

Micro Connect New Markets (NUMA) White Paper

October 3, 2026 Edition

HAT over MASS

Micro Agents · Simple & Short tasks

Measurement Infrastructure for the Contract Economy:
Theory, Factory, and Samples



Micro Connect New Markets (NUMA) White Paper

HAT over MASS——Micro Agents · Simple & Short tasks

Measurement Infrastructure for the Contract Economy: Theory, Factory, and Samples

Govern the Individual | Unleash the Collective

Executive Summary

HAT over MASS is the third white paper from Micro Connect New Markets, dated October 3, 2026. Each of the three white papers stands on its own and advances the discussion: the August 3 white paper (8-03) sets out the market framework for investment and financing based on contractual cash flows, answering "What kind of market should be built?"; the September 3 white paper (9-03) sets out the measurement rules for the two Compasses, answering "How do we understand and measure contracts and businesses?"; the October 3 white paper (10-03) addresses the next question: Who performs this professional work, and how is it organized, inspected, and continuously improved?

I. Who Is Micro Connect New Markets (NUMA)?

Micro Connect New Markets provides Registration and Measurement Infrastructure for the contract securitization market. This white paper describes six core tools within that infrastructure as "instruments," while the organization that builds and maintains these instruments is described as an "AI factory": HAT over MASS. People set the rules and bear responsibility; AI collaborates through assigned tasks so that the development and maintenance of the tools are traceable and open to inspection. The six instruments work together across the lifecycle of a contract, from admission through its ongoing term. The development status and scope of access for each instrument are governed by its official version and actual disclosures.

1. The Contract Terminal is the first stop for contract admission. The system collects the required materials, enters and verifies each item, and requests additional information or clarification where materials are missing, inconsistent, or fail to meet requirements. It organizes nonstandard commercial contracts into structured, verifiable records that meet admission requirements. It continuously collects and checks information, retaining its source, timestamp, and verification status. The system automatically flags missing or conflicting information; AI inference must not be used to fill gaps in the facts. The contract must be lawful and compliant, in force, and genuinely being performed, with cash flows verifiable by a third party. Once these admission conditions are met, investment suitability must still be independently assessed using the specific evidence.
2. The Contract Registrar organizes collected and verified contracts into standardized records that can be registered, queried, and compared, identifying the contracting parties, agreed terms, and current status and linking to the underlying materials. When terms or status change, the basis for the changes and historical versions are retained. Registration does not change contractual rights or obligations or guarantee contract performance.
3. The Contract Compass measures registered contracts: the Core Compass examines contract structure; the Industry Compass examines the business behind the contract. Each has a Return Dimension and a Risk Dimension, presented separately, without addition or offsetting. Factor definitions, data sources, Rule Versions, and scoring methods are explained with the corresponding official version for professional institutions to review.
4. Compass Reading updates Factor data for registered contracts daily and updates readings according to the nature of the data and new evidence. It continuously records changes in business operations and contract performance, specifies timestamps and Observation Periods, and retains records before and after each update. When there are no new facts, the existing readings and their timestamps are retained; daily checks alone do not change the values.
5. Once investors have completed transactions and settlement through licensed institutions, the Contract Portfolio Registry records daily which contracts each portfolio holds, the quantities held, and changes in those holdings. The Order Register retains original trade and holding facts; the Contract Portfolio Registry links those records at the portfolio level. It records holding relationships and changes and does not provide custody of client assets. Simulated allocations and actual holdings are separately identified and kept in segregated records, so they are not mixed under the same holding measure.

6. The Compass Terminal is a professional interface for queries and analysis by licensed institutions or qualified overseas investors. Users retrieve readings, compare materials, and trace their supporting basis under their own methodologies to form their own judgments. Access to the professional terminal requires the appropriate qualifications and authorization and is unavailable to the public. The public may read general industry research materials; these two forms of access are managed separately.

This instrument system serves both investors and financing parties. On one side are operating entities such as shops, factories, and farms acting as contract issuers. They provide contract materials and operating information organized by Micro Connect New Markets so that professional institutions can understand, measure, and verify them. On the other side are Micro Connect New Markets investors, comprising licensed institutions or qualified overseas investors, which use the tools provided by Micro Connect New Markets to extend professional judgment to more businesses of smaller scale. Investment decisions are made by parties with the relevant decision-making authority. Licensed institutions fulfill the obligations associated with their respective functions within the scope of their licenses.

Readings are not ratings. The compass points north; it does not prescribe a route.

Micro Connect New Markets provides verifiable, continuously updated records and data points. It does not rate the credit quality of contracts or recommend assets to investors. Scores cannot be interpreted as rates of return, default probabilities, or credit ratings. Investors may refer to Compass readings to assess asset quality, but they bear the investment risks themselves.

II. Three White Papers and Six Instruments

Micro Connect New Markets uses six instruments to register, measure, and continuously track contract information. The terminals support interaction and access; between them, registries preserve facts, Compasses provide measurement, and readings track change. The six instruments have distinct functions and can call on one another during business processes; they need not be used in a fixed sequence: the Compass Terminal can support professional consultation throughout, while portfolio registration records actual holding relationships after transactions and settlement are complete. The 8-03 white paper sets out the market framework, and the 9-03 white paper develops the registration and measurement functions. This white paper further integrates those functions into a six-instrument system and uses HAT over MASS to explain how the organization supporting their development and maintenance operates.

The instruments are for people to use; the factory gives their development and maintenance rules, inspection, and accountability.

III. The Problem: How Can Professional Expertise Reach Small Businesses?

Corporate securitization organizes the rights and responsibilities that capital needs to participate in business operations at the company level. When investment and financing focus on individual contracts, due diligence, verification, pricing, and monitoring must still be completed item by item, yet a single small business may be unable to bear the professional cost. A contract specifies cash flow arrangements, but whether those commitments can be fulfilled and payments actually recovered still requires ongoing verification. Previously, professionals handled matters one by one. The finer and more numerous the tasks, the harder it was to make the economics work. AI may lower the cost of repetitive work such as organizing materials and checking rules, giving professional expertise an opportunity to reach smaller businesses.

IV. How the Instruments Are Built: HAT over MASS

HAT over MASS: Govern the Individual, Unleash the Collective

The factory that builds and maintains this instrument system is called HAT over MASS. MASS combines Micro Agents with Simple & Short Tasks and also draws on the English word "mass" to convey scale: it organizes vast numbers of Agents to undertake Microtasks that can be independently inspected. HAT stands for Human-Commanded, AI-Operated, Trust: a trust mechanism in which people exercise control and AI performs the work.

People establish, amend, and revoke the rules. Each AI Organizational Position has a named human accountable owner through Named Human Accountability Assignment. AI executes within its authorization. Governance Hats govern boundaries, not actions. When rules are absent, authority would be exceeded, or human discretion is needed, the affected execution pauses and the matter is referred to a human with the appropriate authority.

"Micro" refers to task granularity, not model capability. Breaking work into small tasks makes boundaries clearer and results easier to inspect; collaboration at scale lets full model capabilities realize their potential across different tasks. Governance-to-Agent Control safeguards rules and boundaries; Agent-to-Agent Collaboration handles assignment, execution, and review.

The ambition is "large pools of capital put to small-scale use"; the obstacle is "top talent cannot economically serve small tasks"; the answer is "powerful capabilities put to small-scale use."

The Twelve Sufficiency Conditions constrain the organization in four areas: tasks must be sufficiently small, narrow, and granular; foundations must be sufficiently diverse, cross-checked, and cost-effective; staffing must be sufficiently numerous, fine-grained, and counterbalancing; inspection must be sufficiently comprehensive, long, and deep. If any condition fails, a warning must be issued, the gap identified, and recalibration performed. Rule blind spots, data contamination, Common-Source Bias, and the propagation of erroneous judgments must be identified and addressed continuously. Judgments that cannot be broken into Microtasks are assigned to humans with the appropriate authority; AI may provide information and analytical support within its authorization.

Place powerful AI under a firm Governance Hat established by people.

Two Tensions and the Response




Figure 0-1: Two Sets of Tensions in the AI Era and the Response of HAT over MASS

V. The AI Factory Today

The 5,000 AI employee positions have been entered in the organizational register across four functions: Administration, Business, Control, and Deployment. "Taken up their positions" means that the positions have been entered in the organization's register; it does not mean 5,000 instances are running simultaneously.

AI Employee is the external designation for a position; an AI Organizational Position is a stable organizational identity. When a task arrives, the system starts an Agent according to the requirements of the position and calls on the capabilities needed for that run. The Agent Identifier (Agent ID) links the position to its accountable human,

and the Agent Capability and Performance Profile informs scheduling. Each instance run generates a corresponding Run Identifier linked to its task and to the inputs, outputs, calls, and inspection records for that run. A single task can involve multiple instances and repeated rework. Teams may be disbanded; identities and records are retained.

From task intake to archiving, seven stages form a delivery workflow and preserve the basis and responsibility for each step.

For example, a task to generate the Compass for chain convenience stores first defines its scope and delivery standards, then splits into work units such as Factor identification, specification of measurement bases, and Scoring Rules. It proceeds through seven stages: task intake, Rule Matching, team formation, execution, validation, Deliverable Consolidation, and archiving. Each stage records inputs, rules, outputs, and exceptions. Problems are first located and clarified; where changes are needed, work is returned to the appropriate stage for revision and reinspection. SIM-01 in this white paper is an illustrative simulation, not an actual operating result. The progress and inspection status of real tasks are established by retained execution and acceptance records.

VI. Deliverables to Date: 500 Industry Compasses

Each Industry Has Its Own Scale: Compass Deliverables for Five Hundred Industries

The Industry Compass rule deliverables in this edition cover 500 industries. The main text examines chain convenience stores, Sichuan-Chongqing hot pot restaurants, and freshly prepared tea beverage stores in detail; Appendix I includes 17 additional category cases, making 20 in total. These 20 cases are offered for detailed reading; they are not a statistical sample. Rule coverage for 500 industries does not mean that enterprises in those industries have been scored. The current cases show Indicator and Factor structures. Before enterprise application, users must still check the applicable version, measurement specifications, data sources, Observation Period, weights, and scoring standards, then collect data, conduct actual scoring, and perform backtesting. Released industry deliverables and version materials are available through the links.

For a score to withstand scrutiny, it must first specify what is measured, how it is calculated, and how long it is observed, then state which conditions fall into which score band. Applying the same rules to the same data should reproduce the same result. The Return Dimension and Risk Dimension are expressed separately, without addition or offsetting: readers must see both where operational support is strong and where risks lie.

The Industry Compass uses "3 Industry Perspectives + cross-checking by 3 models": the Ordinary Operator, Industry Veteran, and Adversarial Manipulation Perspectives examine the business from different positions. The Adversarial Manipulation Perspective looks for ways to dress up a business or game its scores. The three models form judgments separately before their assumptions, evidence, and differences are compared.

The Six Review Perspectives and Eligibility and Composition Review conduct design review, eligibility assessment, and Factor combination selection for candidate industry Factors. The six perspectives examine business relevance, reproducibility of calculations and scores, credibility of data sources and collection paths, the basis for distinguishing cases, incremental information relative to the full Factor pool, and resistance to manipulation. Eligibility and Composition Review uses these findings to assess candidate eligibility and then select combinations from the eligible pool. A candidate's "eligible for selection" status does not mean final selection, much less completed enterprise validation.

Industry research is divided into Microtasks and advanced in parallel. Multiple perspectives propose candidate judgments, multiple models cross-check them, the Six Review Perspectives review them, and Eligibility and Composition Review selects them. This accelerates the development of 500 Industry Compasses, with each deliverable retaining the basis for inspection and revision.

VII. Next Steps: Validation, Calibration, and Simulation

The next phase has four areas of work. First, enterprise data collection and actual scoring will ground readings in verifiable operating data. Second, backtesting and calibration will use operating and performance outcomes

to examine Factors, thresholds, and coefficients; problems will be diagnosed before the relevant revisions are made. Third, the entire system will run continuously in an environment explicitly marked as simulated, linking registration, measurement, task handoffs, inspection, alerts, and response, and incorporating adverse scenarios. This will test whether the stages work together and whether exceptions can be detected and handled. Fourth, development updates will be published to the market on the 3rd of each month, stating what has been done, how far it has progressed, and what evidence is still missing, according to actual progress. The rule deliverables for 500 industries are a starting point; enterprise application and ongoing inspection must still establish whether the rules are applicable and the readings reliable.

Simulated contracts, funds, allocations, and results must be clearly labeled and managed separately from real transactions and holding records. The invitation in 9-03 also sets the direction for the subsequent simulation experience:

Your Order. Your Portfolio. Your SHARP.

VIII. Oversight and Inspection of HAT over MASS

Micro Connect New Markets actively embraces and complies with regulation and provides verifiable rules, system records, and execution records for regulatory inspection. This white paper explains inspection of HAT over MASS through oversight of people, examination of the system, and review of standards. Oversight of people checks each position's Named Human Accountability Assignment, authorization, and records of rule establishment, amendment, and revocation. Examination of the system checks task inputs, Rule Versions, execution, inspection, and exception handling records; inspection interfaces, replay capabilities, and their limitations are described according to actual development. Review of standards checks whether evidence supports the Twelve Sufficiency Conditions and whether warnings and recalibration occur when conditions are insufficient. People govern rules and exceptions; the system governs execution and recordkeeping; independent inspection checks whether outputs and processes meet requirements.

People bear responsibility under regulatory requirements and manage the system through rules and authorization; the system organizes execution under those rules and retains records. This division applies system governance to the AI factory. Human work centers on rules, authorization, supervision, and exception handling; professional confirmation and human discretion required by law or institutional rules remain with people who hold the corresponding responsibilities.

IX. Division of Responsibilities with Licensed Institutions

Investors complete transactions through institutions holding the appropriate licenses. Investment decisions are made by parties with the relevant decision-making authority; investor suitability, related services, and compliance obligations are handled by licensed institutions performing those functions within the scope of their licenses. After transactions and settlement are complete, actual holding relationships are entered in the Contract Portfolio Registry and linked to original trade records. The Contract Registrar records the contract itself and its status, Compass Reading is updated on the basis of new facts and evidence, and the Compass Terminal serves licensed institutions or qualified overseas investors with the appropriate qualifications and authorization.

Micro Connect New Markets is responsible for registration, measurement, and the infrastructure for using these materials. It does not select assets or allocate holdings for investors. Contract registration is based on collected and verified contract information. Portfolio registration is based on completed transactions, settlement, and subsequent changes in holdings. Versions and supporting bases are retained for rules and records. Micro Connect New Markets is not a securities company; it does not match trades, list securities, establish pools of client funds, or handle client assets. Micro Connect New Markets is responsible for the soundness of its own system governance and the reliability of its infrastructure: data errors must be corrected with historical records retained; the effects of system exceptions must be identified, the faults corrected, and the affected functions reinspected; problems in Factor definitions and rules must be revised through governance procedures. The

absence of a commitment to investment outcomes does not remove responsibility for registration, data, systems, and services themselves.

The public may read general industry research materials. Formal investment participation is available only to licensed institutions or qualified overseas investors meeting the requirements of the applicable jurisdiction, through appropriately licensed intermediaries, and is unavailable to retail investors. No fundraising is conducted within China, and domestic funds are not accepted for investment participation.

X. How to Read and Help Improve This White Paper

The main text has six parts: the business challenge, Compass tools and rules, HAT over MASS theory, the AI work manual, case deliverables and the outlook. General readers can use Parts I and II to understand the problem and instruments, and Part III to understand how HAT over MASS is organized. The work manual in Part IV sets out responsibilities for the named humans accountable for AI positions and enables technical due diligence parties and AI auditors to inspect organizational and operating arrangements. Operators, industry experts, and investment researchers may start with Part V and Appendix I. Subsequent validation arrangements appear in Part VI.

Trust has no passbook; it is settled daily.

Micro Connect New Markets discloses its methods, deliverables, and boundaries so that the market can raise specific questions.

Whether Factors capture what matters to a business, whether data can be obtained and verified, whether rules allow score manipulation, and what should change when counterexamples arise can all be discussed item by item. Each evidence-based revision should bring the methodology closer to real operations. Industry coverage still cannot replace verification of specific projects and ongoing observation.

The map is gradually unfolding, the compasses continue to be calibrated, and the organization has taken up its positions.

Micro Connect New Markets invites business operators, industry experts, financial institutions, and AI technology partners to contribute, and welcomes guidance from regulators, so that the methodology continues to improve through practice and inspection.

Access to Industry Compass Deliverables

Start with the 3 cases in the main text and 17 in the appendix, then explore Compass deliverables for 500 industries

Readers can scan the QR code below, register the requested information, and complete verification to view Indicator and Factor structures, released rule materials, and version descriptions by industry. Current access for each industry is governed by official disclosures.

Note: Online viewing only; downloads are unavailable.



