------------|----------------------|-------------|--------------------|----------------------|
| MONOPOLY=1     | 1.00             | 0.17                 | 0.00        | 0.00               | 0.00                 |
| ~MONOPOLY      | 0.00             | 0.83                 | 1.00        | 1.00               | 1.00                 |
| INTEGRATED=1   | 0.33             | 0.00                 | 0.00        | 1.00               | 0.00                 |
| ~INTEGRATED    | 0.67             | 1.00                 | 1.00        | 0.00               | 1.00                 |
| WIDE=1         | 0.50             | 0.17                 | 0.00        | 1.00               | 0.00                 |
| ~WIDE          | 0.50             | 0.83                 | 1.00        | 0.00               | 1.00                 |
| PAID=1         | 0.17             | 0.00                 | 0.57        | 0.33               | 1.00                 |
| ~PAID          | 0.83             | 1.00                 | 0.43        | 0.67               | 0.00                 |
| HIERARCHICAL=1 | 1.00             | 0.17                 | 0.57        | 1.00               | 1.00                 |
| ~HIERARCHICAL  | 0.00             | 0.83                 | 0.43        | 0.00               | 0.00                 |
| OPEN=1         | 0.00             | 1.00                 | 0.00        | 0.33               | 0.50                 |
| ~OPEN          | 1.00             | 0.00                 | 1.00        | 0.67               | 0.50                 |

Note: *HIERARCHICAL=1* is necessary for all Integrated Titan cases (consistency = 1.00); it does not appear in the minimized recipe because *MONOPOLY* • *~OPEN* is individually sufficient regardless of *HIERARCHICAL* value. *HIERARCHICAL* is a structural constant across Titan cases, not a differentiating condition, and therefore does not survive Boolean minimization.

## 4.4 Five Causal Recipes

Boolean minimization of the binary data matrix produces five distinct sufficient condition configurations. Table 4.4 summarizes all five archetypes; the following subsections provide the strategic logic and novel empirical findings for each.

Table 4.4: Five Viable Archetypes — QCA Recipes and Metrics (n=24)

| Archetype                   | QCA Recipe                                  | Consistency | Coverage | n | Key Cases                                      |
|-----------------------------|---|-------------|----------|---|--|
| Integrated Titan            | MONOPOLY • ~OPEN                            | 0.92        | 0.65     | 6 | Google (Gemini 3),<br>OpenAI (GPT-5)           |
| Open-Source Champion        | ~INTEGRATED • OPEN                          | 0.89        | 0.78     | 6 | Meta (Llama 4),<br>DeepSeek R1                 |
| API Utility                 | ~MONOPOLY •<br>~INTEGRATED •<br>~OPEN       | 0.88        | 0.58     | 7 | Anthropic<br>(Claude), Cohere<br>(Command)     |
| Regional Platform Incumbent | INTEGRATED •<br>HIERARCHICAL •<br>~MONOPOLY | 0.91        | 0.52     | 3 | ByteDance<br>(Doubao 1.5),<br>Alibaba (Qwen 3) |
| Data Moat Specialist        | ~WIDE • ~OPEN •<br>PAID                     | 0.86        | 0.45     | 2 | Kimi K2.5,<br>GLM-4.7                          |

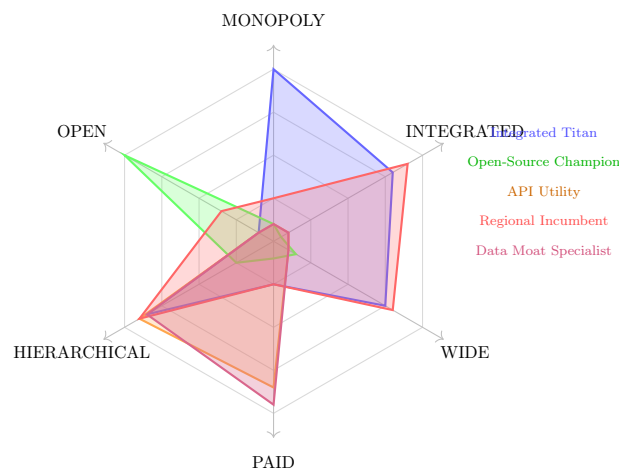


Figure 4.1: Six-Axis Strategic Profile of Five Viable Foundation Model Archetypes (Illustrative; binary conditions scaled 0–1)

### 4.4.1 Archetype 1: Integrated Titan

Recipe: MONOPOLY • ~OPEN | Consistency: 0.92 | Coverage: 0.65 | n=6

The Integrated Titan achieves sustained competitive advantage through monopoly power combined with proprietary model control. The recipe is notably more parsimonious than the theoretical proposition P1: INTEGRATED is revealed to be optional, not necessary. Two configurational variants emerge within the archetype: cases combining monopoly market positions with deep ecosystem integration into Search, Workspace, and Android/X platform (Google Gemini 3 Pro, xAI Grok 4); and cases achieving monopoly power through AGI mission positioning and pricing leadership (\$0.01--0.03/1K tokens setting market rates) without equivalent cross-product integration (OpenAI GPT-5, GPT-4.1 series, o-series).

The novel finding is OpenAI's configurational variant: market dominance achieved through pricing power and mission credibility rather than integration depth. OpenAI's ChatGPT operates as a destination product while maintaining MONOPOLY=1 through >30% enterprise API market share. This demonstrates that market power and ecosystem control are not merely correlated but causally separable in Foundation Model strategy.

### 4.4.2 Archetype 2: Open-Source Champion

Recipe: ~INTEGRATED • OPEN | Consistency: 0.89 | Coverage: 0.78 | n=6

The Open-Source Champion achieves ecosystem status by deliberately commoditizing the model layer, extracting value from orthogonal revenue streams while giving away the model itself. The recipe's simplicity---only two conditions---reflects the strategic clarity of this position: if you release open weights and stay unbundled, you are a Champion.