Industry Compass website: <https://microconnect.com/industry-compass>

Version Relationships and Conditions of Use

Version relationships: *This is the third white paper published by Micro Connect New Markets on October 3, 2026 (the "current edition"). It connects with the first white paper published on August 3 and the second published on September 3. Each has its own focus; together they explain the public framework of Micro Connect New Markets. Where this white paper expressly updates the same matter, the relevant explanation in this white paper prevails.*

Basis of deliverables: *This white paper presents the theoretical framework of HAT over MASS, the initial organizational practice of Micro Connect New Markets, and the current round of Industry Compass deliverables. The methods and rules will be continuously tested and improved using actual operations, enterprise data, and feedback from all parties. "5,000 AI Employees have taken up their positions" and "coverage of 500 industries" refer respectively to positions entered in the organization's register and industry rule deliverables. They do not equate to specific task run records or validation of individual enterprises. The chain convenience store task chain is an illustrative simulation.*

Conditions of use: *This white paper, its accompanying Industry Compasses, and the NUMA Adjustment Factor are solely for general industry research, business analysis, and methodological discussion. They do not constitute investment advice, recommendations, any transaction offer, product offering, or solicitation of investment, nor any commitment to returns, legal opinion, or compliance approval. Scores, coefficients, and rankings do not represent investment returns, default probabilities, credit ratings, or investment priorities. The Industry Compass cases in this edition primarily show Indicator and Factor structures; they do not indicate completed data collection, scoring, or validation for specific enterprises. Illustrative simulations do not represent actual operating performance. Market participants should make decisions based on their own independent judgment, due diligence, and professional advisors' opinions. AI only organizes evidence, performs checks, and develops recommendations within its authorization. Final scoring and investment, suitability, compliance, or legal judgments must be made or confirmed by people with the appropriate authorization and qualifications. Services that legally require a license must be provided by appropriately licensed parties within the scope of their licenses. Internal authorization or human confirmation cannot replace those licensing requirements.*

Scope of participation: *This white paper and the accompanying Industry Compasses provide general research information to the public. Public access does not constitute the provision of investment services to readers or confirmation of their eligibility to invest. Investment participation in Micro Connect New Markets is available only to professional investors meeting applicable jurisdictional requirements, through appropriately licensed intermediaries, and is unavailable to retail investors. Micro Connect New Markets serves assets worldwide; on the funding side, it serves qualified professional investors outside China. Micro Connect New Markets conducts no fundraising within China and accepts no domestic funds for investment participation.*

Contents

Part I The Business Challenge	15
Chapter 1 From the Scale of Companies to the Scale of Contracts	16
1.1 The Contribution and Processing Scale of Corporate Securitization	16
1.2 Contracts as Smaller Units for Understanding Business Operations	16
1.3 The Ambition of "Large Pools of Capital Put to Small-Scale Use" and the Cost Constraint of "Top Talent Cannot Economically Serve Small Tasks"	16
1.4 Evidence and Responsibility Must Hold Even as Costs Fall	17
Part II Compass Tools and Rules	18
Chapter 2 The Uses and Division of Responsibilities of Registration and Measurement Instruments	19
2.1 Contract Terminal.....	19
2.2 Contract Registrar	19
2.3 Contract Compass	19
2.4 Compass Reading.....	20
2.5 Contract Portfolio Registry	20
2.6 Compass Terminal	20
2.7 From Instrument Functions to This Edition's Deliverables	20
Chapter 3 How the Industry Compass Explains a Business	22
3.1 A Three-Level Structure	22
3.2 Quantity and Selection Rules	22
3.3 Scoring Rules	22
Chapter 4 What the 500 Industry Compasses Have Delivered	24
4.1 Deliverable Coverage and the Index of 500 Industries	24
4.2 Factor Definitions, Units, Observation Periods, and Scoring Standards	26
4.3 Deliverable Versions, Evidence Records, and Subsequent Calibration.....	26
4.4 How to Find Your Industry's Deliverables and Help Improve Them.....	27
Chapter 5 How Industry Deliverables Enter Micro Connect New Markets	29
5.1 The Core Compass Examines Contract Structure; the Industry Compass Examines the Specific Business.....	29
5.2 Core Codes and Industry Codes: Their Respective Return and Risk Scores	29
5.3 The NUMA Adjustment Factor and Cross-Industry Calibration.....	30
5.4 A Reading Path from Within-Industry Judgment to Cross-Industry Comparison	31
5.5 Conditions of Use: Industry Deliverables, Project Verification, and Ongoing Observation	32
Part III HAT over MASS Theory	33
Chapter 6 Two Sets of Tensions Between AI Capability and Human Responsibility	34
6.1 The Dilemma of Letting AI Act and Watching AI Closely	34
6.2 The Financial Industry's Essential Questions: Compliance, Investor Protection, and Risk Control	34
6.3 How HAT over MASS Addresses Both Sets of Tensions	35
Chapter 7 The HAT Governance Hat: People Control the Rules; AI Executes Within Them.....	37
7.1 The Meaning of HAT: Ultimate Human Responsibility and the Release of AI's Potential	37
7.2 Rules Before Capability: Permitted Actions, Prohibitions, and Escalation Conditions	37

7.3 A Named Human Behind Every Governance Hat: Accountability Assignment and Rule Establishment, Amendment, and Revocation.....	38
7.4 Governance Hats Govern Boundaries, Not Actions.....	38
7.5 Human Responsibility for Governance, Supervision, and Discretion.....	38
Chapter 8 MASS: How Powerful Capabilities Enter Microtasks.....	40
8.1 The Meaning of MASS: Micro Agents, Simple and Short Tasks, and Collaboration at Scale.....	40
8.2 Small Tasks, Strong Agents: Full Model Capability and Limited Task Authority.....	40
8.3 Three Sufficiency Conditions for Task Decomposition: Sufficiently Small, Narrow, and Granular, with Business Meaning Intact.....	41
8.4 Establish Verifiability First, Scale Next: Independent Acceptance Review, Error Location, and Local Rework.....	42
8.5 Three Sufficiency Conditions for Supporting Foundations: Diversity, Cross-Checking, and Cost-Effectiveness.....	42
Chapter 9 How Microtasks Form Reliable Collaboration.....	44
9.1 Governance-to-Agent Control and Agent-to-Agent Collaboration: Two Channels for Rule Governance and Task Collaboration.....	44
9.2 From Decomposition to Assembly: Task Dependencies, Handoff Interfaces, and Deliverable Integration.....	45
9.3 From Disagreement to Resolution: Clarification, Review, Return, and Escalation to an Authorized Human.....	45
9.4 Qualified Components and Complete Delivery: Separation of Authority and Checks and Balances Among Scheduling, Acceptance Review, and Backtesting.....	46
9.5 Every Agent Is Assigned to a Named Human: Temporary Collaboration and Continuing Responsibility.....	47
9.6 Three Sufficiency Conditions for Expanding Staffing: Sufficiently Numerous, Fine-Grained, and Counterbalancing, Within Supporting Capacity.....	47
Chapter 10 How Every Delivery Can Withstand Inspection.....	49
10.1 Full-Coverage Inspection Rather Than Sampling: Every Output Is Inspected.....	49
10.2 Ongoing Backtesting: Identifying Short-Term Luck and Environmental Change.....	49
10.3 End-to-End Traceability: From Results Back to Inputs, Calls, Rules, and Responsibility.....	50
10.4 Three Sufficiency Conditions for Inspection: Sufficiently Comprehensive, Long, and Deep.....	51
Part IV AI Operating Manual.....	52
Chapter 11 How 5,000 AI Employees Fit into the Organization [Operating Procedures].....	53
11.1 The 5,000 AI Employees Already in Position.....	53
11.2 Three Structural Layers: Rules, Responsibilities, and Identification.....	54
11.3 Four Functions: Administration, Business, Control, and Deployment.....	55
11.4 Multi-Perspective Roles: Three Industry Perspectives + Cross-Checking by Three Models, Three Investor Perspectives, and "6+1" Review Perspectives.....	56
11.5 Separating Perspectives from Capabilities: Calling on Models, Skills, and Tools by Task.....	57
Chapter 12 How Responsibility Is Distinguished between a Position and a Run [Operating Procedures].....	59
12.1 AI Employee Positions and Task Instances.....	59
12.2 AI Employee Agent Identifiers, Capability and Performance Profiles, and Run Identifiers.....	60
12.3 Separately Managing the AI Position Register, Status Updates, and Instance Information.....	61
12.4 Records Linked to AI Employee Run Identifiers and Evidence Retention.....	62
Chapter 13 How Industry Compass Tasks Are Decomposed and Teams Formed [Operating Procedures].....	63

13.1 Defining an Industry Compass Task: Input Materials, Delivery Standards, and Constraints	63
13.2 Decomposing an Industry Compass Task: Work Units for Factor Identification, Measurement Specifications, and Scoring Rules.....	64
13.3 Allocating AI Employees: Matching Positions by Capability and Performance Profiles, Permissions, and Availability	64
13.4 Forming AI Employee Teams: Establishing Task-Specific Teams, Handling Disagreements, and Ending Collaboration.....	65
13.5 The Seven-Stage Task Workflow for AI Employees: Task Intake, Rule Matching, Team Formation, Execution, Validation, Deliverable Consolidation, and Archiving	66
13.6 Example: A Chain Convenience Store Industry Compass: From a Task Sheet to a Candidate	67
Task Intake: Narrowing the Meaning of "Chain Convenience Store"	67
Team Formation and Execution: Multiple Perspectives Begin Work Together	68
Chapter 14 How Continuing Positions Support Dynamic Collaboration [Operating Procedures]	69
14.1 The Departmental Affiliation Line and the Task Workflow Line.....	69
14.2 Temporary Teams Cover Four Functions, and Continuing Responsibility Records Remain Intact	70
14.3 The Scheduling Agent Handles Routine Team Formation, Handoffs, and Overall Delivery Coordination.....	71
14.4 Named Humans Handle Rules, Authorization, and Escalated Matters Outside the Rules	72
14.5 Overall Delivery and Component Responsibilities Coexist, with Acceptance Review and Control Remaining Independent	72
14.6 Example: A Chain Convenience Store Industry Compass: Who Clarifies, Revises, and Decides When Disagreements Arise	73
Locate First, Then Clarify, and Return Work Only When Revision Is Clearly Needed	73
Chapter 15 How a Deliverable Is Inspected and Revised [Operating Procedures]	76
15.1 Acceptance Review Is the Admission Gate for Task Progression	76
15.2 The "6+1" Review Perspectives: Responsibilities of the Six Review Perspectives and Eligibility and Composition Review, Inspection Outputs, and Returns for Nonconformance	77
15.3 Checking Inputs, Outputs, Rules, and Responsibility through Run Records	79
15.4 Backtesting against Facts Drives Capability Calibration and Rule Revision.....	80
15.5 Authorization Checks, Version Management, and Record Completeness	80
15.6 Example: A Chain Convenience Store Industry Compass: Why V0.2 Is Pending Clarification after the Six Design Reviews	81
What Acceptance Review Examined	82
Deliverable Consolidation and Archiving: Record Where Work Stops	82
Traceability and Backtesting: Subsequent Implementation Takes Up the Next Stage of Inspection	82
Part V Case Deliverables.....	84
Chapter 16 Close Readings of Three Categories and How to Read the Cases	85
16.1 Close Readings of Three Categories	85
Chain Convenience Stores: Ten Indicators and Seventeen Factors.....	85
Return Dimension: From Selling More to Sustaining Earnings	85
Risk Dimension: Costs, Losses, and Expansion	86
What the Table Does Not Show: Why One Candidate Is Absent.....	86
Sichuan-Chongqing Hot Pot Restaurants: Twelve Indicators and Sixteen Factors	86
Return Dimension: Table Output and Customer Loyalty	87
Risk Dimension: Advance Receipts, Supply, Fixed Costs, and Cash Verification	87

Freshly Prepared Tea Beverage Stores: Eleven Indicators and Seventeen Factors	87
Return Dimension: Cup Output, Transaction Value, and Single-Store Returns	88
Risk Dimension: Cup Output Authenticity and Operating Requirements	88
16.2 Selecting and Reading the Seventeen Cases.....	89
Part VI Future Outlook.....	90
Chapter 17 From This Round's Deliverables to the Next Delivery.....	91
17.1 This Round's Theory, Practice, and Deliverables	91
17.2 Enterprise Data Collection, Calibration, Simulation, and Progress Disclosure	91
17.3 Continuous Operation in a Simulated Environment	92
17.4 Industry Feedback and Joint Inspection.....	92
Appendix I Seventeen Category Cases and Navigation to Research Deliverables for 500 Industries.....	93
Case Scope and Reading Guide.....	93
Limits on Use	93
Navigation to the 500-Industry Research Results	93
Case Index.....	93
01 Burger Fast-Food Stores	95
Return Dimension 7 Indicators 9 Factors.....	95
Risk Dimension 7 Indicators 9 Factors.....	95
02 Bakery Stores.....	96
Return Dimension 5 Indicators 9 Factors.....	96
Risk Dimension 7 Indicators 9 Factors.....	96
03 Car Rental.....	97
Return Dimension 4 Indicators 8 Factors.....	97
Risk Dimension 5 Indicators 9 Factors.....	97
04 Discount Snack Retail	98
Return Dimension 6 Indicators 9 Factors.....	98
Risk Dimension 6 Indicators 8 Factors.....	98
05 Chain Pharmacy Stores	99
Return Dimension 5 Indicators 9 Factors.....	99
Risk Dimension 6 Indicators 8 Factors.....	99
06 Residential Elder Care Facilities	100
Return Dimension 4 Indicators 6 Factors.....	100
Risk Dimension 5 Indicators 7 Factors.....	100
07 Economy Business Hotels.....	101
Return Dimension 3 Indicators 8 Factors.....	101
Risk Dimension 4 Indicators 7 Factors.....	101
08 Teahouse Operations	102
Return Dimension 5 Indicators 9 Factors.....	102
Risk Dimension 5 Indicators 7 Factors.....	102
09 Shared Massage Chair Operations	103
Return Dimension 5 Indicators 9 Factors.....	103

Risk Dimension 6 Indicators 8 Factors.....	103
10 Comprehensive Luxury Goods Care	104
Return Dimension 6 Indicators 7 Factors.....	104
Risk Dimension 7 Indicators 8 Factors.....	104
11 In-Home Pet Care Services	105
Return Dimension 4 Indicators 5 Factors.....	105
Risk Dimension 7 Indicators 7 Factors.....	105
12 Small Immersive Performance Venues.....	106
Return Dimension 5 Indicators 9 Factors.....	106
Risk Dimension 4 Indicators 6 Factors.....	106
13 Instant Retail Fulfillment Warehouses.....	107
Return Dimension 6 Indicators 9 Factors.....	107
Risk Dimension 6 Indicators 9 Factors.....	107
14 Esports Event Services	108
Return Dimension 6 Indicators 8 Factors.....	108
Risk Dimension 6 Indicators 8 Factors.....	108
15 Short-Form Drama Production	109
Return Dimension 5 Indicators 8 Factors.....	109
Risk Dimension 6 Indicators 7 Factors.....	109
16 AI Computing Center Hosting	110
Return Dimension 6 Indicators 8 Factors.....	110
Risk Dimension 6 Indicators 8 Factors.....	110
17 AI Model Deployment and Operations.....	111
Return Dimension 4 Indicators 9 Factors.....	111
Risk Dimension 4 Indicators 7 Factors.....	111
Appendix II Overview of the Twelve Sufficiency Conditions.....	112
I. Overview of the Twelve Sufficiency Conditions	112
II. Limits on Use of the Overview	114
Appendix III Term Definitions and Chinese–English Correspondence	116

Part I The Business Challenge

Chapter 1 From the Scale of Companies to the Scale of Contracts

1.1 The Contribution and Processing Scale of Corporate Securitization

Corporate securitization enables capital of different scales to form investment relationships around the same company. Companies organize business activities; financial statements, disclosures, and professional services bring fragmented information together into objects that can be examined and compared. This institutional framework enables capital to participate in economic activity under clear arrangements of rights and responsibilities without directly operating every business.

Using companies as the primary unit of observation also establishes a scale of processing. Understanding a company requires verification of materials, analysis of operations, identification of risks, and ongoing monitoring; these activities have fixed costs. Larger companies with more complete materials are better placed to bear those costs. Professionals struggle to understand the vast number of small and micro businesses across the economy to the same depth. The framework can serve a broader economy, but the cost and capacity of professional work may not keep pace.

1.2 Contracts as Smaller Units for Understanding Business Operations

Business activity ultimately takes shape in specific contracts: who provides what to whom, which conditions trigger payment, and when and along what path cash flows arise. A company's overall performance matters, but the basis of rights, conditions of performance, and operating evidence in an individual contract also have independent observational value. Moving down to the contract level reveals differences easily obscured in aggregate data.

Cash flows specified in a contract are not equivalent to cash flows already realized or certain to be realized. Real operations change, payments may be delayed, and information may be incomplete. Finer granularity therefore requires separate recording of contract terms, operating facts, evidence sources, and Observation Periods to distinguish what has been established from conditions, forecasts, and matters still requiring verification.

The earlier difficulty was not that professionals failed to understand small businesses. The cost of understanding, verifying, and continuously observing each contract could exceed what a small transaction could bear. Extending services from companies to vast numbers of contracts requires more than repeating the same work more often. The cost of repetitive processing must fall at the same time, so that evidence and responsibility can connect at a finer level.

1.3 The Ambition of "Large Pools of Capital Put to Small-Scale Use" and the Cost Constraint of "Top Talent Cannot Economically Serve Small Tasks"

This cost problem is particularly concrete in small and micro businesses. Our ambition of "large pools of capital put to small-scale use" means dividing substantial wealth into smaller funding units and directing them to small factories, shops, and farms, reaching the capillaries of the real economy that traditional finance finds difficult to serve. "Small-scale use" changes the scale at which funding enters real operations; it does not lower the requirements for using that funding. We want more small, dispersed operators to have an opportunity to be understood seriously and served effectively.

Smaller funding units do not make judgments simpler. A small shop still has revenue, costs, performance obligations, and cash collection. A small factory still requires examination of orders, capacity, inventory, and operating risks. Every small, dispersed, nonstandard task may still require verification of materials, operating analysis, pricing judgment, and continuous monitoring. The more numerous the businesses and the greater the differences among industries, the finer the information to process and the more fragmented the judgments become.

Traditional professional services rely on experienced people to handle complex matters individually, and that capability is irreplaceable. The difficulty is that devoting the same amount of senior professional time to every small transaction creates a mismatch between talent cost and task scale. Professionals can certainly perform small tasks. Yet an expensive, sequential, item-by-item way of working cannot cover vast, dispersed demand. Each individual transaction can be handled with great care, yet the overall economics fail to work. This is the constraint behind "top talent cannot economically serve small tasks."

For the first time, AI reveals another possibility: dividing powerful analytical and processing capabilities into finer units and putting them into these tasks. Organizing materials, comparing facts, matching rules, and conducting structured checks can all be divided into units with clearer boundaries and executed repeatedly by machines at lower marginal cost. Some capabilities previously affordable only for large transactions may thus become available to more, smaller businesses.

Assigning tasks to AI addresses cost and scale, but it does not yet establish responsibility. Allowing AI to take on an oversized task and exercise judgment freely within vague boundaries may cost less, yet make errors harder to locate and responsibility harder to assign. "Large pools of capital put to small-scale use" requires more than faster, cheaper automation. It needs an organizational method that places powerful capabilities inside manageable, inspectable tasks that can be escalated to a human when necessary. Only then can the obstacle of "top talent cannot economically serve small tasks" be addressed through "powerful capabilities put to small-scale use."

1.4 Evidence and Responsibility Must Hold Even as Costs Fall

AI changes the processing cost of some professional work. Work whose inputs, rules, outputs, and inspection standards can be specified in advance can be executed repeatedly by machines. Fact verification, information structuring, and continuous comparison may thereby extend to more businesses. A machine's ability to process more tasks does not grant it authority over issuance, transactions, investment judgments, or legal discretion.

Moving from the business challenge to tools first requires answers to three questions: how to record dispersed contract facts clearly, how to measure them under consistent specifications, and how to bring new facts continuously into readings. Part II explains the uses and rules of these instruments; Part III then explains why the organization that manufactures and maintains them adopts HAT over MASS. Discussing the problem, tools, and organization separately makes clear what each layer actually addresses.