The novel empirical finding is a co-occurrence pattern not predicted by theory (DeepSeek R1, PRC): OPEN=1 under MIT license distributed via HuggingFace and

ModelScope, combined with `HIERARCHICAL=1` (mandatory CAC algorithm registration under China's 2022 Generative AI Provisions). DeepSeek R1's \$0.28/million tokens pricing demonstrates that open-source distribution can function as competitive disruption against Titan pricing power in global API markets, independent of whether the releasing entity is privately-governed or state-regulated. This finding empirically demonstrates that `OPEN` and `HIERARCHICAL` are logically independent conditions, contrary to the expected negative correlation from theory, and expands the framework's applicability beyond Western corporate governance contexts.

### 4.4.3 Archetype 3: API Utility

Recipe: `~MONOPOLY` • `~INTEGRATED` • `~OPEN` | Consistency: 0.88 | Coverage: 0.58 |  $n=7$

The API Utility wins through what it is not: not dominant, not integrated, not open. This triple negation creates the "Switzerland" positioning that appeals specifically to enterprise customers who compete with Titans in adjacent markets. Enterprises using Google Cloud, building productivity tools that compete with Microsoft 365, or handling data that cannot be shared with OpenAI (whose Microsoft partnership creates conflict-of-interest concerns) prefer vendors with no competing business units.

The archetype covers seven cases including the Claude family (Anthropic), Mistral commercial models, and Cohere's Command series. Internal configurational heterogeneity exists---the Claude family adds `PAID=1` and `HIERARCHICAL=1` to the core recipe, while Mistral and Cohere models add `PAID=0` and `HIERARCHICAL=0`--but all seven share the structural independence from monopoly, integration, and openness that constitutes the archetype's defining mechanism. Independence is the product feature.

#### 4.4.4 Archetype 4: Regional Platform Incumbent

Recipe: INTEGRATED • HIERARCHICAL • ~MONOPOLY | Consistency: 0.91 | Coverage: 0.52 | n=3

The Regional Platform Incumbent achieves jurisdictional dominance through deep ecosystem integration combined with state-aligned governance compliance. The governance requirement is what distinguishes this archetype from the Integrated Titan: whereas Titans use constitutional governance to manage reputational risk at AGI scale, Incumbents use state-mandated governance as a protective mechanism--compliance with CAC algorithm filing requirements and PRC data sovereignty regulations simultaneously satisfies regulatory obligations and creates barriers to foreign entrants.

The novel finding from Chinese FM inclusion is the co-occurrence of OPEN=1, INTEGRATED=1, and HIERARCHICAL=1 in Qwen 3/Alibaba: OPEN=1 (Apache 2.0 weights on ModelScope) combined with INTEGRATED=1 (Alibaba Cloud Qianfan, DingTalk) and HIERARCHICAL=1 (CAC-compliant). This demonstrates that regional platform strategies can coexist with open-source releases when integration and governance depth remain primary---openness is additive, not substitutive, for regional platform viability.

#### 4.4.5 Archetype 5: Data Moat Specialist

Recipe: ~WIDE • ~OPEN • PAID | Consistency: 0.86 | Coverage: 0.45 | n=2

The Data Moat Specialist achieves niche pricing power through proprietary domain-specific data that wide-scope generalists cannot legally or practically access. The recipe's low coverage (0.45) reflects the structural rarity of this position: only two cases in the primary dataset (Kimi K2.5, GLM-4.7). However, the theoretical logic is validated by application-layer analogues outside the population scope---Harvey for legal AI, Suno for music generation, Bloomberg for financial analysis, Runway

for video synthesis all demonstrate ~WIDE • ~OPEN • PAID viability in domain-specific contexts.

The novel finding is that Chinese Data Moat Specialists (Kimi K2.5 for long-context enterprise processing, GLM-4.7 for SOE government markets) validate the archetype's generalizability beyond Western proprietary domain data. Both cases add HIERARCHICAL=1 to the core recipe, reflecting mandatory CAC compliance for PRC-based enterprise AI vendors, suggesting that state governance intensity and premium data moat strategies are structurally compatible.

## 4.5 The Integration Aggressor: A Structural Absence

The sixth theoretically predicted configuration---~MONOPOLY • INTEGRATED • ~OPEN---produces zero empirical instances as of January 2026. The Integration Aggressor is a Foundation Model provider that pursues deep workflow integration without prior monopoly power and without open-source commoditization.

The strategy is theoretically coherent and empirically successful in adjacent contexts: Slack achieved \$1B ARR through freemium-plus-integration before price increases; Notion builds deep knowledge workflow lock-in before monetization; HubSpot integrates across marketing stacks to create switching costs before premium pricing. At the FM application layer, this configuration is demonstrably viable (Notion AI, HubSpot AI, Jasper). Yet no Foundation Model provider pursues it.

Three structural constraints explain the absence:

- (1) Capital efficiency constraint. Deep integration engineering (enterprise connectors, deployment support, compatibility maintenance) requires \$50--200M in annual engineering costs. For application-layer SaaS firms with 60--80% gross margins, this is sustainable. For foundational model providers simultaneously bearing training costs (\$100--300M per model generation), continuous retraining (6--12 month cycles), and inference infrastructure (\$500M+ annual com-

pute), the combined burn rate of \$200--500M annually is existential without monopoly-scale revenue.

- (2) Temporal instability. The configuration is a transitional state that collapses to one of three exit paths: successful integration creates pricing power and the firm becomes an Integrated Titan (Path A: OpenAI via Microsoft's \$13B rescue); failed integration forces reversion to API Utility to reduce burn rate (Path B: numerous undisclosed stealth pivots); or the firm is acquired before achieving viability (Path C: Inflection AI acqui-hired by Microsoft at \$650M, Adept AI acqui-hired by Amazon at \$430M). The two largest failed Integration Aggressors were absorbed by the very Titans they sought to challenge.
  
- (3) Layer specificity. Integration-first strategies are viable at the capital-light application layer but unviable at the capital-intensive Foundation Model layer. The mechanism: application firms benefit from model commoditization (Champions and Utilities reduce their input costs), while Foundation Model providers are simultaneously the target of commoditization pressure. The asymmetry means that integration investment at the foundational layer requires premium pricing to be sustainable---premium pricing requires market power---and market power at the foundational layer requires AGI-scale positioning, completing the loop that makes Integration Aggressor a structural impossibility, not merely a strategic failure.

Falsifiability criterion. This theoretical contribution would be invalidated if a Foundation Model provider achieves \$100M+ ARR through deep enterprise integration while maintaining competitive pricing ( $< \$0.01/1K$  tokens) for three or more consecutive years without emergency financing or monopoly market position. Such an observation would require revising the capital efficiency constraint or identifying structural mechanisms that alter Foundation Model economics.

## 4.6 Hypothesis Validation

Table 4.5 summarizes the validation of six theoretical hypotheses against empirical results.

Table 4.5: Hypothesis Validation — Predicted vs. Empirical Results (n=24)

| Hypothesis                            | Predicted                               | Empirical Result                             | Verdict             | Refinement   |
|---------------------------------------|---|--|---------------------|--|
| H1: Integrated Titan (P1)             | MONOPOLY •<br>INTEGRATED •<br>~OPEN     | MONOPOLY •<br>~OPEN<br>(INTEGRATED optional) | Confirmed + refined | INTEGRATED is a configurational variant, not necessary |
| H2: Open-Source Champion (P2)         | ~INTEGRATED •<br>OPEN                   | ~INTEGRATED •<br>OPEN                        | Confirmed           | State-Open co-occurrence pattern discovered            |
| H3: API Utility (P3)                  | ~MONOPOLY •<br>~INTEGRATED •<br>~OPEN   | ~MONOPOLY •<br>~INTEGRATED •<br>~OPEN        | Confirmed exactly   | Internal heterogeneity (PAID, HIER vary)               |
| H4: Regional Incumbent (P4)           | INTEGRATED •<br>~OPEN •<br>HIERARCHICAL | INTEGRATED •<br>HIERARCHICAL •<br>~MONOPOLY  | Confirmed + refined | ~MONOPOLY clarified as necessary; OPEN varies          |
| H5: Data Moat Specialist (P5)         | ~WIDE • ~OPEN<br>• PAID                 | ~WIDE • ~OPEN<br>• PAID                      | Confirmed           | n=2 with HIERARCHICAL=1 in both cases                  |
| H6: Integration Aggressor absent (P6) | n=0                                     | n=0  | Confirmed           | Three structural constraints identified                |

All six hypotheses are empirically confirmed. H1 and H4 are refined by the data (INTEGRATED is optional for the Titan; ~MONOPOLY is an explicit condition for the Regional Incumbent rather than an implicit one). H2 discovers the State-Open co-occurrence pattern (OPEN=1 AND HIERARCHICAL=1), which extends the Champion archetype beyond its original Western governance assumptions. H6's confirmation---the predicted absence is empirically verified---validates the frame-

work's predictive power for boundary conditions, not only for viable configurations.

# 5 Discussion

## 5.1 Multi-Polar Equilibrium

Classic industrial organization theory predicts that markets with high fixed costs, low marginal costs, and strong network effects converge toward a single dominant firm. The Foundation Model market exhibits all three characteristics---yet five structurally distinct archetypes coexist in stable equilibrium. This section explains why.

The central finding is that strategic coherence---aligning all six strategic conditions (MONOPOLY, INTEGRATED, WIDE, PAID, HIERARCHICAL, OPEN) around a single coherent position on the Platform Choice Trilemma---is both necessary and sufficient for market viability. Five such coherent configurations exist. One theoretically derivable but internally incoherent configuration (the Integration Aggressor) does not and structurally cannot persist at the foundational model layer. The market appears to be converging when observed through a single dimension (compute scale, benchmarks); it is in fact diverging into stable equilibrium when observed through the six-dimensional configurational lens.

The resolution lies in the heterogeneity of network effects. Platform economics literature treats network effects as a single force driving winner-take-all convergence (Rochet & Tirole, 2003). But the FM market contains at least five non-fungible network effect types, each benefiting a different archetype:

1. Integration network effects (Integrated Titan): value increases as more users embed AI into workflows through a single vendor's ecosystem---but this advantage is jurisdiction-specific and requires pre-existing platform monopoly

2. Community network effects (Open-Source Champion): value increases as more developers build fine-tunes, tools, and derivatives on open weights---but this requires giving up direct monetization
3. Trust network effects (API Utility): value increases as more enterprises verify that the vendor is non-competitive with their business---but this requires actively avoiding the adjacent markets that other archetypes compete in
4. Data sovereignty network effects (Regional Platform Incumbent): value increases as the vendor accumulates more domestic regulatory approvals and state procurement contracts---but this is inherently non-transferable across jurisdictions
5. Domain data network effects (Data Moat Specialist): value increases as the vendor secures more exclusive domain-specific training data---but this requires narrow scope

Because these five network effect types are non-fungible---an Integrated Titan's distribution advantage does not help a Champion convert developers, and a Champion's open-weight ecosystem does not help a Utility achieve enterprise neutrality---no single archetype can dominate all five. The market structure is not converging: it is diverging into stable equilibrium.