Part II Compass Tools and Rules

Chapter 2 The Uses and Division of Responsibilities of Registration and Measurement Instruments

Micro Connect New Markets is an infrastructure provider for the contract securitization market, responsible for registration and measurement. Our instruments serve licensed institutions and human financial experts. They connect contract terms, operating evidence, and applicable rules so that users can retrieve facts, review readings, and form their own professional judgments. HAT over MASS manufactures and maintains these instruments; it does not replace licensed institutions in making investment judgments and conducting transactions.

The 8-03 white paper sets out the market framework for investment and financing based on contractual cash flows; the 9-03 white paper establishes rules for measuring contracts and businesses and functions for registration, observation, and portfolio records. Following a contract from admission through its ongoing term, this white paper organizes those functions into six instruments: the Contract Terminal, Contract Registrar, Contract Compass, Compass Reading, Contract Portfolio Registry, and Compass Terminal. This describes functions and responsibilities. The development status, users served, and scope of access of each instrument must be assessed separately under its own official version and actual disclosures.

2.1 Contract Terminal

The Contract Terminal is the first stop for contract admission: an interactive system that maintains a continuing dialogue with the issuer or the issuer's AI Agent. It proactively requests materials, enters and verifies each item, assesses them in real time, and rejects materials that fail to meet the rules or requests further information. It organizes nonstandard contracts into structured records with sources, timestamps, and verification status. AI inference cannot fill gaps in missing facts; satisfactory collection of materials does not mean that a contract has been approved for issuance, listing, or trading.

The Contract Terminal interacts continuously with the issuer and the professionals responsible for verification, checking whether the contract is in force, is genuinely being performed, and has cash flows verifiable by a third party. AI collects and checks information under the rules. Contract interpretation, legal characterization, and judgments not covered by rules remain with people who have the appropriate authority and qualifications. Admission only confers eligibility to be measured; it does not mean the contract is worth holding.

2.2 Contract Registrar

The Contract Registrar records a contract's identity, terms, related parties, initial profile, operating status, and underlying evidence. Its inputs are facts collected and verified under the relevant rules, along with subsequent updates. Its outputs are standardized records that can be queried, compared, tracked, and traced to their sources. It retains historical versions so that users can distinguish initial conditions, subsequent changes, and current status.

The Contract Registrar records; it does not issue. Registration does not create assets, change the rights and obligations of the original contract, or endorse the contract or its future performance. Registration status and evidence status should be shown separately. "Registered" must not be taken to mean "all facts verified."

2.3 Contract Compass

The Contract Compass provides a measurement framework for registered contracts. The Core Compass examines contract structure; the Industry Compass examines the specific business behind the contract. Each has a Return Dimension and a Risk Dimension, presented separately, without addition or offsetting. Inputs comprise facts with clearly identified objects, sources, and periods, together with applicable Factor definitions, weights, and Scoring Rules. Outputs are readings that retain evidence, versions, and conditions of use.

The Compass measures; it does not provide credit ratings or investment recommendations. It does not output "buy" or "sell," and its scores do not represent investment returns, default probabilities, or credit ratings. The same facts under the same Rule Version should produce consistent calculations and score-band assignments. Where evidence is insufficient, the result is left blank. The compass points north; it does not prescribe a route.

2.4 Compass Reading

Compass Reading consists of Factor data updated daily for registered contracts, continuously tracking changes in operations, performance, and evidence. Inputs are newly obtained facts, contract events, and corresponding evidence. Outputs are updated records carrying data timestamps, Observation Periods, and Rule Versions. Updates should preserve the relationship between earlier and later records. New data cannot overwrite the earlier conditions on which historical results relied.

Compass Reading is data, not a recommendation for investment action. A daily update does not mean that every field is mechanically rewritten when there is no new evidence. Each field is updated according to its nature, data sources, and actual changes. Missing data, conflicting sources, and delayed updates must be clearly marked; old readings cannot be treated as new facts.

2.5 Contract Portfolio Registry

After investors complete transactions and settlement through licensed institutions, the Contract Portfolio Registry records daily which contracts are in which portfolios, which contracts each portfolio contains, and how those relationships change with subsequent events. In the 9-03 white paper, the Order Register records trade, holding, distribution, and exit facts. The Contract Portfolio Registry links these facts at the portfolio level; it does not replace or rewrite original transaction records.

Inputs are valid trade, settlement, and holding records; outputs are portfolio relationships traceable to the corresponding transactions and underlying contracts. Pre-investment screening and simulated allocations should be separated from actual holdings; a draft is not a portfolio whose trades have been completed. Observation and calculation of portfolio constituents follow explicit rules. The platform does not select assets for users. The Contract Portfolio Registry records only relationships and facts and does not handle client assets. Recordkeeping does not constitute custody of funds.

2.6 Compass Terminal

The Compass Terminal brings registration facts, Compass rules, readings, and evidence to licensed institutions or qualified overseas investors. Inputs are materials authorized for access and observation conditions set by the user; outputs are professional views that can be expanded, compared, and traced to their supporting basis. Users retrieve readings and form judgments under their own methodologies. Adjustments to personal parameters do not rewrite official facts, default rules, or other users' results.

The terminal serves licensed institutions or qualified overseas investors with the appropriate qualifications and authorization. Public access to general research materials through the Industry Compass website and access to the professional terminal are separate arrangements. Public reading does not confirm investment eligibility, and providing an observation interface does not grant the terminal authority to make judgments or trade for clients. Specific functions and access arrangements are governed by the corresponding service descriptions.

2.7 From Instrument Functions to This Edition's Deliverables

The six instruments connect collection, registration, measurement, daily updates, portfolio records, and professional use. They systematically organize underlying contract data so that it can be read and understood. Their reliability depends on whether factual sources, Rule Versions, authorization boundaries, and operating records can be checked against one another. Investment decisions are made by parties with the relevant decision-making authority. Licensed institutions handle suitability, transactions, and duties to investors according to their respective functions and within the scope of their licenses. The infrastructure provider remains accountable for the registration, data, and system responsibilities it actually undertakes.

This edition focuses on disclosure of Industry Compass structures, rule deliverables, and the organizational method used to produce them. Coverage of 500 industries does not mean completed data collection, scoring, or validation for 500 enterprises. The 3 cases in the main text and 17 in the appendix, 20 in total, are presented at the level of Indicator and Factor structures. Actual use requires checking the official version. The next three chapters explain the Industry Compass, deliverable versions, and conditions for using the Twin Compasses, respectively.

Chapter 3 How the Industry Compass Explains a Business

The Industry Compass provides scoring rules specific to a single asset category within the NUMA SHARP Twin Compasses, alongside the Core Compass that is uniform across the market. The Core Compass examines contract structure; the Industry Compass examines the specific business. It identifies which measurable industry facts determine returns, which determine risks, and the corresponding industry anchors. Rules are publicly released after adversarial testing and human review.

3.1 A Three-Level Structure

The Industry Compass uses the dimension–indicator–factor structure (Table 3-1). Dimensions are fixed as the two directions of returns and risk, scored and presented separately; return support does not offset risk exposure. Indicators define "what to ask" within each Dimension; Factors are operating variables for which data can be obtained and verified, answering "what specifically to examine."

Level	Definition	Quantity Rules
Dimension	Two directions of observation: the Return Dimension and Risk Dimension	Two fixed Dimensions for all categories
Indicator	An evaluation item within a Dimension that defines "what to ask"	3 to 9 Indicators per Dimension, with content determined by the category
Factor	A verifiable operating variable that supports judgment of an Indicator	3 to 9 Factors per Dimension; at least 1 Factor per Indicator

Table 3-1: The Industry Compass's Three-Level Structure

3.2 Quantity and Selection Rules

The content and number of Indicators and Factors may vary. Within one Dimension of a category, there are 3 to 9 Indicators and 3 to 9 Factors. Each Indicator corresponds to at least 1 Factor and may correspond to only 1. Key Indicators and Factors are selected for each category according to its business complexity and degree of standardization, so that the Industry Compass can accommodate the diversity of industries.

Industry Factor selection follows four constraints. Each Factor belongs to only one Indicator; the same information is scored only once. Factors must have clear operating meanings and verifiable data sources; they cannot be included if they rely on subjective judgments that cannot be reviewed. Indicator and Factor composition and default weights are disclosed with the Rule Version, and changes must go through governance procedures. Investors may adjust Factor selection and weights in their personal interfaces.

Common Return Dimension Factors concern demand and revenue realization, operating efficiency and output per unit, and revenue quality and continuity. Common Risk Dimension Factors concern cost and execution risk, concentration and dependency risk, and cyclical and compliance shocks. These are common options, not a fixed list.

3.3 Scoring Rules

Each Factor must specify its data definition and source, category scale and anchors, and the conditions for assignment to scores from 1 to 3. Scoring directions are consistent across the two Dimensions: stronger return support yields a higher score for Return Dimension Factors; higher risk yields a lower score for Risk Dimension Factors. During the current launch phase, a score of 2 is a reference point or estimated anchor. It may be described as the category median only after sufficient samples and independent validation.

Scoring requires evidence. Scores with no supporting evidence or with less credible evidence are marked gray. This indicates only the strength of evidence and does not change the score.

Factor scores form Indicator scores under publicly disclosed weights, then aggregate into Return Dimension and Risk Dimension scores. They are normalized to a common score range so that categories with different numbers of Indicators and Factors can be presented on the same scale. Dimension scores form industry SHARP after Cross-Industry Calibration using the NUMA Adjustment Factor; together with the Core Compass, they form NUMA SHARP for an individual contract. Industry Compass performance cannot offset the Core Compass's minimum admission standards. Applicable conditions and the scope of comparability are set out in Chapter 5, Section 5.3.

Chapter 4 What the 500 Industry Compasses Have Delivered

Once the method is clear, the deliverables need to be examined. This chapter explains which operating activities the 500 Industry Compasses cover, what each set of rules contains, how readers can check definitions, units, periods, and score bands, and how to find their sources and versions.

An industry deliverable should take readers from the industry name to specific rules, showing which facts support returns and risk. A common structure makes content easy to find; Factors explain the business accurately. Coverage counts show how broad the work has been. Clear rules, versions, and conditions of applicability explain how the deliverables can be used.

Figure 4-1 summarizes the path through which multiple perspectives come together as Industry Compass deliverables.




Figure 4-1: Multiple Perspectives Combine into Industry Compass Deliverables (the 20 cases in the figure comprise 3 in the main text and 17 in the appendix)

4.1 Deliverable Coverage and the Index of 500 Industries

This research round covers 500 industries and produces research deliverables on industry identification. This edition develops 17 formal category cases, included in Appendix I, with three additional categories—chain convenience stores, Sichuan-Chongqing hot pot restaurants, and freshly prepared tea beverage stores—examined in detail in the main text. The scope of operating activities, quantitative specifications, data access paths, weights, score bands, and review status of each industry are disclosed by deliverable version. Coverage counts do not mean that 500 enterprises or 500 contracts have been scored and validated.

Level	Presentation in This Edition	Limits of Interpretation
Organizational configuration	5,000 AI Employees have taken up their positions; four functions, with positions and instances managed separately	Position counts do not mean 5,000 instances running simultaneously; a disclosable register and statistical basis remain to be provided
Industry rules	Coverage of 500 industries; this white paper presents Indicator and Factor structures for 17 formal category cases, with three additional category close readings in the main text	Coverage does not mean completed scoring or validation of 500 enterprises
Process illustration	The chain convenience store SIM-01 runs through tasks, collaboration, inspection, and revision	An illustrative simulation, not actual backend run records
Enterprise application and calibration	Conditions for data collection, scoring, backtesting, and adjustment	Actual enterprise measurements, coefficient versions, and approval and validation records are disclosed separately according to actual evidence

Table 4-1: Deliverables and Validation Status in This Round

Find the industry first, then confirm applicability. Businesses with similar names may have different revenue, cost, and risk sources; a single enterprise may also operate multiple businesses. The index therefore needs to connect identifiers, standard names, aliases, scopes, and versions so that users can find the rules that truly match.

A usable industry index needs to answer the questions in the following table.

Index Field	Question Answered	Purpose for Readers
Industry Identifier	How the deliverable is consistently identified	Avoiding loss of historical correspondence after a name changes
Standard name and necessary aliases	What the market usually calls this business	Helping users reach the correct entry through familiar names
Scope of applicability	Which operating model the rules address	Preventing rules from being mixed across businesses with similar names but different economic structures
Rule Version and status	Which version is being viewed and its current status	Distinguishing historical drafts, current versions, and restrictions on use
Update date and effective date	When content was amended and when it became applicable	Supporting review of historical results and comparison over time
Deliverable contents	Which Indicators, Factors, and Scoring Specifications are included	Taking readers from the index into specific rules rather than leaving them at an industry list

Table 4-2: Fields and Reading Purposes of the Index of 500 Industries

After finding an industry, users should be able to open its Return Dimension and Risk Dimension, examine Factors, definitions, data specifications, units, periods, and scoring standards, and then return to the version and amendment records. Industry splits, mergers, and renaming must also retain original identifiers, succession relationships, and applicable dates so that historical references can still locate the rules in force at the time.

Industry rules explain how to judge; specific readings must be linked to the business being observed, factual evidence, observation time, and Rule Version. Deliverables for all 500 industries do not mean that every enterprise already has Industry Codes, or that each industry has undergone equal sampling and backtesting. The index should accurately state maturity so that a uniform appearance does not conceal different statuses.

4.2 Factor Definitions, Units, Observation Periods, and Scoring Standards

Each Factor deliverable comprises elements such as its definition, quantitative specification, data sources, unit, Observation Period, and standards for scores from 1 to 3 (Table 4-3). The same name can mean different things under different specifications or periods; these must be aligned before scoring.

Scoring direction and evidence requirements follow Chapter 3, Section 3.3. Appendix I includes 17 formal category cases, showing only the Indicator and Factor structures for each category. Quantitative specifications and score-band standards for individual Factors are established in the official rules for the corresponding category.

Deliverable Field	Questions Readers Need to Check	Common Misinterpretations
Factor name and definition	What exactly is observed, and whether exclusions are included	Inferring the calculation from the name alone
Dimension and Indicator to which it belongs	Which operating question in the Return Dimension or Risk Dimension the Factor answers	Scoring the same fact repeatedly in multiple places
Quantitative specification	How the numerator, denominator, scope, and outliers are handled	Treating similar formulas as the same specification
Data sources	Which system, platform, document, or other evidence supplies the value	Replacing facts with subjective impressions that cannot be reviewed
Unit	Whether the unit is an amount, ratio, count, number of days, or something else	Assuming equal values have equal meanings
Observation Period	Whether it is monthly, quarterly, annual, rolling, a peak season, or a point in time	Directly comparing results from different periods
Scoring direction and standards for scores from 1 to 3	Which performance falls into which band, and how endpoints are assigned	Judging performance solely by whether a value is high or low
Rule Version	When and to which objects the specification applies	Using current rules to explain historical results

Table 4-3: Fields for Checking Factor Deliverables

4.3 Deliverable Versions, Evidence Records, and Subsequent Calibration

Businesses change; Factor definitions, windows, and thresholds may also require adjustment. For past conclusions to remain explainable, every deliverable should retain industry identification, Rule Version, status, update date, effective date, and scope of evidence.

The Industry Identifier locates the object; the version number distinguishes successive rules for the same object. When content was amended and when it took effect must also be stated separately. A new version may be used going forward, but the old version must remain; otherwise, historical scores cannot be traced to their original basis.

"Developed," "reviewed," "approved," "effective," and "usable for specific scoring" each mean different things. A complete table only proves that the content has been written. Derived fields, fields awaiting confirmation, and fields with restricted use should retain their corresponding statuses so that the basis of their credibility can be examined.

Reviewing a Factor should reveal which industry facts the rule relies on, where its specification comes from, what evidence data collection requires, why intervals are divided as they are, and how revision comments were handled. A version number helps locate the materials; the actual materials determine how deeply they can be examined. A number cannot establish that the entire research and approval process is complete.

Subsequent operations, contract performance, overdue payments, defaults, recoveries, and market use will bring new information and counterexamples. Factors that fail to explain change, unavailable data, thresholds that can be evaded, and overly broad scopes should all enter review. Revisions must explain the changes, reasons, supporting basis, effective date, and effects. Simply changing the score at the output end is insufficient.

Coverage of 500 industries means that the scope of the first round of rules has been established; subsequent calibration will continue to test whether those rules are applicable in practice and need revision. Initial deliverables give the market a specific subject for discussion; new versions turn feedback into more accurate definitions and boundaries. Coverage and calibration are successive steps. Completed coverage need not be repeatedly called into question.

4.4 How to Find Your Industry's Deliverables and Help Improve Them

The 3 cases in the main text and 17 in the appendix, 20 displayed cases in total, show readers how industry rules are organized. The common deliverable portal serves 500 industries, helping readers locate their industry materials and corresponding versions. Case numbers in this edition are only for navigation within the white paper; they are not Industry Identifiers. Correspondence between each case and the 500-industry directory will be listed in the complete deliverable index. These 20 cases were selected by operating scenario. They are not a statistical sample and cannot replace the full set of deliverables.

Level	Content Shown	What Readers Can Do
20 cases throughout the white paper (3 in the main text + 17 in the appendix)	Case numbers, category introductions, Dimensions, Indicators, and Factors	Read Chapter 16, Section 16.1 and the appendix index; propose revisions to specific fields
Common Index of 500 Industries	Search by industry; Industry Identifiers, deliverable versions, and availability status	Find your industry and assess the current extent of access to its deliverables
Specific industry page	Compass, Factor rules, data collection checklist, and version descriptions	Access materials, provide feedback on judgments, subscribe to updates, or book a discussion

Table 4-4: Navigation and Access Arrangements for Deliverables Across 500 Industries

The Industry Compass website is live: <https://microconnect.com/industry-compass>. The QR code and text link lead to the same page. On first entry, one contact method is required, either email or WeChat. The optional identity field offers "operator / investment research / industry expert / other."

Action	What Readers Receive	Follow-Up Connection
View your industry's deliverables online	Released Indicator and Factor structures, rule materials, and version descriptions, for online viewing only	Enter the corresponding industry page after registration and verification; update subscriptions and discussion bookings are chosen separately
Subscribe to updates for your industry	Notifications of subsequent version changes and explanatory updates	Ongoing feedback is received, and revisions to the rules are explained.
Book an Industry Compass discussion	Discuss application to your own operations or investment research	Advance discussion and collaboration around clearly defined application needs

Table 4-5: Three Participation Actions Chosen Separately

Deliverable portal: <https://microconnect.com/industry-compass>. The complete index of 500 industries, deliverable versions and update dates, online viewing content list, and scope of access are governed by current website disclosures.

Chapter 5 How Industry Deliverables Enter Micro Connect New Markets

Deliverables for 500 industries give different businesses their own forms of expression. Another question follows: how can judgments formed separately about convenience stores, manufacturing businesses, and aquaculture businesses be understood together? Their revenue formation, input structures, and cycles differ. Industry-specific language alone makes comparison difficult; a common scale alone can lose the details that determine operations.

The NUMA SHARP Twin Compasses preserve both layers. The Core Compass examines contract structure; the Industry Compass examines the specific business. Each has a Return Dimension and a Risk Dimension. Core Codes and Industry Codes record readings; the NUMA Adjustment Factor supports calibration across categories on the industry side. Retaining scales, facts, periods, and evidence together gives comparison a basis.

The NUMA SHARP Twin Compasses are the twin-compass system in this white paper; NUMA SHARP is the central Indicator for expressing and comparing the risk-return structure of contractual cash flows. Reading specific results still requires distinguishing the Indicator framework, raw scores, and adjusted readings and checking the applicable version.

5.1 The Core Compass Examines Contract Structure; the Industry Compass Examines the Specific Business

Whatever the business, a contract must explain where returns come from, how contractual arrangements, related parties, and performance mechanisms support realization, and what structural risks they introduce. The Core Compass places these recurring cross-industry questions in a uniform framework so that different contracts can be read in a common language.

The Industry Compass goes further into operations: how demand becomes revenue, how inputs generate output, whether results can be sustained, and how costs, execution, dependencies, and external changes affect them. Convenience stores require examination of actual store conditions; other industries require production units, cycles, and data. Under the common questions, facts that truly explain the industry must be preserved.

Clear contract structure does not establish operating quality; good business performance does not replace examination of contract and performance structure. Placing the two compasses alongside each other lets users see both common features and specific differences, preventing conclusions on one side from preempting the other.

Each Compass contains a Return Dimension and a Risk Dimension. The two sides of the Core Compass examine contract structure; the two sides of the Industry Compass examine actual operations. These form four areas of observation that should be understood separately. Contractual arrangements may be relatively sound, yet industry risk remains high; operating performance may be good, yet gaps in contract verification remain. All of this information should be retained.

5.2 Core Codes and Industry Codes: Their Respective Return and Risk Scores

Compasses provide observation frameworks; Codes record the readings formed under those frameworks. Core Codes correspond to the returns and risks of contract structure; Industry Codes correspond to the returns and risks of the specific business. Industry Identifiers serve a different purpose: they locate rules among the 500 industries.

Core Codes bring different contracts onto a common scale, identifying the contract's return support and risk load. They preserve structural questions across industries. Industry Codes are still needed to explain where the specific business stands within its industry.

Industry Codes bring operating differences back into view. The corresponding Factors must explain how convenience store conditions affect demand, efficiency, revenue quality, and risks; in another business, the same

questions may require entirely different evidence. Every reading should carry its definition, unit, period, and Scoring Rules so that its precise meaning is clear.

Both types of Codes present returns and risks separately. Stronger return support does not automatically eliminate risk, and lower risk does not guarantee higher returns. Scores all look like numbers but carry different meanings; they cannot be directly added, offset, or merged into a single direction.

Nor should Core Codes and Industry Codes be compressed into an arbitrary total score with no traceable source. When NUMA SHARP for an individual contract is formed under established rules, the corresponding Compass, Return Dimension or Risk Dimension, Rule Version, and Observation Period must still be retained. Readers can then return to Indicators, Factors, and evidence to understand why the conclusion holds.

5.3 The NUMA Adjustment Factor and Cross-Industry Calibration

The Industry Compass first assesses how a business performs within its own industry. Even if two businesses have similar scores in their respective industries, their industry environments may differ. Simply placing the scores side by side can cause readers to mistake "how it performs within its industry" for "where it stands among all industries." The NUMA Adjustment Factor adds this industry-environment reference. It does not change the Industry Compass's Factors or Scoring Rules. After an individual enterprise completes within-industry scoring, industry business conditions are considered to observe its relative cross-industry position.

The NUMA Adjustment Factor is not limited to a single data source. This edition primarily references relevant data from the National Bureau of Statistics of China. Available data sources will gradually expand to accommodate operating environments and data conditions across industries, countries, and regions. Micro Connect New Markets also welcomes market participants to develop applicable adjustment coefficients based on their own research and use cases, specifying data sources, calculation methods, scope of applicability, and versions. Users may choose among different approaches according to their needs. Cross-industry or cross-region comparisons still require checking specifications and comparability.

This edition primarily uses the business climate index published by the National Bureau of Statistics of China¹. The business climate survey is conducted quarterly; the index aggregates operators' judgments about current business conditions and future trends. Under the National Bureau of Statistics of China's Indicator definitions and compilation methodology, the index ranges from 0 to 200, with 100 as the business climate threshold. It reflects the perceptions and expectations of groups of enterprises regarding their operating environment. It does not directly represent industry revenue, profit, or the return on a specific project.

The 500 detailed industries covered by the Industry Compass in this white paper do not always correspond one-to-one with the statistical industry classification used by the business climate index. Micro Connect New Markets maps detailed industries to relevant underlying industry indices based on their business composition. When a business is affected by multiple underlying industry environments, those indices are combined using preset weights. The mappings and weights express the relevance of different industry environments to that business. Their rationale will be attached, and they will be retained with the calculation version for subsequent review.

The agriculture, forestry, animal husbandry, and fisheries sector and the financial sector have no directly corresponding indices in the business climate index series used in this approach. Micro Connect New Markets therefore uses the preliminary quarterly Gross Domestic Product (GDP) accounting data published by the National Bureau of Statistics of China for statistical estimation. The data separately report the absolute amount of industry value added for the quarter and its year-over-year growth rate for the quarter. Absolute amounts are calculated at current prices and reflect the scale of industry value added in that quarter; year-over-year growth rates are calculated at constant prices and reflect growth in industry value added excluding price changes. This

National Bureau of Statistics of China official explanation of the "business climate index":
<https://www.stats.gov.cn/zs/tjws/zytjzbqs/qyjqzs/1>

approach uses quarterly values and does not substitute cumulative half-year or annual values for single-quarter data.

To estimate the NUMA Adjustment Factor for the agriculture, forestry, animal husbandry, and fisheries sector and the financial sector, samples are first drawn from industries with both business climate indices and quarterly industry value-added data. Data are paired for the same industry and quarter. The model uses published business climate indices as the dependent variable and the quarterly year-over-year growth rate and quarterly absolute amount of value added as explanatory variables. This establishes a linear regression relationship across industries and quarters. Concurrent value-added data for the agriculture, forestry, animal husbandry, and fisheries sector and the financial sector are then entered to estimate their respective business climate indices. Testing examines both the model's performance in subsequent quarters and its ability to estimate indices for known industries temporarily held out of the sample. Test results inform model adjustment and assessment of the stability of industry rankings. The resulting values are model estimates and should be distinguished from indices published directly by the National Bureau of Statistics of China.

Each calculation round uses only the latest quarter for which complete data have been obtained. Published business climate indices, GDP industry value added used to estimate missing items, and subsequent mapping, combination, and ranking must all correspond to the same quarter. More recent business climate indices cannot be combined with GDP data from the preceding quarter. This edition uses data for the second quarter of 2026. When data for a new quarter can be completely aligned, calculation and ranking will be updated together.

After obtaining industry business climate indices for the same quarter, Micro Connect New Markets forms adjustment coefficients based on the relative positions of the industry indices. The coefficient uses 1 as the neutral baseline; industries ranked higher or lower are adjusted correspondingly under uniform parameters. The coefficient expresses the relative position of the industry environment in that period, not a precise scale of operating differences between adjacent industries. When indices are very close, overinterpreting small ranking changes is not recommended.

This edition presents Industry Compass rule deliverables and the design methodology for the NUMA Adjustment Factor. Actual measurement and scoring of individual enterprises and Cross-Industry Calibration remain for subsequent application. In actual use, each coefficient should be traceable to its data, observation period, industry mappings and weights, and model and parameter versions. The NUMA Adjustment Factor adds an environmental reference to judgments on the industry side. Contract structure remains subject to the Core Compass; the portfolio-level risk retention coefficient, denoted by DR, has its own independent purpose. The industry environment provides a reference; the specific business must ultimately be assessed using enterprise operations, contractual arrangements, cash flows, and evidence of performance.

5.4 A Reading Path from Within-Industry Judgment to Cross-Industry Comparison

When reading the Twin Compasses, first examine the return support and risk load in contract structure, then how Factors explain the specific business. Check the materials and periods used on each side. First understand the sources, then examine subsequent ranking comparisons.

Next, check industry-side calibration: whether the corresponding NUMA Adjustment Factor is used, which objects and periods the benchmark covers, and which Rule Version applies to the result. After calibration, users should still be able to return to the original Factors, evidence, and benchmarks to distinguish whether differences arise from operating facts, industry conditions, or changes in periods and versions.

Before comparison, check the objects, industries, periods, versions, data scopes, and calibration conditions one by one. Monthly and annual data carry different facts; new and old rules may change definitions and boundaries. Identical numerical formats only mean that both are written as scores. Comparability of their conditions requires separate examination.

Missing materials, unclear industry classification, core operations spanning industries, and conflicting sources should all remain explicit limitations. Low scores, not yet scored, and insufficient evidence are different statuses. Mixing businesses, periods, or specifications merely to complete readings makes results harder to explain.

The reading order is therefore clear: contracts first, then businesses; examine returns and risks separately; check industry-side calibration; then compare when periods, versions, and evidence conditions match. Each step preserves the origin of the readings.

5.5 Conditions of Use: Industry Deliverables, Project Verification, and Ongoing Observation

Deliverables for 500 industries provide a foundation for understanding businesses and Scoring Specifications. Cases help readers learn how to read; specific industries must still be found through the index. Coverage has been established; the performance of Factors, scoring, and calibration methods in different environments still requires continuous validation.

For a specific project, users must still check the enterprise's business model, operating stage, data quality, and contractual arrangements. Where multiple businesses coexist, the part covered by the current reading must be explained. Similar industry names help with searching; they cannot replace verification of materials or judgments of applicability.

The Twin Compasses help organize information, identify return support and risk load, and indicate what should be verified next. Parties, contracts, cash flows, and changes in ongoing operations still require specific examination. The compasses do not constitute commitments to returns, legal opinions, compliance approvals, or automatic trading conclusions. Users make decisions based on complete materials, applicable rules, and their own responsibilities.

Operations, external benchmarks, and methods all change. New results must specify new conditions; historical results retain their earlier basis so that differences over time can be explained. Ongoing observation allows judgments to be reviewed and corrected as facts develop, rather than treating one score as a permanent conclusion.

The Core Compass preserves the contract scale; the Industry Compass preserves business differences. Core Codes and Industry Codes separately record returns and risks, and the NUMA Adjustment Factor supports industry-side comparison. Delivering objects, periods, versions, evidence, and applicable conditions together makes the deliverables for 500 industries easier to understand, inspect, and use continuously.

Part III HAT over MASS Theory

The Design Philosophy for Manufacturing and Maintaining Registration and Measurement Instruments

HAT Boundaries. MASS Collaboration. Twelve Sufficiency Conditions.

Part II explained what we are building. This part asks what organization manufactures and maintains these instruments. The answer is the AI factory of HAT over MASS: it puts full model capability into manageable, inspectable Microtasks and brings large numbers of tasks together to perform professional work.

AI in real financial work encounters two sets of tensions at once. Machine processing capacity can expand, but responsibility must rest with specific people and organizations.

Human judgment and governance responsibility cannot be absent, yet human capacity cannot keep pace with every machine action. The following chapters explain why the factory adopts HAT over MASS and how boundaries, collaboration, and inspection work together.

Here, the cost constraint in Part I becomes an organizational design question.

Chapter 6 Two Sets of Tensions Between AI Capability and Human Responsibility

6.1 The Dilemma of Letting AI Act and Watching AI Closely

Once AI enters real business, the question quickly moves from "Can it do this?" to "How far do we dare let it act on its own?" In terms of capability alone, it can process large amounts of information simultaneously, advance many tasks in parallel, and continue extending its reach as models, data, and tools are updated. It is not subject to the same constraints of fixed working hours and a single position's capacity. Naturally, we want to give it more work and make tasks previously prevented by speed and cost possible for the first time.

The faster capability advances, the less room there is for ambiguous responsibility. AI can generate judgments, make recommendations, and execute steps, but one correct result does not automatically make it a legally or organizationally accountable party. Who defines what it may do, sets the prohibitions, handles exceptions and disputes, and bears responsibility for results? These questions must ultimately resolve to specific people and organizations. "AI has unlimited capability but limited responsibility" describes this asymmetry: processing capacity and execution scale can expand rapidly, but ultimate responsibility does not automatically arise with more calls.

The human side is the reverse. Judgment, attention, and working time have limits and cannot multiply alongside vast numbers of machine tasks. Yet people cannot withdraw from authorization, signing, supervision, and accountability merely because there are too many tasks. "Human capability is limited, but responsibility is unlimited" does not require an individual to assume responsibility without boundaries. It reminds us that an organization choosing to use AI must fully assign responsibility rather than leave unexplained, uninspected, and unassigned areas for the system to absorb on its own.

One common approach is "AI leads, people support": let AI produce at scale, then have people handle a few exceptions or provide final signatures. Scale emerges quickly, but all pressure converges on the final human gate. AI can generate thousands of results in parallel; human review and signing capacity do not increase proportionately. If reviewers cannot return to inputs, rules, and processes in time, signatures can become formalities. Repeated correct performance by AI can build experience in its use, but probabilistic correctness cannot replace process control, much less a clear accountable party.

The other approach is "people lead, AI supports": AI organizes materials, performs preliminary analysis, and provides prompts; people judge key steps individually. This arrangement remains necessary for major decisions, matters heavily dependent on discretion, or matters difficult to express in rules in advance, and makes accountability more direct. But when tasks become frequent, numerous, and continuous, human review blocks the entire process. AI has entered, yet the system still operates at human speed. Expandable machine capability never becomes the organization's actual processing capacity.

Each approach safeguards one side but loses the other. One runs short of people; the other hesitates to use AI. Letting AI act raises concern about uncontrolled responsibility; placing people before every step constrains capability to human speed. From the outset, the answer we seek must release capability and assign responsibility within the same structure. Sacrificing either side merely postpones the problem.

6.2 The Financial Industry's Essential Questions: Compliance, Investor Protection, and Risk Control

Many industries face these tensions; finance especially cannot avoid them. Financial activity connects funds, operations, and risk. A result affects more than process efficiency: it affects participants' rights and interests and market trust. Technology can change, but compliance, investor protection, and risk control cannot be added after the business has run, nor can their requirements be reduced because processing volume increases.

AI's real attraction in finance is precisely its ability to handle frequent tasks at scale. Materials need organizing, information needs checking, operations need ongoing observation, and existing judgments need repeated testing against new facts. If a few professionals perform all of this work item by item, service coverage and update speed will inevitably be limited by human capacity. Giving divisible, repeatable work to AI may bring tasks previously excluded by excessive costs within reach.

Introducing AI into finance must never amount simply to replacing a manual button with an automatic one. A model's ability to do something does not mean the organization has authorized it to do so. Every execution must make clear where inputs come from, which Rule Version applies, what inspection the result undergoes, and when execution must stop and return to people. Without these connections, faster automation is more likely to replicate errors, biases, and omissions in batches.

The release of capability and control of responsibility must therefore be designed together within one structure. Tasks with clear inputs, outputs, and qualification standards may be executed continuously by AI within established boundaries. Each output must still enter the corresponding inspection stage so that errors can be found, returned, and corrected. When a new judgment is needed, a rule gap is encountered, or authorization is exceeded, the system must refer the issue to people with the corresponding responsibilities. Human value lies not in following the machine and redoing every task, but in establishing rules, defining authority, handling exceptions, supervising operations, and assuming ultimate responsibility.

"Review" consequently has a stricter meaning. Nodding after seeing a result does not count as review. Responsibility is established by being able to determine which materials were used, which steps were followed, which Rule Version applied, and who handled exceptions. Automatic execution, inspection of every item, and explicit accountability must form a single chain so that AI's scale does not come at the expense of compliance, investor protection, or risk control. The financial industry's real question is how efficiency and responsibility can expand together.

6.3 How HAT over MASS Addresses Both Sets of Tensions

HAT over MASS seeks to address both sets of tensions between AI capability and human responsibility and the cost constraint of professional services in Part I. Rule boundaries connect AI capability to human responsibility; Microtask collaboration brings professional capability to vast numbers of small tasks. Together, they create an opportunity for "powerful capabilities put to small-scale use" to become the organizational foundation for "large pools of capital put to small-scale use."

"Powerful capabilities put to small-scale use" means applying powerful model capabilities to specific tasks with clear scopes. A model may have full capabilities for understanding, analysis, and execution, but the objective, scope, and authority of every task must be defined in advance. We divide complex work into independently deliverable, inspectable units and connect their collaboration according to dependencies. Outputs enter the next stage after inspection. Problems outside the rules or authorization are referred to people with the authority to judge them.

This is the origin of HAT over MASS. The "Governance Hat (HAT)" is the rule boundary established and maintained by people; the "organization (MASS)" is the collaboration of large numbers of Agents within those boundaries. An individual Agent takes one Microtask: a focused objective, clear inputs and outputs, a limited duration, an inspectable result, and necessary process records. "Micro" constrains the scope and authority of a single task; it does not reduce model capability.

We govern the individual so that each execution can be examined clearly. The larger the task and the vaguer its boundaries, the harder it is to distinguish problems in materials, rules, execution, or Acceptance Review. Decomposition into traceable, inspectable tasks locates exceptions at the appropriate stage. Matters outside the rules or requiring discretion stop automatic progression and are escalated to a human. Constraining individuals prevents capability from reaching areas where no responsibility can support it; it does not render AI ineffective.

The collective is unleashed through the parallel execution and collaboration of these manageable units. Once boundaries, deliverables, and inspection conditions are clear, the organization can confidently increase task

volume, with different Agents producing, checking, and collaborating. Vast numbers alone do not create reliability; quantity does not automatically become correctness. Scale is viable because each unit can be inspected, divisions of labor provide checks and balances, and problems can be traced back to the corresponding rules and responsible people.