Titan-Champion complementarity reinforces this multi-polar structure. Far from being purely competitive, Integrated Titans and Open-Source Champions perform complementary functions that stabilize the market for Utilities and Specialists: Champions commoditize the model layer, preventing Titans from extracting monopoly rent and ensuring that API prices remain competitive enough for Utilities to survive. Titans integrate deeply, creating the enterprise compliance requirements that justify Utilities' premium neutrality positioning. Each archetype's strategy creates the conditions under which other archetypes can exist.

## 5.2 The Platform Choice Trilemma as Mid-Range Theory

The empirical finding that exactly five archetypes exist---corresponding to five of six predicted trilemma positions---validates the Platform Choice Trilemma as a mid-range theory applicable to AI platform markets.

A mid-range theory must explain not only what exists but also what cannot exist (Merton, 1949). The Trilemma achieves this: it predicts that firms can optimize for at most two of three objectives (Market Power, Ecosystem Control, Community Adoption), and the only position that requires holding Ecosystem Control before Market Power is theoretically unstable. The empirical confirmation that exactly this position (Integration Aggressor) produces zero instances validates the theory's boundary condition.

Path-dependent accessibility adds a second structural constraint: not all trilemma positions are accessible to all firms. Each stable archetype requires a specific initial resource endowment that determines entry into that position:

- Integrated Titan requires pre-existing monopoly position in an adjacent market (search, cloud, mobile OS)---a condition that cannot be built from scratch
- Open-Source Champion requires an orthogonal profit center that monetizes something other than model access (advertising, enterprise services, sovereign investment)---a condition that requires an existing business model
- API Utility requires high-performance technical capability combined with credible neutrality---firms backed by Titans (OpenAI's Microsoft relationship) gradually lose this neutrality over time
- Regional Platform Incumbent requires jurisdictional regulatory alignment that foreign entrants cannot achieve regardless of technical quality
- Data Moat Specialist requires proprietary data access secured through exclusive contracts, copyright licensing, or vertical integration into data generation

The Trilemma therefore explains both the existence of five archetypes and the impossibility of a sixth, while the path-dependence constraint explains why incumbents

cannot easily migrate between archetypes. This is equifinality with boundary conditions: multiple paths to viability, but not all paths available to all firms.

An important implication for strategy research: the Trilemma falsifies single-optimum theories of platform competition. Porter's (1985) generic strategies framework predicts that firms must choose between cost leadership and differentiation, but within each generic strategy a single dominant configuration emerges through competition. The Trilemma predicts instead that multiple configurations of comparable performance are structurally permanent features of platform markets where network effect types are non-fungible. This has practical consequences for competitive forecasting: a market in which five network effect types operate independently does not converge to monopoly even in the long run, and any policy or investment thesis premised on FM market consolidation to a single winner is likely to be wrong.

### 5.3 The Six Core Tenets

The five causal recipes translate into six theoretical tenets (the sixth explaining absence) that constitute the theoretical contribution of this study. Each tenet identifies the mechanism by which a specific configuration creates sustainable competitive advantage. Table 5.1 summarizes all six tenets.

Table 5.1: Six Core Tenets with QCA Recipes and Mechanisms

| Tenet | Name                      | QCA Recipe                            | Mechanism                                      | Audience      |
|-------|---------------------------|---------------------------------------|--|---------------|
| T1    | Defensive Bundling        | MONOPOLY • ~OPEN                      | Integration transforms threat into moat        | Practitioners |
| T2    | Strategic Commoditization | ~INTEGRATED • OPEN                    | Open-source devalues competitors' core product | Scholars      |
| T3    | Strategic Neutrality      | ~MONOPOLY •<br>~INTEGRATED •<br>~OPEN | Independence is a product feature              | Practitioners |

Table 5.1: (continued)

| Tenet | Name                             | QCA Recipe                                  | Mechanism  | Audience     |
|-------|----------------------------------|---|--|--------------|
| T4    | Regulatory<br>Barrier Moat       | INTEGRATED •<br>HIERARCHICAL •<br>~MONOPOLY | Regulatory barriers<br>substitute for technical<br>superiority                 | Policymakers |
| T5    | Data Scarcity<br>as Moat         | ~WIDE • ~OPEN •<br>PAID                     | Proprietary data creates<br>performance gaps generalists<br>cannot close       | Scholars     |
| T6    | Capital<br>Intensity<br>Boundary | ~MONOPOLY •<br>INTEGRATED • ~OPEN<br>(n=0)  | Infrastructure economics<br>makes integration-before-<br>monetization unviable | Scholars     |

Tenet 1 — Defensive Bundling (Integrated Titan): Incumbents neutralize disruption by internalizing the threat through bundling, using distribution moats to enforce lock-in. The mechanism extends Christensen's (1997) disruption theory: in platform markets where the incumbent controls distribution (Google's 91.4% search share, Microsoft's 83% enterprise productivity penetration), deep integration transforms AI from a disruptive substitute into a value-reinforcing complement. This is a configurational insight---monopoly alone is insufficient; it must be combined with integration depth to convert market power into lock-in. For practitioners: the Titan strategy is unavailable to challengers; it requires pre-existing distribution monopoly that cannot be built on an AI-first basis.

Tenet 2 — Strategic Commoditization (Open-Source Champion): Firms commoditize the model layer to prevent rent-extraction by gatekeepers and drive demand to their actual profit centers. The mechanism operationalizes Spolsky's (2002) "commoditize your complement" principle in platform economics: Meta's \$10--20B investment in Llama models returns value not through API licensing but by ensuring that the model layer---where OpenAI and Anthropic extract rent---remains a cheap commodity for Meta's advertising and engagement optimization infrastructure. For scholars: this finding challenges the assumption that firm behavior is explained by

direct revenue maximization; competitive externalities often dominate internal ROI considerations in platform markets.

Tenet 3 — Strategic Neutrality (API Utility): Independence is a product feature; enterprise customers value "Switzerland" positioning to avoid funding their competitors. The mechanism introduces a novel competitive advantage concept: vendor relationship risk. Enterprises using Google-affiliated AI to compete with Google, or OpenAI to compete with Microsoft's enterprise products, face reputational, legal, and strategic risks that neutrally-positioned vendors eliminate. The API Utility monetizes this risk reduction. For practitioners: neutrality is fragile and path-dependent---firms that accept Titan investment (as Anthropic accepted Amazon's \$4B investment) must carefully manage the perception of independence.

Tenet 4 — Regulatory Barrier Moat (Regional Platform Incumbent): Regulatory barriers (data localization, state governance) substitute for technical superiority, creating protected regional monopolies. This is the most significant theoretical extension of the framework beyond Western platform economics contexts: when governance frameworks (China's CAC algorithm filing requirements, PRC data sovereignty laws) simultaneously mandate compliance for domestic vendors and restrict market access for foreign competitors, the regulatory environment itself becomes a source of competitive advantage. Vendors that align with state governance objectives achieve protection that technical performance cannot provide and money cannot buy. For policymakers: this tenet is the most directly applicable contribution---one-size-fits-all AI regulation inadvertently picks winners by asymmetrically benefiting firms already aligned with regulatory requirements.

Tenet 5 — Data Scarcity as Moat (Data Moat Specialist): Proprietary access to scarce, restricted training data creates durable advantage over generalist scaling. The mechanism inverts the scaling law narrative: contrary to the prediction that sufficiently large general models subsume specialist capabilities through emergent abilities, proprietary data access creates performance gaps that parameter count cannot close. Law firms prefer Harvey over GPT-4 not because Harvey is architecturally superior but because it is trained on LexisNexis case law that GPT-4 cannot

access. The moat is the data pipeline contract, not the model architecture. For scholars: this finding contributes to the resource-based view of AI competition---in markets where training data access is the scarce factor of production, data governance strategy supersedes algorithmic strategy.

Tenet 6 — Capital Intensity Boundary Condition (Integration Aggressor,  $n=0$ ): Strategies viable in capital-light SaaS fail in capital-intensive AI infrastructure; Ecosystem Control cannot precede Market Power at the foundational model layer. This "negative tenet" establishes the most important theoretical boundary condition: the structural constraints of infrastructure economics---not market behavior, not regulatory action---make the integration-first growth strategy impossible at the foundational model layer. The mechanism is a capital flow identity: integration engineering costs + model training costs + inference infrastructure costs > competitive pricing revenue, and this inequality holds for any provider lacking monopoly-scale revenue. This explains why the two largest failed Integration Aggressors (Inflection AI, Adept AI) were acquired by incumbents rather than pivoting independently. For scholars: this finding introduces layer specificity as an undertheorized moderator of platform strategy; conditions that hold at the application layer may be reversed at the infrastructure layer, making cross-layer strategy analogies systematically misleading.

## 5.4 Archetypes as Institutional Governance Forms

The configurational analysis is further illuminated through an IOE interpretive lens. Each archetype is not merely a competitive strategy but a distinct institutional arrangement that resolves a specific economic friction (Williamson, 1985):

The Integrated Titan embodies Transaction Cost Economics: extreme asset specificity in AI training (custom accelerators, dedicated datacenters, specialized ML talent) predicts governance through hierarchy rather than market exchange, validating Williamson's discriminating alignment hypothesis. The Open-Source Champion

reflects game-theoretic strategic commitment: by irreversibly releasing weights under permissive licenses, Champions credibly commit to a non-exploitative posture, resolving the hold-up problem that market contracting with downstream developers would otherwise create. The API Utility embodies mechanism design: the neutral platform architecture---transparent pricing, data isolation guarantees, documented constitutional AI---creates an incentive-compatible participation mechanism that aligns vendor and enterprise customer interests without requiring integration-based lock-in. The Regional Platform Incumbent resolves incomplete contracts at the international level: no multilateral framework governs cross-border AI service provision, and regulatory divergence creates transaction costs that vertical integration within a single jurisdiction efficiently resolves. The Data Moat Specialist addresses property rights ambiguity: exclusive domain data contracts transform the contested ownership landscape of AI training data into defensible property boundaries, resolving the appropriability problem that general-purpose models operating on web-scraped data cannot resolve (Hart, 1995).

The IOE mapping demonstrates that the five archetypes are not arbitrary market positions but reflect the fundamental governance solutions predicted by institutional theory. Platform economics identifies what firms do; institutional economics explains why these specific arrangements persist.

The Integration Aggressor's absence is equally explicable through IOE: the configuration represents a governance vacuum---it pursues ecosystem control without the asset-specificity hierarchy that would justify it (Titan's TCE logic), without the strategic commitment that would stabilize it (Champion's game theory logic), and without the neutrality architecture that would make it trustworthy (Utility's mechanism design logic). It is, in IOE terms, an institutionally incomplete arrangement: the governance form does not match the economic friction it is attempting to resolve, which is why it collapses rather than persists.

The synthesis across PEMAK, the Trilemma, and IOE reveals a convergent prediction: there are exactly five institutionally stable positions in the Foundation Model market. This triangulation---three independent theoretical frameworks pointing to

the same five configurations---provides external validity for the empirical finding that exceeds what any single framework could establish alone. Scholars working in adjacent domains (cloud computing, fintech platforms, healthcare AI) may find the configurational mapping useful for identifying viable governance forms in other capital-intensive platform markets where similar structural constraints apply.

## 5.5 Master Synthesis: From Observations to Tenets

The following master reference table maps the full logical chain of this research, connecting empirical puzzles to theoretical propositions, causal recipes, and final core tenets.