Within this structure, people and AI shoulder different duties. AI executes large numbers of divisible tasks within the rules. People establish and amend rules, determine authorization boundaries, handle exceptions and disputes, supervise system performance, and bear ultimate responsibility that cannot be assigned to machines. People need not perform every action themselves, but responsibility does not disappear. It rests where it matters more: rules have owners, boundaries have guardians, exceptions have handlers, and results can withstand review.

Figure 6-1 brings together the two sets of tensions and the responses of HAT and MASS.




Figure 6-1: Two Sets of Tensions and the Response of HAT over MASS (repeated from Figure 0-1 in the Executive Summary)

This answer has strict prerequisites. Reliability depends on whether tasks can be divided for independent inspection, rules can cover actual capabilities, records can reconstruct processes, and different foundations and positions can provide real inspection and checks and balances. Scale may further amplify rule blind spots, data contamination, Common-Source Bias, and propagation of erroneous judgments. HAT over MASS provides organizational principles for boundaries, collaboration, and responsibility. It has never endorsed any automated result in advance.

Three concepts now connect: the ambition is "large pools of capital put to small-scale use"; the obstacle is "top talent cannot economically serve small tasks"; the answer is "powerful capabilities put to small-scale use." We place full model capability in manageable, inspectable Microtasks, govern the individual, and unleash the collective, making every instrument output traceable to rules, evidence, and responsibility. The next chapters explain how Governance Hats keep people in control of rules, how MASS builds collaboration at scale through Microtasks, and how collaboration and inspection determine whether this factory is reliable.

Chapter 7 The HAT Governance Hat: People Control the Rules; AI Executes Within Them

If every AI step awaits human approval, vast numbers of frequent tasks can never get moving. If only automatic execution matters, responsibility falls into a gap with no identifiable accountable party. The HAT Governance Hat resolves this dilemma. Rather than asking whether people or AI do more, it first separates two things: people control rules and the rights and responsibilities for changing them; AI executes continuously within the areas those rules already cover.

A Governance Hat brings together rules, authorization boundaries, and escalation arrangements. It tells an individual AI what it may do, what it may not touch, and when it must return an issue to people. Clear individual boundaries remove the need for human approval of every action within them. Keeping rules in human hands ensures that automatic execution does not become a transfer of responsibility.

7.1 The Meaning of HAT: Ultimate Human Responsibility and the Release of AI's Potential

HAT has three components: Human-Commanded, AI-Operated, and Trust. It is fundamentally a trust mechanism that is human-controlled and AI-operated. It meets two requirements simultaneously: human rules, authorization, and ultimate responsibility must be implemented, and AI has sufficient room to execute within clear boundaries. Trust is established jointly through accountability arrangements, process records, and deliverable inspection. Emphasizing control alone may limit effective AI use; emphasizing capability alone may allow execution beyond what the organization can supervise and take responsibility for.

Human control means that people establish, amend, and revoke rules. AI operation means that once rules, authority, and task scope are clear, AI may autonomously and continuously complete work within the rules. Human work therefore moves to rules, authorization, and exception handling, and AI capabilities unfold in inspectable tasks.

Trust in HAT never rests on a claim that "AI will get it right," nor treats one correct delivery as long-term reliability. It rests on three inspectable conditions: responsibility has an identifiable party, the process is visible and reconstructable, and results can be reviewed under the same specification. Trust thus becomes a structural judgment of responsibility, process, and results rather than merely a feeling about capability.

7.2 Rules Before Capability: Permitted Actions, Prohibitions, and Escalation Conditions

The stronger the capability, the more necessary it is to define the scope of action first. Before introducing an Agent, we answer three questions: what it is allowed to do, what is explicitly prohibited, and when it must stop and escalate to a human. Permitted actions provide room to act, prohibitions establish red lines, and escalation conditions handle situations that rules do not yet cover.

A principle that is sound in general terms cannot yet be used directly as an operating rule. Rules must at least be executable and inspectable and create actual separation in the system. Executable means that objects, tasks, and applicable conditions are clear enough that Agents need not fill key gaps with ad hoc discretion. Inspectable means results can be checked item by item against rules, not passed because they "look reasonable." Actual separation means that unauthorized data, tools, actions, and processes are genuinely restricted during operation; a prohibition written on paper is insufficient.

Rules are not written once to remain effective forever. When models, tools, data or task methods change, the organization must reassess whether the existing Governance Hat still covers the new capabilities. The emergence of a new capability does not mean that it suits the current task or has approval to enter formal

operation. Its applicability, reliability, and permitted actions, prohibitions, and escalation conditions all require renewed confirmation.

For high-consequence actions such as funds transfers, contract signing, and external commitments, and operations subject to licensing, qualification, or security constraints, a preset suspension condition must trigger immediate suspension of the affected automatic execution. The trigger must be recorded, evidence retained, and the matter referred to authorized personnel. After remediation and reinspection, an authorized party must approve resumption. Resumption cannot rely solely on "the error has been corrected"; reinspection findings and approval records must also be retained.

7.3 A Named Human Behind Every Governance Hat: Accountability Assignment and Rule Establishment, Amendment, and Revocation

Rules can govern only if they are clear and the people responsible for them can be identified. Assigning each Agent to a named human lets the organization specify its applicable rules, authorization scope, and who handles rule changes. Naming the person goes beyond defining organizational and departmental responsibilities: it identifies the specific people responsible for rules, authorization, and exception handling, preventing responsibility from ending at a generic system or department name.

Accountability assignment is never complete merely by attaching a name. Identifiable responsibility and records must explain why rules were established, what later changed, which tasks a new version covers, and when use should stop. If rules have changed but no one can identify who changed them when or what the old and new versions cover, subsequent execution cannot be reviewed. Inspection of a Governance Hat ultimately returns to three questions: whether rules are clear, changes are recorded, and execution can be replayed.

The responsible person handles rules, authorization, and exceptions, rather than redoing every AI task. If Named Human Accountability Assignment ends in manual review of every item by the accountable person, responsibility has a name but processing capacity remains confined by human capacity. Effective assignment ensures that someone is responsible for rule decisions, maintains authorization boundaries, and handles matters outside the rules, with execution within the rules remaining with AI.

7.4 Governance Hats Govern Boundaries, Not Actions

Governance Hats govern boundaries, not every action within them. They specify permitted and prohibited scopes and the consequences of crossing boundaries, allowing Agents to execute under the rules within their authorization. When a task exceeds authorization, enters a prohibited area, or encounters a judgment not covered by rules, the system restricts action or escalates to a human under preset conditions. Agents cannot expand their own authority or rewrite rules, much less exercise human discretion in people's place.

"Not governing actions" does not mean "not seeing actions." Execution within the rules must still retain process information and undergo inspection; input errors, execution deviations, and environmental changes may still render results unqualified. The Governance Hat changes how people intervene. For routine tasks covered by rules and with clear authorization, people need not approve every item; they govern through rules, authorization, supervision, and escalation handling. Actions requiring prior approval under law or organizational policy still follow the corresponding requirements.

The object of direct human management thereby narrows from vast numbers of individual tasks to rules that can be maintained, reviewed, and revised. Tasks may continue increasing, but their underlying rules must always be identifiable, visible, and linked to accountable parties. How far scaling can proceed depends not only on increases in execution capability, but also on whether the organization can maintain rules, identify boundary changes, and handle matters beyond boundaries.

7.5 Human Responsibility for Governance, Supervision, and Discretion

When AI takes over repetitive execution within the rules, human work does not disappear; it moves upward from individual execution and review to rule governance, authorization supervision, and accountability. People must

decide which tasks can be governed by rules, which risks must be isolated, when automatic execution must stop, and which judgments still require human experience, authority, and professional responsibility. The system's ability to act does not make these questions dispensable.

Governance cannot end on the day rules are released. After rules enter real tasks, people must continuously understand how they operate. This "operational awareness" cannot be only a general impression; it must rest on specific information: which tasks repeatedly trigger escalation, which results are often returned, which exceptions do not fit existing rules, and which persistent biases review and backtesting uncover. Operating information must continually return to the Rules Layer so that Governance Hats can be revised as environments, capabilities, and tasks change.

Discretion beyond the rules remains a human responsibility. When facts are insufficient and cannot be supplemented through established procedures, objectives conflict in ways that rules cannot resolve, or judgments require authority beyond existing authorization, the system should suspend affected automatic execution and submit established facts, execution records, and rule gaps to authorized people. If a human resolution is repeatable, it should also become a candidate basis for the next round of rule revisions; operations cannot rely indefinitely on ad hoc approvals.

HAT is a continuously operating relationship. People set boundaries for AI that they can take responsibility for; AI exercises its capabilities within them; operating results return to human supervision and rule revision. People need not perform every action personally, but retain control of rules, responsibility, and discretion. Once boundaries are firmly established, Microtasks can safely increase, and vast numbers of Agents can share a common order.

For operations that may have major effects, such as funds transfers, contract signing, and external commitments, it should be specified in advance who may make professional judgments, stop operations, and approve resumption. Once a suspension condition is triggered, operations must stop, evidence must be retained, and the matter must be referred to authorized people under the procedure in Section 7.2. Resumption requires approval after remediation and reinspection. Accountability for an AI Organizational Position or participation in a task does not itself confer these powers.

Chapter 8 MASS: How Powerful Capabilities Enter Microtasks

HAT first establishes boundaries: what may be done, what may not be done, and when a matter must return to people. The question then moves from governance to organization: how to let AI execute at scale within these boundaries without allowing scale itself to create new uncertainty.

MASS answers by organizing full model capability into Microtasks. An Agent may have strong capabilities for understanding, analysis, and execution, but in a specific run it takes only a task of limited scope and relatively short duration that can be fully inspected. Large numbers of Microtasks can run concurrently, connect, and check one another, turning individual constraints into collective execution capacity.

"Micro" and "mass" are therefore mutually enabling. Micro constrains individual tasks so that they can be specified, inspected, and redone; mass organizes those units to cover work previously beyond reach because of human working-time limits. MASS does not simply increase Agent counts. Task granularity, authority boundaries, validation conditions, and foundation configuration make execution at scale manageable.

8.1 The Meaning of MASS: Micro Agents, Simple and Short Tasks, and Collaboration at Scale

MASS: Micro Agents; Simple & Short Tasks.

Every Agent in the organization is built on the capabilities of a frontier large model, but full model capability does not mean unlimited task authority. It takes one Microtask: simple, short, with boundaries clear enough for its result to be inspected and granularity fine enough for key steps to be traced. Assigning powerful AI only one narrowly scoped task makes the task organizable and inspectable, and the parallel collaboration of thousands of Microtasks builds capability at scale. Small tasks let people manage vast numbers; vast numbers let AI realize its potential.

"Simple" requires a single task to carry only one clear judgment or deliverable; "Short" requires completion in one run with clear boundaries, without relying on untraceable, long-running ad hoc judgments. There is no uniform number of minutes, but inputs, outputs, completion conditions, and inspection standards must all be specified in advance.

Breaking complex objectives into clearly bounded units gives large-scale concurrency a real foundation. Different Agents can process independent tasks simultaneously or connect sequentially along explicit dependencies. Only identifiable inputs, outputs, and inspection conditions allow concurrent results to be compared, combined, or returned. If tasks themselves are unclear, adding Agents merely produces more outputs that cannot be judged at the same time.

The "mass" in MASS is the ability to organize thousands of Microtasks into a collaborative network: every unit has boundaries, units have relationships, and results undergo inspection. It does not mean any system can run unlimited concurrent tasks at any time. It describes a direction for expansion: provided unit design and inspection conditions can support it, task volume is no longer confined by the speed at which one human team executes items individually.

8.2 Small Tasks, Strong Agents: Full Model Capability and Limited Task Authority

"Small tasks, strong Agents" first separates capability from authorization. The "soldier" is MASS's metaphor for an Agent; it does not mean the Agent possesses only one low-level skill. An Agent may instead be built on full, advanced model capability, but in one run may use only the authority, materials, and tools needed for the current task and is responsible for a clear deliverable.

Boundaries are a prerequisite for organizing powerful capabilities. If a model's ability to do more leads to material searches, judgments, calculations, revisions, and even approvals all being packed into one run, the result may look complete, but errors become difficult to locate and responsibilities indistinguishable. The breadth of model capability cannot replace clarity of task boundaries.

Small tasks give errors an identifiable scope. Handling one class of problem at a time makes errors easier to isolate at the relevant stage. Sufficiently narrow boundaries let inspectors determine whether a result passes against standards specified in advance. Sufficiently short duration makes complete retention and reinspection of inputs, actions, and outputs possible. These conditions cannot guarantee that Agents are always correct, but they can prevent errors from entering subsequent work inside an indivisible "complete answer."

This is what "powerful capabilities put to small-scale use" means at the task level. "Powerful capabilities" are strong model capabilities callable at scale; "small-scale use" means focusing each run on one matter that can be judged, inspected, and redone. Rather than asking one Agent to play every role simultaneously, we prefer to put powerful capabilities into clear tasks and complete complex objectives through collaboration. Capability has not been reduced; it has been converted into manageable execution units.

8.3 Three Sufficiency Conditions for Task Decomposition: Sufficiently Small, Narrow, and Granular, with Business Meaning Intact

Microtasks cannot be measured solely by word count, duration, or number of steps. Admission to MASS requires three conditions simultaneously: sufficiently small, sufficiently narrow, and sufficiently granular. They answer three questions respectively: how much work is undertaken at once, how much ad hoc discretion remains during execution, and whether outputs can be independently inspected and replayed.

Sufficiently small means doing only one thing at a time. It may involve multiple steps but must have only one clearly identifiable delivery objective. Sources of inputs, permitted resources, required outputs, and qualification standards must all be specified before execution. A task requiring fact collection, rule establishment, risk judgment, and result approval simultaneously is not yet small enough for clear Acceptance Review.

Sufficiently narrow means the task does not depend on discretion arbitrarily added during execution. Agents may select methods within a predetermined scope, but may not change objectives, expand authority, or make unauthorized judgments in place of people. If facts are insufficient, rules conflict, results exceed boundaries, or the task requires new discretion, execution stops within known boundaries and escalates to a human. Escalation does not indicate failed task design; it proves that authority has not silently expanded.

Sufficiently granular means each output can be independently inspected and key steps replayed. Granularity is not merely dividing an article by paragraphs, nor may it split a judgment until its meaning is lost. If understanding a unit requires the executor to orally supply a large amount of unrecorded context, it is not an independent Microtask, however short its form. Sound decomposition preserves the full meaning of each unit and allows it to be individually returned, corrected, or redone.

All three conditions—small, narrow, and granular—must hold. A small task with unlimited discretion is not sufficiently narrow. A narrowly bounded task whose output cannot be independently inspected is not sufficiently granular. Work split into many pieces, each still carrying multiple delivery objectives, is not sufficiently small. Only when all three conditions hold does a Microtask become an organizational unit that can be replicated, assigned, and inspected at scale.

Condition	Specific Meaning	Inspection Points
Sufficiently small	Only one clear delivery objective at a time	Whether inputs, outputs, and qualification standards can be specified in advance
Sufficiently narrow	Execution only within authorized boundaries	Whether it relies on ad hoc discretion or self-expansion of authority
Sufficiently granular	Results can be independently inspected; key steps can be replayed	Whether it can be individually returned, located, and redone

Table 8-1: Three Sufficiency Conditions for Task Decomposition

8.4 Establish Verifiability First, Scale Next: Independent Acceptance Review, Error Location, and Local Rework

Once a complex objective is divided into Microtasks, their number often grows rapidly. Work originally completed sequentially by a few people may expand into many tasks for obtaining information, making independent judgments, cross-checking, and combining results, assigned to different Agents and creating larger staffing requirements. But greater staffing only shows that tasks have been split; it does not prove the decomposition is sound.

Without clear standards for an individual task, greater quantity means more unknown outputs. Without independent validation of an individual task, a combined result cannot reveal which step failed. If an output depends on large amounts of unrecorded context supplied ad hoc by its executor, another Agent cannot reliably take over inspection or continue execution. Quantity cannot fill gaps in boundaries and evidence.

Before increasing volume, the organization must answer basic questions clearly: whether every unit has explicit inputs and outputs; whether qualification standards can be stated before execution; whether the process can be replayed; whether failed work can be returned only to the relevant unit; and whether a need for discretion triggers stopping and Escalation to an Authorized Human. Only with assessable answers can an individual task support replication and concurrency.

The cost of creating an Agent itself is relatively limited. Organizational costs that grow with scale also include model calls, data, tools, scheduling, inspection, and exception handling. No universal number needs to be assumed in advance. Expandability depends first on whether each execution unit can be examined, judged, and returned, rather than how many Agent names the system can create. The more unvalidatable units are replicated, the harder it is to contain risk.

MASS follows a simple sequence: achieve verifiability first, then increase task volume. Establish that a task can be clearly delivered and inspected before replicating it more broadly. Execution scale, inspection capacity, rule coverage, and accountability capacity must grow together. Increasing tasks when inspection and exception handling remain unchanged does not constitute manageable collaboration at scale.

8.5 Three Sufficiency Conditions for Supporting Foundations: Diversity, Cross-Checking, and Cost-Effectiveness

Clear task boundaries alone are insufficient to support collaboration at scale. If all tasks depend on the same model sources, training biases, and problem-solving paths, individual tasks may look independent, yet large-scale outputs may still collectively fail at the same point. MASS therefore requires not only sufficiently small tasks but also model foundations that are sufficiently diverse, cross-checked, and cost-effective.

Sufficiently diverse means not placing all tasks and inspections on the same foundation. Different model foundations and suppliers within their applicable scopes can reduce repeated replication of Common-Source Bias. Diversity seeks meaningful differences, not the largest possible number of models. Models from different brands may still share similar data, search results, or toolchains; different suppliers do not inherently mean complete independence. Diversity can reduce Common-Source Bias, but cannot establish that bias has disappeared.

Sufficiently cross-checked means having different foundations independently execute the same task where parallel inspection is appropriate, bringing both agreement and disagreement into view. Its value lies not in a guarantee that majority voting among models is correct, but in turning previously hidden differences in judgment into inspectable signals. When substantive disagreement occurs, Agents may not erase differences or select favorable answers on their own. Preset conditions should automatically trigger escalated inspection, followed by further review or human handling under the rules.

"Triggering escalation" is first an organizational principle here. For a specific task, the corresponding rules must state which differences constitute substantive disagreement, who handles them, how they

are recorded, and whether subsequent processes are suspended. Without disagreement standards, producing multiple answers simultaneously remains parallel production; it has not yet become real cross-checking.

Sufficiently cost-effective means configuring models of different capabilities and costs according to task difficulty, risk, and inspection needs. Calling the most expensive model for every task makes collaboration at scale difficult to sustain. Selecting underpowered models based only on price may raise failure, rework, and inspection costs. Real cost-effectiveness matches capability, cost, and scale while meeting task and inspection conditions.

Diversity, cross-checking, and cost-effectiveness must be considered together. Diversity without cross-checking leaves different foundations working separately, without observable checks and balances. Cross-checking without diversity may reproduce the same bias over repeated executions. Even with both, a lack of cost-effectiveness makes it hard to extend validation to vast numbers of tasks. Together, the three conditions safeguard two things: foundations must not become a single source of failure, and resource allocation must not undermine execution at scale.

MASS ultimately assesses not how many "soldiers" the organization contains, but whether each unit is constrained to an inspectable scope and its supporting foundations are sufficiently diverse, capable of cross-checking, and financially sustainable. The three task conditions establish the units themselves; the three foundation conditions keep those units comparable and sustainable at scale. With both sets in place, small tasks can truly support vast numbers, and releasing capability need not come at the expense of management control.

Condition	Specific Meaning	Inspection Points
Sufficiently diverse	Retain different sources of models and capabilities	Whether sources and methods genuinely differ
Sufficiently cross-checked	Form judgments independently, then compare supporting bases and disagreements	Whether inspections repeat the same assumptions and disagreements are retained
Sufficiently cost-effective	Match capabilities to task difficulty and cost	Consider the costs of calls, data, inspection, and rework together

Table 8-2: Three Sufficiency Conditions for Supporting Foundations

Chapter 9 How Microtasks Form Reliable Collaboration

Making tasks small addresses how to govern an individual Agent, but not how thousands of tasks run together. As scale grows, three questions emerge immediately: who assigns work, who determines whether results pass, and who handles areas beyond the rules. Without clear relationships, more Agents may accelerate handoffs, conflicts, and error propagation. Yet returning every step to human approval constrains collaboration at scale to human speed again.

HAT over MASS therefore needs more than task decomposition; it needs an organizational method capable of supporting scale. Rule governance and task execution follow separate channels. Scheduling, Acceptance Review, and backtesting each retain their responsibilities, and every Agent is assigned to a named human. The purpose is clear: enable large numbers of Agents to collaborate within boundaries and ensure every output is inspected, can be reviewed, and can hold up when tested against subsequent facts, rather than adding another layer of organizational structure.

9.1 Governance-to-Agent Control and Agent-to-Agent Collaboration: Two Channels for Rule Governance and Task Collaboration

The organizational principle of HAT over MASS can be condensed into one sentence: people govern Governance Hats; AI manages tasks. Governance Hats first make boundaries clear: what may be done, what may not be done, and when execution must stop or escalate. People establish, amend, and revoke the rules; discretion beyond the rules also returns to people. Agents may bring strong capabilities to the work, but may act only within the boundaries set by their Governance Hats.

This creates two channels. Governance-to-Agent Control governs rules and boundaries: whether a task may enter, which materials and capabilities it may call, and which lines trigger restriction or escalation. Agent-to-Agent Collaboration governs collaboration within boundaries: who accepts tasks, who receives results, and how Acceptance Review and subsequent review occur. One channel confines capability to an inspectable scope; the other removes the need for ad hoc human instructions at every routine step.

Why must there be two channels? Once people relay and approve every task, they become the narrowest bottleneck in the entire collaborative network, preventing release of Agents' concurrent capabilities. Yet task collaboration detached from Governance Hats allows rule gaps, erroneous inputs, and unauthorized actions to run along the chain together. We seek free collaboration within boundaries, with restriction, return, or escalation when a boundary is reached.

People leave the role of executing every item; they cannot leave their governance responsibility. Someone must always answer whether rules are clear, processes retain evidence, exceptions are handled, and who decides when discretion is needed. Collaboration among Agents must also rest on clear task objectives, inputs and outputs, and qualification conditions. People govern Governance Hats and AI manages tasks because rules reconnect responsibility with execution.

Figure 9-1 shows how rule boundaries, Microtask collaboration, and independent inspection connect into a delivery chain.




Figure 9-1: Microtask Collaboration and Delivery Within Rule Boundaries

9.2 From Decomposition to Assembly: Task Dependencies, Handoff Interfaces, and Deliverable Integration

After tasks are split apart, they must connect again. One unit's result is often another unit's input. Until the preceding step is confirmed, the next cannot treat it as an established fact. Scheduling first identifies these dependencies: which tasks may run in parallel, which must await preceding results, and which need rescheduling when conditions change. Concurrency is determined by dependencies, not simply by the number of available Agents.

Handoff interfaces must clearly specify what the next recipient needs. A delivery should contain more than a conclusion: it should also carry input bases, applicable rules, output specifications, versions, inspection status, and unresolved matters. The next stage uses these to confirm whether it can accept the delivery, without relying on ad hoc context from the preceding executor. Clear interfaces prevent division of labor from cutting a complete problem into mutually unintelligible fragments.

Assembly has its own Acceptance Review requirements. Even if every component qualifies, it is still necessary to inspect whether their definitions, units, periods, and versions align, whether they contradict one another, and whether together they cover the original delivery objective. A complete result is not established automatically by joining several qualified answers. The relationships after combination also require independent inspection.

When an interface fails to connect, first locate the affected stages, then determine which subsequent results need rework. Scheduling rearranges dependencies, execution positions correct the corresponding outputs, and Acceptance Review positions examine revisions and their downstream effects. Local rework must be possible, with its effects visible, so that Microtask flexibility does not become repeated restarts of the entire chain.

9.3 From Disagreement to Resolution: Clarification, Review, Return, and Escalation to an Authorized Human

When different Agents give different answers, first determine whether they are answering the same question. Differences in input materials, concept definitions, Observation Periods, or applicable rules can create apparent conflicts. Putting these bases together for clarification distinguishes differences in specifications or evidence from differences in judgment itself.

If substantive disagreement remains after clarification, review follows. Reviewers examine facts, calculations, and reasoning under established rules, retaining each party's opinions and supporting bases. Model counts, expressed confidence, and delivery speed cannot replace this step. If evidence is insufficient, unresolved questions should remain; disagreement cannot be erased merely to assemble one answer.

A task is returned to the corresponding work unit only once the need for revision is confirmed. The return must state the problem, supporting basis, revision scope, and conditions for renewed Acceptance Review, so the executor knows what needs correction. An error at one stage does not invalidate all qualified outputs. Changes to prerequisites or common specifications also require checking their effects on downstream tasks.

When a task goes beyond existing rules, requires expanded authorization, or introduces new discretion, the Agent should pause automated execution of the affected part and escalate the matter. It passes facts, existing rules, disagreements, and possible effects to an authorized named human, who decides whether to continue, revise, or terminate. The resolution returns to task and rule records so that routine collaboration can resume within clear boundaries.

9.4 Qualified Components and Complete Delivery: Separation of Authority and Checks and Balances Among Scheduling, Acceptance Review, and Backtesting

Once collaboration scales, assigning work, admitting results, and ongoing inspection cannot be merged into one kind of authority. Scheduling, Acceptance Review, and backtesting belong to the same task chain but answer different questions: how tasks flow, whether an individual output may enter the next stage, and whether existing judgments and rules hold up when tested over longer periods, across more situations, and against subsequent facts.

Scheduling gets work moving. It assigns tasks, time, and resources, confirms inputs and prerequisites, and passes work to suitable execution units. When results are returned or conditions change, it reschedules. Control over sequence does not confer final authority over professional conclusions, nor may scheduling bypass Acceptance Review to meet a delivery deadline.

Acceptance Review guards the gate through which results enter the next stage. Under preset standards, it checks whether outputs can be understood, verified, and used further, rather than accepting an executor's claim that "I am done." Only qualified results may proceed. Results that do not meet admission conditions should be held back while problems are explained and supporting bases checked. Where revision is clearly needed, they are then returned to the corresponding work unit. Executors cannot approve their own work, and scheduling cannot substitute "task completed" for Acceptance Review.

Qualification of components and complete delivery must be judged separately. Independent Acceptance Review examines both individual outputs and whether interfaces, dependencies, specifications, and objective coverage are complete after consolidation. Scheduling may organize assembly but cannot, on its own authority, release the assembled output as a result. If contradictions or gaps remain, the work should be returned to the stages affected.

Backtesting extends the time horizon. An output passing at the time does not prove that rules remain effective in another environment, or that correctness reflects stable capability rather than short-term luck. Ongoing inspection extends observation beyond a single Acceptance Review to different times and scenarios. Historical backtesting examines the performance of a specified version of a method using existing materials; validation against subsequent facts observes performance on new data; stress testing examines adverse situations; regression checks and exception monitoring help identify effects of changes and operating deviations. These inspections complement one another. Each retains its own objects, versions, samples, and conclusions.

Separating these three responsibilities does not mean every organization must mechanically establish three fixed positions. What must be preserved is independence of judgment: scheduling cannot replace Acceptance Review; Acceptance Review cannot replace long-term backtesting; and producers cannot end review with their own explanations. Position arrangements may vary with task scope and risk. Separation of production and review, retention of dissent, and the ability to return problems to the correct place must remain.

Separation of authority aims to let a task chain correct itself, rather than make the three parties obstruct one another. Scheduling advances work, Acceptance Review protects admission of results, and backtesting examines whether rules and capabilities remain effective. Problems found in Acceptance Review are first clarified and located; those requiring correction enter rework. Biases found through ongoing inspection are first distinguished by input, execution, process, or rule causes, then addressed accordingly. Rule changes require approval by an authorized party before returning to scheduling and Acceptance Review.

9.5 Every Agent Is Assigned to a Named Human: Temporary Collaboration and Continuing Responsibility

Agents may execute, but cannot be the endpoint of organizational responsibility. Every Agent is assigned to a named human, so its Governance Hat, authorization, and issues beyond the rules can return to an identifiable person. When exceptions occur, the organization is not left with only a model name, an instruction, and a result that nobody takes responsibility for.

The accountable person's core work is to govern the Governance Hat, not do the Agent's work. That person needs to know whether rules are clear, authorization is appropriate, processes leave records, and backtesting and exception signals indicate that the Governance Hat needs adjustment. Routine tasks covered by rules, with explicit authority and inspectable results, should be completed by Agents under the rules. Escalation triggers, rule gaps, or a need for discretion return the issue to people authorized to handle it.

Named Human Accountability Assignment is most vulnerable to two extremes. One is item-by-item approval: every action awaits human confirmation, and the organization is again constrained by human capacity. The other is nominal assignment: a name is registered, but no one truly understands operating status or addresses weakened rules and exception signals. Effective assignment requires people to step away from execution of every step while staying close to rules, records, and backtesting. They should also explain failures within their responsibility to identify weakened rules and abnormal conditions.

Problems recurring in actual operation should inform inspection and revision of rules. Repeated Agent disagreement at the same stage, changed input conditions, or rules that cannot clarify the next action cannot be resolved only through repeated rework. The accountable person should identify whether the issue lies in execution, authorization, or failure of rules to cover the current situation, then advance the corresponding adjustments. Executors cannot amend rules themselves, but should submit identified problems, related evidence, and handling records to people authorized to revise them.

Named Human Accountability Assignment has boundaries of its own. The accountable person is responsible for Governance Hats and authorization and does not thereby automatically become every task's scheduler, much less replace independent Acceptance Review or backtesting. The organization must separately identify who assigns work, produces outputs, judges qualification, and handles matters outside the rules. Concentrating all these roles in one person may appear to centralize responsibility but actually makes it unclear again.

Collaborative teams may form for tasks and disband after delivery, but Named Human Accountability Assignments and work records cannot disappear with them. Position identities carry continuing responsibilities, and every run leaves corresponding evidence. If problems emerge later, the rules and authorization then applicable and the people responsible must still be identifiable.

9.6 Three Sufficiency Conditions for Expanding Staffing: Sufficiently Numerous, Fine-Grained, and Counterbalancing, Within Supporting Capacity

Staffing in HAT over MASS is more than a headcount list. It must address whether tasks are covered, errors can be located at specific stages, and judgments face challenge. Staffing must therefore be sufficiently numerous, fine-grained, and counterbalancing at the same time. These three conditions correspond to coverage, attribution, and checks and balances. Without any one, vast numbers of Microtasks are difficult to organize effectively.

Sufficiently numerous staffing first ensures that necessary tasks have executors and key judgments have room for cross-review. With insufficient staffing, one Agent may cross too many task boundaries

and important stages may rely on a single judgment. But "sufficiently numerous" never means unlimited expansion. Every new Agent needs corresponding rules, inspection, and accountability. However large the workforce, it is only a number if its scale does not match tasks and inspection relationships.

Sufficiently fine-grained means dividing responsibilities finely enough to locate errors. Even if tasks are already small, a vague role encompassing input selection, judgment, Acceptance Review, and revision still leaves the source of a problem unclear. Finer divisions clarify inputs, outputs, and responsibilities at each stage and let unqualified results return to the place that actually needs correction. The cost is more handoffs, so rules and scheduling must reconnect these fine units into one chain.

Sufficiently counterbalancing means that different positions truly provide checks and balances. "Counterbalancing" does not mean creating organizational conflict or arguing for its own sake. It prevents all executors and inspectors from following the same assumptions and arriving uniformly at the same bias. Different positions must form judgments independently, retain dissent, and resolve disagreements through evidence and rules. Only then may systemic bias be seen before it spreads.

"Sufficiently numerous, fine-grained, and counterbalancing" address task coverage, division of responsibilities, and independent checks and balances, respectively. Enough positions with unclear responsibilities still make errors difficult to locate. Fine divisions without independent checks from different positions may still repeat biases across stages. Ample disagreement without positions and scheduling arrangements to handle review still leaves problems unresolved. Only when all three conditions hold can large numbers of Agents move beyond simultaneous execution to clearly divided, orderly, mutually inspectable collaboration.

The limits of staffing expansion are jointly determined by capacity to support rules, inspection, and responsibility. More task types require maintenance of permitted actions, prohibitions, and escalation conditions; more outputs require supporting evidence and review capacity to keep pace. New Agents are only potential execution capacity. They become effective collaboration only when rule coverage and inspection capacity can support them.

Exceptions also need people able to handle them. Unclear accountable parties cause problems to circulate among tasks. Concentrating all matters beyond the rules on a few people creates another bottleneck in human handling capacity. Staffing expansion must check capacity for rule maintenance, evidence review, and Escalation to an Authorized Human together; headcount cannot run ahead of them.

A staffing design does not remain valid indefinitely once it has been drawn up. Tasks and capabilities change, and factual feedback changes the requirements for coverage, attribution, and checks and balances. The organization must continuously identify uncovered tasks, stages with overly broad responsibilities, and judgments lacking independent challenge, then adjust accordingly. The next chapter examines how Full-Coverage Inspection, Ongoing Backtesting, and complete traceability provide fresh evidence for these arrangements every day. Without ongoing inspection, even complete staffing remains only a design plan.

Condition	Specific Meaning	Inspection Points
Sufficiently numerous	Cover necessary tasks and cross-review	Whether execution, inspection, and exception handling capacity expand together
Sufficiently fine-grained	Divide delivery and responsibility into identifiable units	Whether problems can return to specific positions and tasks
Sufficiently counterbalancing	Different positions and responsibilities provide checks and balances	Whether production, inspection, and coordination retain independent opinions

Table 9-1: Three Sufficiency Conditions for Expanding Staffing

Chapter 10 How Every Delivery Can Withstand Inspection

Agents extend human capability to task volumes that were previously unaffordable and raise another question: as outputs increase, undiscovered errors within them may also increase. Expansion cannot therefore be judged only by how many Agents are working. It also requires examining whether every task has rules, every result has evidence, and every exception has an identifiable handling responsibility. The preceding chapter separates scheduling, execution, Acceptance Review, and backtesting ultimately to answer: what grounds do we have to trust this result?

Trust here attaches to a specific result and carries explicit conditions. It must withstand successive questions: whether the result complies with current rules, inputs hold, the process can be reconstructed, subsequent facts continue to support the original judgment, and who handles problems. Inspection, backtesting, and traceability each answer part of these questions and together determine how far trust can extend.

10.1 Full-Coverage Inspection Rather Than Sampling: Every Output Is Inspected

Traditional sampling reflects a practical constraint: limited inspection resources require inferring the whole from part of a sample. It can help assess overall risk, but cannot prove that every individual output has passed. Once Agents run concurrently at scale, outputs arrive faster and in greater numbers. If inspection remains limited to a few samples, any error not sampled may enter the next stage directly, be cited and copied, and ultimately amplified. For the person using that result, "generally reliable overall" is insufficient; that person needs to know that this specific result has been inspected.

Full-Coverage Inspection in HAT over MASS means that every output enters an Acceptance Review or inspection process matched to the nature of its task, leaving no default exemption from inspection. It does not require all tasks to use the same method or people to redo every result from the beginning. For tasks with clear rules and stable structures, Agents may check completeness, consistency, and boundaries. Matters with significant effects, beyond rule coverage, or requiring discretion must undergo deeper inspection and escalation under applicable conditions. After inspection of all outputs is complete, authorized human spot checks may still examine whether inspection mechanisms, high-risk results, and exception handling are genuinely effective. Human spot checks provide additional supervision; they do not replace Full-Coverage Inspection.

AI can also participate in inspection. Repetitive comparison, format checks, Rule Matching, exception detection, and cross-review are suitable for AI, making inspection of every output possible at scale. Yet lower costs must not mean relaxed boundaries. Executors cannot declare their own results qualified, and inspectors cannot depart from established rules to invent standards ad hoc. Production and review must remain separate. Inspection authority, escalation conditions, and ultimate human responsibility remain defined by the rule system represented by Governance Hats.

Every inspection must also leave mutually consistent records. Its conclusion must connect to the task, Rule Version, output inspected, and person handling exceptions. A bare "pass" with no explanation of what was checked, under which rules, or to whom problems were referred is not complete evidence. Full-Coverage Inspection does not aim to accumulate formal records; it gives every output a clear route into inspection and a basis for judgment.

10.2 Ongoing Backtesting: Identifying Short-Term Luck and Environmental Change

Passing Acceptance Review establishes only that the result passed under the inputs, rules, and inspection conditions of the time. It does not prove that capability remains stable or rules remain effective in another environment. A correct answer may reflect a reliable method or merely a favorable sample. Rules effective in one period may lose explanatory power in another market, business cycle, or boundary situation. Ongoing Backtesting distinguishes stable capability from short-term luck.