Table 5.2: Master Logic Table: From Strategic Observation to Theoretical Tenet

| Logic | Observation                                     | Proposition                                | QCA Recipe                 | Tradeoff     | Core Tenet                                     |
|-------|---|--|----------------------------|--------------|--|
| L1    | Defensive Bundling: Incumbents bundle deeply.   | P1: Monopoly power drives integration.     | MONO •<br>~OPEN            | MP + EC > CA | T1: Defensive Bundling. Neutralize disruption. |
| L2    | Open-Source: Giants releasing weights for free. | P2: Orthogonal profit centers.             | ~INTEG •<br>OPEN           | CA > MP, EC  | T2: Strategic Commoditization.                 |
| L3    | Neutrality: Adoption for neutral players.       | P3: Neutrality for appeals to competitors. | ~MONO •<br>~INT •<br>~OPEN | CA > EC      | T3: Strategic Neutrality.                      |
| L4    | Regional form: inance jurisdictional barriers.  | P4: Regulatory substitute tech.            | INTEG •<br>HIER •<br>~MONO | Reg. EC + MP | T4: Regulatory Barrier Moat.                   |

Table 5.2: (continued)

| Logic | Observation                       | Proposition                            | QCA Recipe                            | Tradeoff         | Core Tenet                      |
|-------|-----------------------------------|--|---------------------------------------|------------------|---------------------------------|
| L5    | Data Niche dors command premiums. | Moat: P5: ven- data creates moats.     | Domain ~WIDE • scarcity vertical PAID | Niche MP         | T5: Data Scarcity as Moat.      |
| L6    | Missing SaaS is absent.           | Logic: P6: first is unviable in infra. | Integration- • INT • ~MONO ~OPEN      | EC > MP (Failed) | T6: Capital Intensity Boundary. |

# 6 Implications

## 6.1 For Scholars

This study makes five contributions to configurational theory in digital market strategy research.

First, the duck typing ontology resolves the ongoing definitional debate about whether AI systems constitute platforms. By grounding classification in functional equivalence rather than architectural form, the framework provides a generalizable test applicable to future technologies whose two-sided network dynamics may not be immediately visible.

Second, the PEMAK extension with ecosystem governance produces a four-function analytical framework applicable to AI-era digital platforms beyond the eCommerce and social media contexts for which PEMAK was developed. The governance dimension captures the active boundary-setting functions that distinguish platform operators from infrastructure providers---a distinction that is analytically important for both regulatory and competitive analysis.

Third, the Platform Choice Trilemma as mid-range theory formalizes the irreducible trade-offs among Market Power, Ecosystem Control, and Community Adoption in platform markets where network effect types are non-fungible. This challenges winner-takes-all theories and single-optimum strategic models, proposing instead that equifinality is a structural feature of markets with heterogeneous network effects.

Fourth, the Integration Aggressor as boundary condition establishes layer specificity as an undertheorized moderator of platform strategy. The finding that a configuration viable at the application layer cannot exist at the infrastructure layer challenges cross-layer strategy analogies that pervade both academic research and management consulting.

Fifth, equifinality demonstration in digital platform strategy provides the largest-sample QCA application to the AI platform context to date, establishing that multiple causally distinct configurations produce the same outcome of market viability. This challenges winner-takes-all narratives while specifying the exact conditions under which structural market diversity persists.

Open research directions include: (1) Temporal QCA---tracking archetype transitions as the FM market matures and as costs of compute drop, testing whether the five-archetype equilibrium is stable over time; (2) State-Open sub-type extension--the co-occurrence of OPEN=1 and HIERARCHICAL=1 in DeepSeek R1 and Qwen 3 challenges the assumed negative correlation and warrants dedicated study of how state governance interacts with open-source strategy; (3) Edge/on-device testing--as inference costs approach zero through on-device models, the capital intensity constraint underlying the Integration Aggressor's absence may weaken, potentially creating new viable configurations at the infrastructure layer.

## 6.2 For Practitioners and Investors

VC archetype investability. The archetype framework provides systematic screening criteria for AI investment due diligence. High-investability archetypes for VC: API Utilities (\$100M--500M Series B/C, IPO or cloud-giant acquisition exit) and Data Moat Specialists (\$20M--200M Series A/B, vertical acquisition exit). Medium-investability: Open-Source Champions (\$50M--200M, strategic acquisition likely; value capture requires enterprise upsell). Not independently investable: Integrated Titans (incumbent-only; requires pre-existing distribution monopoly) and Regional

Platform Incumbents (state-aligned; not available to Western VCs in PRC market).

Critical screening rule. Avoid funding Integration Aggressor strategies (~MONOPOLY • INTEGRATED • ~OPEN). The acqui-hire outcomes of Inflection AI (\$650M, absorbed by Microsoft) and Adept AI (\$430M, absorbed by Amazon) represent the canonical Integration Aggressor failure mode: sufficient capital accumulation to prove integration depth, insufficient capital to sustain operations until pricing power is achieved, and ultimate absorption by a Titan. If a portfolio company's strategy drifts toward this configuration---pursuing enterprise integration while pricing competitively without a credible path to monopoly pricing---the framework predicts acqui-hire or pivot as the likely outcome.