Backtesting places a judgment of a defined version into a longer history, different scenarios, and boundary conditions. There is no single measure of "sufficiently long" for all tasks. Short-cycle tasks may need to cover

different inputs and abnormal states; tasks affected by seasonality, operating cycles, or market changes must span the corresponding phases. Only when capability repeatedly holds under multiple conditions do we have grounds to believe that it mainly reflects a stable method rather than temporary luck.

Backtesting must also identify when answers become outdated. Data sources, operating methods, market environments, and external rules change. A previously correct conclusion cannot continue under an old Acceptance Review if its supporting conditions have changed. Comparing new facts with the original judgment reveals whether deviations arise from the environment, inputs, execution capability, or rules failing to cover new situations. Different causes require different handling: updating materials, calibrating capabilities, revising rules, or referring the matter to a human for renewed judgment.

Backtesting should therefore not end after a single pre-release exercise; it is a condition for trust to continue. Frequencies, samples, and thresholds need not be forced into one standard, but every backtest must specify the version inspected, periods and factual conditions used, and how deviations are handled. Without this information, Ongoing Backtesting remains an appealing principle with no practical basis for distinguishing capability, luck, and environmental change.

10.3 End-to-End Traceability: From Results Back to Inputs, Calls, Rules, and Responsibility

Once inspection finds a problem, the next step is to locate where it began, rather than attributing it generally to "the model was wrong." Looking only at the Agent that submitted the final result is far from enough. An output is often affected jointly by input materials, model calls, Rule Versions, task handoffs, and inspection methods. End-to-End Traceability follows the result backward and recovers the key conditions that formed it, step by step, so that errors can be located, explained, and corrected.

A reviewable evidence chain connects the result to task inputs, model or tool calls, applicable Rule Version, execution path, and inspection findings. Inputs reveal which facts were available at the time; call records show which capabilities were used; the Rule Version explains the boundaries applied; the execution path shows the task's stages; inspection records explain why the result was accepted, returned, or escalated. Linking the Agent Identifier and Named Human Accountability Assignment then lets the organization answer both how the result was formed and who handles exceptions.

Reconstructing the process distinguishes the nature of the problem. Incomplete or distorted materials are input problems; rules failing to cover a new scenario are rule gaps; conditions lost during handoff are process problems; unsuitable capabilities or versions call for calibration of capability selection. Calling all of these "the model was wrong" overlooks the places that actually need correction. Finding a convenient person to blame will not remove the problem either. Traceability places responsibility where the problem can actually be corrected.

Traceability does not end after revision. Updated inputs, adjusted calls, and changed rules or processes must be retested with similar problems to determine whether the original error is genuinely prevented. Versions, inspection findings, and scopes of applicability before and after revision must also be retained. Later users can then see how the problem was discovered and handled and under which conditions the new method holds again. Proof that something "was changed" does not prove it "was changed correctly."

End-to-End Traceability is an evidence requirement, not a promise about system capability. Every Agent Instance run generates a Run Identifier that links it to its task. The Run Identifier can help locate that run, but is not a task number and cannot replace evidence. The eventual depth of tracing depends on which input, call, rule, process, and responsibility records were actually retained. Missing portions cannot be filled with "the system has passed Acceptance Review" or "the records should theoretically exist."

10.4 Three Sufficiency Conditions for Inspection: Sufficiently Comprehensive, Long, and Deep

The three inspection conditions support trust from three directions: sufficiently comprehensive examines coverage; sufficiently long examines time; sufficiently deep examines reconstructability. They must work together, and none may remain only a slogan.

Condition	Specific Meaning	Inspection Points
Sufficiently comprehensive	Every output enters the corresponding Acceptance Review or inspection	Whether inspection relies only on a few samples or leaves default exemptions
Sufficiently long	Continuously cover different periods, scenarios, and boundaries	Whether short-term luck is mistaken for capability or outdated answers remain in use
Sufficiently deep	Trace back through inputs, calls, rules, and responsibility	Whether error sources can be located and revision effects validated

Table 10-1: Three Sufficiency Conditions for Inspection

Sufficiently comprehensive requires every output to enter inspection. Inspection depth may vary with task risk, scope of impact, and rule maturity; authorized human spot checks may continue after Full-Coverage Inspection. But however many tasks there are, some results cannot be allowed to bypass inspection by default. Human spot checks provide additional examination; they do not return Full-Coverage Inspection to a sampling approach.

Sufficiently long requires backtesting to span environments and boundaries genuinely relevant to the task. Correctness at only one time and under one condition neither excludes chance nor reveals deviations after environmental change. The task and its patterns of change determine how long is enough. A uniform timeframe cannot answer this question.

Sufficiently deep requires going beyond output text, one Run Identifier, or a "pass" status. A Run Identifier links one Agent Instance run to its task but is not a task number. Explaining the source of deviations, which stage should correct them, and whether revisions genuinely prevent similar problems still requires returning to inputs, calls, rules, inspection, and responsibility records.

Without any one of these three conditions, a structural gap remains. Comprehensive inspection without sufficient time may mistake repeated short-term good luck for stable capability. Long backtesting without deep enough tracing may reveal deviations without showing where to intervene. Deep tracing without sufficient coverage still leaves outputs that never entered inspection. The four sets of conditions—tasks, foundations, staffing, and inspection—together form the complete Twelve Sufficiency Conditions. Part IV turns these design conditions into positions and operating procedures.

Part IV AI Operating Manual

The Workshop Manual for 5,000 AI Employees

For Named Human Accountable Owners, Technical Due Diligence Teams, and AI Auditors
Named Accountability Assignment, Task-Based Team Formation, Independent Acceptance
Review

This part is the AI operating manual for HAT over MASS and the workshop manual for manufacturing the registration and measurement instruments. It explains the positions assigned to 5,000 AI employees, task decomposition, team formation rules, and the Seven-Stage Task Workflow, enabling named human accountable owners, technical due diligence teams, and AI auditors to examine responsibilities, procedures, and records.

General readers can begin with each chapter's "Chapter Overview" and return to specific steps as questions arise. The AI Position Register, individual tasks, and individual instance runs are managed separately. This white paper uses an illustrative chain convenience store simulation to demonstrate handoffs and inspection; it does not present the simulated sequence as evidence of actual backend performance. The manual makes execution and accountability verifiable.

Chapter 11 How 5,000 AI Employees Fit into the Organization [Operating Procedures]

Chapter Overview | This chapter answers: How do 5,000 AI employees fit into the organization?

Key rules:

- An AI employee is a relatively stable position assigned to a named human who assumes accountability. The position carries ongoing responsibilities and accumulates work records;
- When a specific task arrives, an Agent Instance is launched to meet the position's requirements, calling on models, skills, and tools for that execution;
- Each run is recorded separately and linked to its position and task;
- After the task ends, the position's identity, accountability assignment, and historical records remain in place.

**General readers may stop here. The detailed text is intended for professionals such as organizers, practitioners, technical due diligence teams, and AI auditors.*

Part III requires humans to retain control over rules, authorization, and ultimate responsibility. Through three structural layers, four functions, and named accountability assignments, this chapter translates those requirements into identifiable relationships among positions (see Section 7.3 of Chapter 7 and Section 9.5 of Chapter 9).

11.1 The 5,000 AI Employees Already in Position

The 5,000 AI employees are already in position across four functions: Administration, Business, Control, and Deployment. "AI employee" is the public-facing name for a position: each position carries ongoing responsibilities, a named accountability assignment, and work records. An Agent Instance performs a specific execution. The positions have entered the organizational structure; the stage reached by specific task runs and validation still needs to be demonstrated by corresponding records.

As the workforce grows, the first distinction to clarify is between positions and runs, rather than how many models there are. An AI Organizational Position is a relatively stable identity within the organization, carrying continuing responsibilities, rules, organizational affiliation, and records. When a specific task arrives, the system launches an Agent Instance according to the position's requirements and task conditions, calling on the necessary capabilities to complete a run. A position may participate in multiple tasks over time and use different capabilities at different times. When the instance ends, the position's identity remains.

Being "already in position" and "currently executing a task" therefore cannot be treated as the same thing. The former means that a position has entered the organizational register and has an identifiable identity and responsibility relationships; the latter means that an Agent Instance is performing a specific run. Having 5,000 AI employees in position does not mean that 5,000 instances are always working simultaneously, nor can it establish how many Agents a task has called on. The scale of positions, task organization, and individual runs must be recorded separately so that one set of data is not passed off as another.

When a position enters the register, its responsibilities, accountability assignment to a named human, and actual task records must all be put in place. The human accountable owner is responsible for rules, authorization, supervision, and escalated matters requiring discretion; routine work assignment, team formation, and inspection are entrusted to Agents collaborating under the rules. We do not send every step back for human approval, because that would reduce collaboration at scale to human speed. We retain human responsibility for rules and escalation.

Records must also connect a stable position to each task. The Agent Identifier, Agent ID, identifies the position and links it to its accountability assignment to a named human; the Agent Capability and Performance Profile records the sources, versions, and performance of capabilities actually used; each Agent Instance run generates

a Run Identifier linked to its task, permissions, inputs, outputs, and inspection records. The next chapter develops these three categories of information. This chapter first answers an earlier question: how 5,000 AI employees are organized into HAT over MASS, with rules, defined responsibilities, and the ability to collaborate as tasks require.

11.2 Three Structural Layers: Rules, Responsibilities, and Identification

A list of positions cannot sustain a large AI organization. Without common rules, general capabilities may cross authorization boundaries; with unclear responsibilities, even when multiple Agents can do the work, it is unclear who should undertake it; if identities are disconnected from run records, the process cannot be reconstructed after a task ends. NUMA HAT over MASS therefore establishes a Rules Layer, a Responsibilities Layer, and an Identification Layer, bringing action boundaries, organizational responsibilities, and execution evidence into one structure.

The Rules Layer first defines action boundaries: what may be done, what may not be done, and which circumstances require pausing the affected automated execution or escalating the matter. The permitted task scope, permission conditions, delivery standards, and exception-handling requirements must become constraints that Agents can execute when accepting tasks, calling on capabilities, and submitting results. Where rules do not cover a matter, evidence is insufficient, or a value judgment is required, no level of capability permits an Agent to expand its own authorization. It should stop and refer the matter to a human under the established conditions.

The Responsibilities Layer identifies who undertakes the work and how participants work together. The standing organization includes Administration, Business, Control, and Deployment; specific tasks also contain different responsibility roles. "Can do" describes capability; "should do" concerns responsibility. An Agent's ability to generate a result does not authorize it to approve that result; a position that has coordinated a process cannot replace a quality judgment with "the process is complete." The Responsibilities Layer separates coordination, production, inspection, and support and provides interfaces for collaboration across positions.

The Identification Layer connects a stable identity to every run. The Agent Identifier identifies the position that participated, the Agent Capability and Performance Profile explains the capabilities on which this allocation was based, and each Agent Instance run generates a separate Run Identifier linked to its task and process records. A position's identity does not end with a task; capability and performance records accumulate through actual use, and a single task may contain multiple instances and multiple Run Identifiers. Identification enables a reviewer to trace a result back to the position, capability configuration, and specific run.

Figure 11-1 shows how the three layers connect and how inspection and backtesting feed back to humans.

The Three Layers of HAT over MASS




Figure 11-1: The Three-Layer Structure of HAT over MASS

Each layer governs one aspect and constrains the others. The Rules Layer defines boundaries but cannot assume a position's responsibility; the Responsibilities Layer assigns work but cannot demonstrate how it was completed; the Identification Layer connects records but does not automatically make rules reasonable or results correct. Only when all three operate together can the organization retain stable identities, form teams flexibly around tasks, and locate the stage that actually needs attention when a problem arises.

A Governance Hat brings these three relationships into an operational organizational unit: it connects a position's role, rules, permissions, and responsibility, but is neither an administrative rank nor another name for a model. The same underlying capabilities can serve different positions, provided that the Governance Hat worn for this task, the permissions granted, and the deliverables required are sufficiently clear and the calls are recorded. We can thereby share a capability foundation and retain real and necessary differences among positions.

11.3 Four Functions: Administration, Business, Control, and Deployment

The three-layer design provides the framework; the four functions explain how the 5,000 AI employees divide their work. These are not four groups of Agents that never interact, nor four approvals through which every task must mechanically pass. The organization has four continuing responsibilities: organizing objectives into tasks, producing business results, independently checking boundaries and quality, and preparing models, data, tools, and the operating environment.

A: Administration turns objectives into executable, inspectable work arrangements. It identifies task scope, dependencies, and delivery requirements, then organizes decomposition, position selection, handoffs, and validation. Coordination brings the right positions into a task at the right time; it does not make judgments on behalf of specialist positions. Later chapters explain in detail how an Industry Compass is decomposed into work units and how AI employees form teams around tasks.

B: Business directly produces the analysis, processing, and deliverables corresponding to task objectives. Industries differ in the knowledge, data, and working methods required, so business positions need specialized responsibilities. They are the main producers of results but cannot grant final acceptance to their own results. Important judgments must still undergo the appropriate checks, and the inputs, evidence, and processes supporting those judgments must also be retained.

C: Control continually asks three questions: Is the work within authorization boundaries? Is the evidence sufficient? Does the result meet the requirements? Risk, rules, legal, financial, compliance, and necessary independent checks fall within this function. Business positions cannot perform all validation of their own results, and coordinating positions cannot substitute "the process is complete" for a quality conclusion. When problems are found, Control must require clarification, reexamination, revision, or escalation.

In this white paper, "Control" means the AI function of gathering evidence, checking, flagging exceptions, conducting design reviews, and inspecting quality within its authorization. Compliance approvals, investment judgments, suitability judgments, final enterprise scores that take effect externally, and legal opinions are made or confirmed and recorded by humans with the appropriate authorization and qualifications. AI may still calculate and propose scores under the rules; internal calculation, independent inspection, and final approval carry different responsibilities.

D: Deployment prepares and maintains models, data, skills, tools, systems, and operating environments. An AI position's ability to work consistently depends on the availability of data paths, tool integration, permission configuration, and version control, as well as model capability. Deployment makes these conditions available as tasks require and records changes. It can demonstrate whether technical conditions are available, but cannot determine for Business or Control positions whether a conclusion is correct.

The four functions form a chain of responsibility: Administration organizes tasks, Business produces results, Control checks boundaries and quality, and Deployment provides operating conditions. Depending on difficulty, risk, and dependencies, a specific task may proceed in parallel or in stages; it need not be fixed into an unchangeable approval pipeline. Continuing functions keep responsibility stable, and Dynamic Team Formation enables 5,000 AI employees to collaborate quickly around an objective.

11.4 Multi-Perspective Roles: Three Industry Perspectives + Cross-Checking by Three Models, Three Investor Perspectives, and "6+1" Review Perspectives

Organizational responsibilities establish who undertakes each type of work; an Industry Compass must also establish how to view the same business. Different positions see different facts and raise different questions. NUMA HAT over MASS therefore establishes Multi-Perspective Roles, separating operational observation, Cross-Model Cross-Checking, investment judgment, and design inspection. A perspective is a responsibility role that determines what a position focuses on and what judgments it provides; it is neither a model count nor a model's name.

The Three Industry Perspectives are defined as the Ordinary Operator Perspective, Industry Veteran Perspective, and Adversarial Manipulation Perspective. The Ordinary Operator Perspective begins with daily operations and assesses how a factor affects revenue, cost, efficiency, and business continuity. The Industry Veteran Perspective brings long-standing industry experience to identify gaps between rules on paper and actual operations, as well as cycles, conventions, and implicit constraints. The Adversarial Manipulation Perspective specifically searches for ways to repackage indicators, exploit specifications, or inflate results. The "adversarial manipulation" role here is an adversarial position in design review; it does not label any specific operator. Bringing the three views together reduces the bias of understanding a business solely through an idealized process.

Cross-checking by three models concerns calls on underlying intelligence and is distinct from the Three Industry Perspectives. For the same task, we have different models work independently, then compare their assumptions, reasoning, and conclusions. When differences emerge, we continue checking inputs, specifications, and evidence. Multiple models can help expose omissions in a single approach, but do not automatically provide effective cross-checking. Results are comparable only when task conditions, input evidence, output requirements, and the treatment of disagreements are consistent. The formal description is therefore "Three Industry Perspectives + Cross-Checking by Three Models." They cannot be combined into nine industry perspectives, and models need not be bound one-to-one to positions.

The Three Investor Perspectives focus on business, price, and risk. Business assessment asks how an enterprise creates value and whether its operating mechanism is sustainable; price assessment examines how value judgments translate into pricing conditions; risk assessment checks adverse circumstances, vulnerabilities, and potential losses. Industry perspectives first develop an understanding of the business; investor perspectives then ask what that business, price, and risk mean for capital decisions. Their subjects are connected, but their responsibilities differ.

After operational observation and investment judgment produce candidate factors, the "Six Review Perspectives + Eligibility and Composition Review" are responsible only for this round's design review, eligibility determination, and composition selection. Business Relevance Review examines a candidate factor's operating relevance; Calculation and Scoring Specification Review checks whether its definition, calculation, and scoring can be reproduced unambiguously; Data Source and Access Review checks credible data paths and necessary conditions; Discriminative Validity Review examines the basis for distinguishing operating conditions; Incremental Information Review checks what information it adds relative to the Full Factor Pool; Manipulation Resistance Review tests whether specifications or inputs can be manipulated to inflate scores without improving the actual state. Eligibility and Composition Review then classifies the findings as Excluded, Pending Clarification, Pending Revision, or Eligible, and selects compositions from the Eligible Pool. Eligible means only that a candidate qualifies for composition discussion, not that it has been finally selected, much less that enterprise data collection or actual validation is complete. Enterprise data collection, actual scoring, calibration, deployment, and publication belong to subsequent implementation.

Multiple perspectives illuminate different responsibilities. Industry perspectives understand the business; cross-checking by three models compares assumptions and judgments produced by underlying capabilities; investor perspectives examine capital and decisions; the Six Review Perspectives and Eligibility and Composition Review examine this round's candidate factor designs and composition eligibility. Different views need not be rushed into a single answer. Evidence should first be retained and disagreements clarified, followed by a decision under the rules on whether revision or escalation is needed. These checks and balances arise from clear responsibilities and handling procedures, rather than simply increasing the number of Agents.

11.5 Separating Perspectives from Capabilities: Calling on Models, Skills, and Tools by Task

Multiple perspectives work by separating a position's perspective from its execution capabilities. A perspective specifies the type of judgment the position must make and its corresponding responsibility; models, skills, and tools provide the capabilities called on to make that judgment. The responsibility role may remain stable as specific capabilities are adjusted to the task object, evidence conditions, and permission scope. The same model may also serve different positions, but using the same model does not turn different responsibilities into one responsibility.