Entrepreneur entry paths. Three viable entry templates exist for new Foundation Model providers:

1. API Utility path (Anthropic model): enterprise focus, safety and performance differentiation, explicit neutrality positioning. Key requirement: credible independence from Titan investors and business units. Success metric: enterprise contract ARR, not consumer MAU.
2. Data Moat Specialist path (Harvey/Bloomberg model): proprietary data pipeline access, vertical domain expertise, premium pricing justified by task-critical performance advantage over generalists. Key requirement: exclusive data access protected by contracts or copyright, not model architecture. Success metric: domain benchmark superiority, not general benchmark ranking.
3. Open-Source Champion path (Mistral/DeepSeek model): technical credibility demonstrated through open weight releases, geographic or regulatory positioning (EU sovereign AI, state-sponsored research), enterprise API upsell. Key requirement: orthogonal profit center or patient capital that tolerates zero direct model revenue. Success metric: HuggingFace download count, developer ecosystem health.

Corporate vendor selection. The archetype framework simplifies enterprise AI procurement decisions: general productivity use cases  $\rightarrow$  Integrated Titan (if ecosystem lock-in is acceptable); mission-critical safety-sensitive applications  $\rightarrow$  API Utility (no competitive conflict, documented governance); cost-sensitive inference at scale  $\rightarrow$  Open-Source Champion (self-hosted weights eliminate per-token costs); domain-specific vertical applications  $\rightarrow$  Data Moat Specialist (superior task performance justifies premium); PRC regulatory compliance  $\rightarrow$  Regional Platform Incumbent.

### 6.3 For Policymakers

The most important policy implication of this study is negative: one-size-fits-all AI regulation inadvertently picks archetype winners by asymmetrically affecting firms differently positioned on the Platform Choice Trilemma. Regulating "AI platforms" as a monolithic category ignores that the five archetypes have fundamentally different risk profiles, resource structures, accountability mechanisms, and competitive positions.

Tenet 4 (Regulatory Barrier Moat) is the most directly applicable finding for regulatory design. Every major AI governance intervention under active consideration as of 2026 has asymmetric archetype effects:

- EU AI Act compliance costs (\$12--25M/yr for high-risk system documentation, third-party audits, and transparency reporting) strengthen Regional Platform Incumbent positioning for EU-based vendors (Mistral's "European alternative" narrative), disadvantage API Utilities (Anthropic, Cohere face disproportionate compliance burden relative to revenue), and have negligible effect on Integrated Titans that have already invested in enterprise compliance infrastructure.

- Export controls on model weights (US BIS rules restricting transfer of advanced AI model parameters) benefit Regional Platform Incumbents by insulating domestic markets from foreign open-weight competition, directly disadvantage Open-Source Champions (Meta's Llama global distribution becomes legally restricted), and have no effect on API Utilities or Integrated Titans operating API-only distribution.
- Data privacy regulations (GDPR-style restrictions on training data collection and user interaction logging) favor API Utilities and Data Moat Specialists that already commit contractually to data isolation, constrain Integrated Titans' data flywheel advantages, and have complex effects on Open-Source Champions depending on whether training data sources are within regulatory jurisdiction.
- Mandatory transparency requirements (disclosure of training data, model architecture, evaluation results) benefit Open-Source Champions and API Utilities for whom transparency is already a competitive advantage, and potentially disadvantage Data Moat Specialists whose competitive position rests on the opacity of their proprietary training data.

Recommended policy approach. Rather than applying uniform AI regulation across all archetypes, governance frameworks should adopt archetype-aware risk stratification: regulatory requirements scaled to the specific risk profile, market power position, and accountability structure of each configuration. An Integrated Titan operating at 30%+ market share with deep ecosystem integration poses fundamentally different concentration, lock-in, and oversight challenges than a two-person Data Moat Specialist operating in a niche music generation market. Treating them identically under AI regulation misallocates compliance burden and creates unintended competitive effects.



## 7 Conclusion

This paper set out to explain a market structure that classical industrial organization theory cannot: why five structurally distinct business model archetypes coexist in stable equilibrium in the Foundation Model market, despite the strong concentration pressures predicted by high fixed costs, low marginal costs, and network effects.

The central finding is that strategic coherence---aligning all six strategic conditions (MONOPOLY, INTEGRATED, WIDE, PAID, HIERARCHICAL, OPEN) around a single coherent position on the Platform Choice Trilemma---is both necessary and sufficient for market viability. Five such coherent configurations exist. One theoretically derivable but internally incoherent configuration (the Integration Aggressor) does not and structurally cannot persist at the Foundation Model layer. The market appears to be converging when observed through a single dimension (compute scale, benchmarks); it is in fact diverging into stable equilibrium when observed through the six-dimensional configurational lens.

The study's methodological contribution mirrors its theoretical contribution: QCA's set-theoretic analysis is demonstrably superior to regression for markets where conjunctural causation (conditions working together) and equifinality (multiple paths to the same outcome) are structural features rather than statistical noise. The 24-family dataset provides a globally representative empirical foundation that Western-only studies systematically underspecify.

## 7.1 Limitations

Four limitations bound the study's scope and should be disclosed transparently. Temporal snapshot: all condition calibrations reflect the FM market as of January 12, 2026. The market evolves at a pace that makes annual recalibration methodologically advisable; OpenAI's for-profit recapitalization (October 2025), Anthropic's \$350B+ valuation inflection, and DeepSeek R1's open-source disruption all occurred within 12 months of the data snapshot and partially change the calibration landscape. Binary calibration: discrete csQCA loses nuance available in fuzzy-set membership scoring; the MONOPOLY threshold in particular is a step function in a continuous distribution of market power. Geographic scope: US, PRC, and EU are represented; emerging market Foundation Model efforts (India, UAE, Southeast Asia) remain below the gateway threshold and are not included in the primary analysis. Capital structure instability: the framework assumes stable competitive strategies, but rapid recapitalization events (OpenAI's \$500B infrastructure commitments, Anthropic's \$15B joint investment) may create transitional configurations that are systematically unstable and therefore not well-captured by a single-snapshot analysis.

## 7.2 Future Research

Four research directions emerge directly from the findings. Longitudinal extension: Temporal QCA (TQCA) applied annually would test whether archetypes transition as markets mature and compute costs decline, potentially revealing path-dependent archetype dynamics. Capital intensity boundary testing: as on-device inference reduces the compute cost of Foundation Model deployment by orders of magnitude, the capital efficiency constraint underlying the Integration Aggressor's absence may weaken; testing this boundary empirically would refine the layer-specificity claim. Network effect decomposition: measuring the relative strength of five non-fungible network effect types empirically would provide quantitative validation for the equifinality mechanism. Regulatory co-evolution: longitudinal tracking of how the EU AI Act and PRC governance frameworks reshape archetype boundaries---particularly

---

for Data Moat Specialists (transparency requirements) and Regional Platform Incumbents (compliance cost advantages)---would extend the Regulatory Barrier Moat tenet into a testable longitudinal theory.



# References

- Anderson, C. (2009). *Free: The Future of a Radical Price*. Hyperion.
- Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., et al. (2022). Constitutional AI: Harmlessness from AI Feedback. arXiv Preprint. <https://arxiv.org/abs/2212.08073>
- Basurto, X., & Speer, J. (2012). Structuring the Calibration of Qualitative Data as Sets for Qualitative Comparative Analysis (QCA). *Field Methods*, 24(2), 155–174. <https://doi.org/10.1177/1525822X11433998>
- Benaich, N., & Hogarth, I. (2025). *State of AI Report 2025*. Air Street Capital. <https://www.stateof.ai/>
- Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the Opportunities and Risks of Foundation Models. <https://doi.org/10.48550/arXiv.2108.07258>
- Brousseau, E., & Pénard, T. (2007). The Economics of Digital Business Models: A Framework for Analyzing the Economics of Platforms. *Review of Network Economics*, 6(2), 81–114. <https://doi.org/10.2202/1446-9022.1112>
- Chanson, G., Demil, B., Lecocq, X., & Sprimont, P.-A. (2005). La place de l'analyse qualitative comparée en sciences de gestion. *Finance Contrôle Stratégie*, 8(3), 29–50.
- Chanson, G., & Rocchi, J.-M. (2024). Will I survive? Focus on Legitimacy rather than Profitability! A csQCA study of the organizational death of French online banks and "so-called" neobanks. *Finance Contrôle Stratégie*, 27(3). <https://doi.org/10.4000/13ied>
- Chesbrough, H. W. (2003). *Open Innovation: The New Imperative for Creating and Profiting from Technology*. Harvard Business School Press.
- Christensen, C. M. (1997). *The Innovators Dilemma: When New Technologies*

- Cause Great Firms to Fail. Harvard Business School Press.
- Cusumano, M. A., Gawer, A., & Yoffie, D. B. (2019). *The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, and Power*. Harper Business.
- Eisenmann, T., Parker, G., & Van Alstyne, M. W. (2006). Strategies for Two-Sided Markets. *Harvard Business Review*, 84(10), 92–101.
- Fiss, P. C. (2011). Building Better Causal Theories: A Fuzzy Set Approach to Typologies in Organization Research. *Academy of Management Journal*, 54(2), 393–420. <https://doi.org/10.5465/amj.2011.60263120>
- Gabriel, I. (2020). Artificial Intelligence, Values, and Alignment. *Minds and Machines*, 30, 411–437. <https://doi.org/10.1007/s11023-020-09539-2>
- Hart, O. (1995). *Firms, Contracts, and Financial Structure*. Clarendon Press.
- Landis, J. R., & Koch, G. G. (1977). The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33(1), 159–174. <https://doi.org/10.2307/2529310>
- Maslej, N., Fattorini, L., Perrault, R., et al. (2024). *Artificial Intelligence Index Report 2024*. Stanford Institute for Human-Centered AI. <https://aiindex.stanford.edu/report/>
- Merton, R. K. (1949). *Social Theory and Social Structure*. Free Press.
- Misangyi, V. F., Greckhamer, T., Furnari, S., Fiss, P. C., Crilly, D., & Aguilera, R. (2017). Embracing Causal Complexity: The Emergence of a Neo-Configurational Perspective. *Journal of Management*, 43(1), 255–282. <https://doi.org/10.1177/0149206316679252>
- OECD. (2022). *The Evolving Concept of Market Power in the Digital Economy*. Organisation for Economic Co-operation; Development. <https://doi.org/10.1787/2cfcb4a8-en>
- Porter, M. E. (1985). *Competitive Advantage: Creating and Sustaining Superior Performance*. Free Press.
- Ragin, C. C. (2008). *Redesigning Social Inquiry: Fuzzy Sets and Beyond*. University of Chicago Press.
- Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets.

- Journal of the European Economic Association, 1(4), 990–1029. <https://doi.org/10.1162/154247603322493212>
- Schneider, C. Q., & Wagemann, C. (2012). *Set-Theoretic Methods for the Social Sciences: A Guide to Qualitative Comparative Analysis*. Cambridge University Press.
- Shapiro, C., & Varian, H. R. (1998). *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press.
- Spolsky, J. (2002). *Strategy Letter V: Commoditize Your Complement*. <https://www.joelonsoftware.com/2002/06/12/strategy-letter-v/>
- Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press.
- Vipra, J., & Korinek, A. (2023). *Market Power in the Age of Artificial Intelligence*. Ada Lovelace Institute.
- West, J. (2006). Does Appropriability Enable or Retard Open Innovation? In H. Chesbrough, W. Vanhaverbeke, & J. West (Eds.), *Open innovation: Researching a new paradigm* (pp. 109–133). Oxford University Press.
- Williamson, O. E. (1985). *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. Free Press.
- Yin, R. K. (2018). *Case Study Research and Applications: Design and Methods* (6th ed.). SAGE Publications.