This arrangement avoids permanently packaging every model as an "employee." Position identities and responsibility relationships can continue as models are updated, skills adjusted, and tools replaced. When a position enters a new task, suitable capabilities are assembled within the rules. A shared foundation reduces duplication, and differentiated configurations ensure that each Agent receives only the background, tools, and permissions required for its current responsibilities. Capabilities may change; responsibility must always remain identifiable.

Calling on capabilities by task still requires boundaries. The choice of models, skills, and tools must correspond to task objectives, permitted scope, and delivery standards; when multiple models participate in cross-checking, inputs and outputs must be comparable and disagreements identifiable; data and tool calls must comply with the corresponding permissions and conditions of use; high-risk judgments or gaps in rules are escalated to humans under the specified conditions. The organization calls on capabilities; the Governance Hat governs how those capabilities are used.

How capabilities are chosen and how they actually perform must both be recorded. The Agent Capability and Performance Profile must identify the capabilities used, their versions, the tasks completed, the results produced, and the inspections conducted. Each Agent Instance run generates a Run Identifier linked to its task, allowing reviewers to ask whether a problem arose from position responsibilities, capability selection, input evidence, or execution and inspection. Records do not automatically provide the right answer, but give validation, backtesting, and revision an identifiable basis.

The relationship between a stable organization and dynamic capabilities is now clear. The Rules Layer maintains boundaries for permitted actions, prohibitions, and escalation; the Responsibilities Layer assigns responsibilities among positions; the Identification Layer connects identities and runs. The four functions maintain continuing responsibilities, Multi-Perspective Roles undertake observation, judgment, and design review, and models, skills, and tools enter successive runs as tasks require. The collaborative capability of 5,000 AI employees comes both from stable positions carrying clear responsibilities and from dynamic capabilities moving within the rules, with a requirement that every execution can be linked to its corresponding identity and records.

This architecture provides a foundation for subsequent operations. The next chapter explains how AI Organizational Positions, Agent Instances, Agent Identifiers, Agent Capability and Performance Profiles, and Run Identifiers connect, enabling stable positions, accountability assignments to named humans, and individual task records to be managed separately and cross-checked along the same chain of evidence.

Chapter 12 How Responsibility Is Distinguished between a Position and a Run [Operating Procedures]

Chapter Overview | This chapter answers: How is responsibility distinguished between a position and a run?

Key rules:

- The Agent Identifier identifies "which position participated," the capability and performance profile explains "why it was allocated," and the Run Identifier records "what was done this time";
- Each run must be linked to its task, permissions at the time, capability versions, inputs, outputs, and inspection results;
- Rework or reexecution must generate a new Run Identifier;
- When a position's status or capability configuration is updated, historical versions should be retained without overwriting past execution conditions;
- Identifiers are only lookup points. Whether a result can be reviewed depends on the completeness of linked records and evidence.

**General readers may stop here. The detailed text is intended for professionals such as organizers, practitioners, technical due diligence teams, and AI auditors.*

Part III makes End-to-End Traceability a condition for trusted delivery. This chapter connects stable positions to individual executions through Agent Identifiers, Run Identifiers, versions, and evidence (see Section 10.3 of Chapter 10). When capabilities change, the records should support revalidation.

Once 5,000 AI employees enter organizational operations, management must answer a more practical question: how exactly are continuing positions, available capabilities, and individual executions distinguished? If the same name refers to a position, a configuration, and a temporary instance, then identities, permissions, processes, and responsibility become entangled when a result needs review.

This chapter separates information into three levels. An AI Organizational Position carries a stable organizational identity, the Agent Capability and Performance Profile provides the basis for retrieval and allocation, and the Agent Instance records a specific execution. Agent Identifiers and Run Identifiers connect them: temporary teams may disband without interrupting continuing identities, and a result can be traced through records to its task, configuration, and evidence at the time.

12.1 AI Employee Positions and Task Instances

"AI employee" is the name used with the organization and its users; internal management must further distinguish positions from instances. An AI Organizational Position enters the continuing register as a formal, relatively stable organizational identity; an Agent Instance is launched as required when a specific task arises and completes a technical run. A position may participate in multiple tasks over time, each using different models, skills, tools, permissions, and context. The position still exists after an instance ends.

This distinction preserves continuity of responsibility. An AI Organizational Position identifies the execution participant's responsibilities within the organization, its human accountable owner, and its current status; an Agent Instance explains what it did, when, in which task, and under which constraints. Continuing records belong to the position; individual execution information belongs to the instance. If they are mixed, stopping an instance may erase clues to responsibility, and an isolated result may be mistaken for a position's long-term performance.

A model name cannot replace a position either. Models, skills, and tools are capabilities called on to complete tasks; positions carry organizational identities and responsibilities. Technical configurations, including models, skills, and tools, may be updated, and Governance Hats and permissions within tasks may change. These changes should enter the capability and performance profile and run records. After a model is replaced, the organization can still identify the same position and compare its performance under different configurations.

Every Agent Instance must have clear task boundaries. At launch, it receives the task objective, current Governance Hat, permissions, context, and delivery requirements. The run then generates a separate Run

Identifier linked to its task, inputs, outputs, call information, and inspection results. After the task ends, the team may disband, but the position identity, accountability assignment to a named human, and historical records remain. We manage stable identities separately from temporary execution and reconnect them through run records.

Figure 12-1 summarizes the relationships among stable positions, task-based team formation, and delivery records.



Figure 12-1: Stable Positions and Delivery Records in Task Collaboration

12.2 AI Employee Agent Identifiers, Capability and Performance Profiles, and Run Identifiers

The Agent Identifier (Agent ID) corresponds to an AI Organizational Position's stable organizational identity, connecting the position's meaning, main responsibilities, status, accountability assignment to a named human, and historical records. It does not change when an instance starts or ends. It provides a common address for this complex information: work completed, status changes, and human accountability assignments can all be gathered under the same identity. The identifier itself does not replace the substance of responsibilities.

The Agent Capability and Performance Profile is used to explain the current basis on which a position can be retrieved and allocated. Business scenarios, industry knowledge, professional roles, models, languages, regions, skills, and tools may all enter the record, but a static name is far from sufficient. Capability sources, versions, and the results produced under specific task conditions directly affect whether the position can be called on next time. The profile therefore serves allocation and is also subject to ongoing calibration.

Each Agent Instance run generates a Run Identifier linked to its task. The Run Identifier connects the objective, Governance Hat, permissions, context, inputs, outputs, model and tool calls, inspection results, and time information for that run. The same position generates different Run Identifiers when participating in different tasks; one task may also contain multiple instance runs. Rework or reexecution likewise requires a new Run Identifier. One task number cannot cover the entire process.

The three categories of information answer three questions: the Agent Identifier identifies the stable position that participated; the Agent Capability and Performance Profile explains why it was allocated; and the Run Identifier

describes exactly what a particular instance did. Without an Agent Identifier, multiple runs cannot be attributed back to a continuing position; without a capability and performance profile, the allocation basis cannot be assessed; without a Run Identifier, the specific process disappears. Connecting all three enables identity, allocation rationale, and execution evidence to be cross-checked.

Identity and run records also cannot replace the Governance Hat and permissions for the current task. The Agent Identifier remains relatively stable; the current Governance Hat identifies the position's role in this collaboration; and permissions limit what it may read, call on, and submit. A position may take different roles in different tasks, but each change must be recorded under the corresponding Run Identifier and matched to the task and authorization conditions at the time. A stable identity does not automatically confer current permissions.

12.3 Separately Managing the AI Position Register, Status Updates, and Instance Information

The AI Position Register uses Agent Identifiers as its organizing reference, retaining each position's meaning, main responsibilities, accountability assignment to a named human, status, capability and performance profile, and historical records. It is not a static roster. The organization uses the register to determine whether a position can currently be called on, how its responsibilities and capabilities have changed, and whether past task experience supports a new allocation.

Agent Instance records retain a single run: the Governance Hat worn, permissions granted, context, models, skills, and tools used, inputs and outputs formed, and inspections conducted. This information arises with the task and cannot directly overwrite the continuing identity. A temporary tool call does not permanently add the tool to a position's capabilities; a failed inspection does not cancel the position's identity. Whether to update status or the capability and performance profile must be determined from the rules and actual run facts.

A position's status describes its actual situation at the organizational level. It must reflect whether the position remains under effective management, whether it meets the conditions for being called on, and when and on what basis changes such as suspension or resumption occurred. Status cannot be inferred automatically from one task result, nor arbitrarily rewritten apart from accountability assignments, configurations, and run records. Maintaining this distinction prevents "already in position" from being treated as the same as "running at this moment."

The Agent Capability and Performance Profile must also retain sources and versions. Available capabilities may be adjusted when industry knowledge, models, skills, or tools are updated, but historical runs must retain the versions and performance actually used at the time. Otherwise, new records would overwrite old conditions, leaving reviewers unable to determine which capability version produced a result. Each update to status or the capability and performance profile should have verifiable run facts and preserve the relationship between the old and new versions.

Separating the continuing register from instance records also allows responsibility and information boundaries to be managed separately. The accountability assignment to a named human connects to the stable position and identifies who is responsible for rules, authorization, and exception handling; specific task materials, context, inputs, and outputs are governed by the permissions of the current instance. The organization need not recreate an identity for every task, nor presume that an existing position identity permits access to all materials or execution of all actions.

This separation ultimately serves allocation. The system should see an allocation basis formed jointly by the position's current status, capability sources and versions, and actual run performance, rather than an isolated label. After a new task starts, every instance run generates a Run Identifier linked to its task. Actual configurations and results then feed back into the register, enabling position status and capability and performance profiles to be updated on a factual basis.

12.4 Records Linked to AI Employee Run Identifiers and Evidence Retention

A Run Identifier is an index: it connects information distributed across an instance run, allowing the organization to trace a result back to its task, current Governance Hat, permissions, context, inputs, outputs, model and tool calls, inspection results, and time information. An index is not evidence. If records are missing, versions unclear, or inputs and outputs cannot be matched, a Run Identifier alone cannot reconstruct the complete process.

A reviewable record must explain what the instance received, the constraints under which it processed the work, the capabilities called on, the results produced, and the inspections conducted. Changes to material sources or task specifications must also enter the linked records of the same run. Reviewers can then see more than the final text or score and check whether inputs, rules, and authorization conditions remained consistent throughout.

In an Industry Compass task, the Agent Identifier first locates the stable participating position and its accountability assignment; the Agent Capability and Performance Profile explains the industry knowledge, professional roles, models, skills, or tool conditions on which allocation was based; and the Run Identifier corresponds to one instance run, linking the task, materials, configuration, outputs, and inspection results. Only together can these categories distinguish "the position has a capability" from "whether the work was completed under the corresponding conditions this time."

Model and tool calls also belong in run records. A position may use different models in different tasks, and different stages of the same task may use different skills and tools. Call information preserves the actual conditions under which a result was produced; technical names do not replace responsibility judgments. After model, tool, or knowledge versions change, historical runs can still be matched to their configurations at the time, providing a basis for subsequent comparison.

When a task ends, instances may stop and teams may disband, but run records and their evidence links must remain. Subsequent Acceptance Review, returns, backtesting, and accountability tracing all begin with these records. How far a review can reach depends solely on the evidence actually retained. Having a Run Identifier does not mean that the organization already has the ability to reconstruct the complete process.

Chapter 13 How Industry Compass Tasks Are Decomposed and Teams Formed [Operating Procedures]

Chapter Overview | This chapter answers: How are Industry Compass tasks decomposed and teams formed?

Key rules:

- An Industry Compass task must first define its observation object, material scope, delivery standards, and constraints, then be decomposed into work units that can be executed and inspected independently. Positions are matched and teams formed according to capabilities, permissions, and availability;
- The task proceeds through seven stages: task intake, Rule Matching, team formation, execution, validation, Deliverable Consolidation, and archiving. Units work in parallel according to their dependencies. Proposing a candidate does not mean it has passed inspection; operating relevance, statistical specifications, and the scoring basis must still undergo review;
- Disagreements are clarified first; work is returned to the corresponding stage only when modification is necessary;
- A substantive change in definition, purpose, or scoring requires a new version and a repeated Full Six-Perspective Review. Matters outside the rules or authorization are referred to a named human.

**General readers may stop here. The detailed text is intended for professionals such as organizers, practitioners, technical due diligence teams, and AI auditors.*

Part III requires Microtasks to have clear boundaries and independently inspectable results. Through task definitions, work units, and the Seven-Stage Task Workflow, this chapter explains how those requirements enter Industry Compass delivery (see Sections 8.3 and 8.4 of Chapter 8).

An Industry Compass needs design review even after a model has generated an initial draft. We must turn our understanding of a business into a factor framework that can be explained, scored, and reviewed: industry boundaries must be clear, input materials applicable, factors able to explain operations, observation specifications reproducible, and Scoring Rules able to distinguish states. If these judgments are all placed in a black box at once, even a complete final table cannot explain where its conclusions came from or where a problem arose.

HAT over MASS first decomposes an Industry Compass into clearly bounded work units, then selects suitable AI employees from the resource pool, launches Agent Instances, and forms a Task-Specific Team serving only the current objective. Stable positions provide identities, capability and performance profiles, and responsibility records; the current task then supplies context, permissions, knowledge, skills, tools, and delivery requirements. Once the task ends, the team disbands, but positions and records remain. We gain flexibility in team formation and preserve continuity of the process.

13.1 Defining an Industry Compass Task: Input Materials, Delivery Standards, and Constraints

Before decomposing a task, we must specify what "generating an Industry Compass" means. An industry name is only a starting point: the same industry may contain different business formats, revenue models, operating scopes, and observation objects. Without clear boundaries, downstream Agents may interpret the subject as an individual store, a chain network, a franchise system, or supply chain management. Even if all factors and Scoring Rules are complete, they cannot fit into one Compass.

An Industry Compass task ready for execution must specify at least five things: the operating activities covered; the industry observation and scoring framework to be developed; the permitted industry materials, operating records, rule evidence, and existing research; the factor names, definitions, quantitative specifications, units, Observation Periods, Scoring Rules, and explanations included in delivery; and the permissions, deadlines,

material boundaries, applicable rules, and escalation conditions. Together, these five elements define a single task.

For input materials, the key is not quantity but whether sources correspond to intended uses. Industry background materials help explain the business but do not automatically establish that scoring thresholds are reasonable. Operating records may support observation, but must also specify the object, time, and measurement specification. An earlier Compass version may be consulted, but old definitions cannot continue to be copied after a version changes. Confirming the material scope and Rule Version at the start ensures that all work units address the same question and that the objective does not quietly shift during execution.

Delivery standards also cannot be reduced to "submit a complete Compass." "Completeness" means that every output has a defined place, its name, definition, specification, and score can be cross-checked, and evidence supporting key judgments can be found. At the start, the task must also identify which results need independent validation, which disagreements should first be clarified under existing rules, and which matters outside the rules require escalation to a named human. Setting boundaries first gives subsequent collaboration a common language.

13.2 Decomposing an Industry Compass Task: Work Units for Factor Identification, Measurement Specifications, and Scoring Rules

Once a task is clearly defined, it is decomposed into work units that can be executed and inspected independently. The three basic types of work are factor identification, measurement specification, and Scoring Rules. Factor identification explains why an operating fact merits observation; measurement specification defines the observation object, calculation method, unit, and period; Scoring Rules explain how observations become scores under industry standards. These activities depend on one another but cannot be collapsed into one ambiguous judgment.

Work units can be tested against "small, narrow, and granular." Small means that a unit has only one clearly identifiable delivery objective. Narrow means that discretion is confined to predetermined rules and authorization: the unit cannot change objectives or expand permissions on its own, and must stop execution and escalate when boundaries are exceeded. Granular means that each output can be inspected independently and key steps replayed. Every decomposed unit must retain complete meaning and be capable of being returned, corrected, or redone on its own.

Consider the candidate "store density." Proposing it as a candidate factor requires the following work: defining stores and trade areas, choosing an Observation Period, and assessing the relationship between density changes and operations. Operating relevance must be explained, the measurement specification must be reproducible, and the scoring direction must be supported by evidence. After decomposition, the layer containing the problem is clarified first; the measurement-specification unit is returned only when its specification clearly requires revision. If operating relevance itself is not established, there is no need to continue subsequent scoring work.

Every work unit entering allocation must carry its inputs, objective, outputs, dependencies, and inspection conditions. Factor identification begins with industry operating questions and existing evidence, producing candidate factors and their rationale; measurement specification takes confirmed factors and produces definitions, calculation methods, units, and Observation Periods; Scoring Rules then use confirmed definitions and evidence to establish direction and ranges. Specifying inputs, outputs, and acceptance standards together makes each result independently inspectable and a reliable input for the next stage.

13.3 Allocating AI Employees: Matching Positions by Capability and Performance Profiles, Permissions, and Availability

Only after work units are clear does the system know which positions to call on. Factor identification requires an understanding of industry operations; measurement specification addresses definitions, data, and calculations; Scoring Rules convert factual observations into stable scales. Different activities may also require different

material permissions, industry knowledge, tool conditions, and independent inspection capabilities. Matching cannot be achieved merely by looking at position names or distributing work evenly among departments.

Position allocation requires the Agent Capability and Performance Profile, available knowledge, permission scope, historical records, and current availability to be considered together. The profile explains how the position has actually performed on related tasks; available knowledge identifies the industry and rule materials it can use; permissions determine the scope of reading and operations; historical records reveal quality and bias; and current availability shows whether a position is suspended, awaiting calibration, or in another status, preventing the same position from both producing and reviewing its own results. Allocation is therefore based on a set of facts, rather than a claim that a position "is good at research."

Models, skills, and tools belong to the current task configuration; they are not the position itself. The formal organizational description is Three Industry Perspectives + Cross-Checking by Three Models. Industry perspectives specify the operating standpoint from which the business is understood; models generate comparable underlying judgments about the same question. The system first determines the required perspectives from the work units, then chooses suitable models, skills, and tools.

Once positions are selected, the system launches Agent Instances for the current work. Every instance run generates a separate Run Identifier linked to its task, objective, inputs, permissions, Rule Version, outputs, and inspection records. The same position may participate in different tasks over time, and a single task may contain multiple instance runs. When work is redone, both the earlier and later runs are retained; one task number or the final result cannot overwrite every individual process.

This step separates continuing positions from the current execution. A stable identity enables capabilities, experience, and responsibility to accumulate; launching instances by task contains context and permissions within the current boundaries. Dynamic Team Formation is not a matter of temporarily finding a few models to write an answer together; it accurately connects defined work, suitable positions, and the conditions for each run.

13.4 Forming AI Employee Teams: Establishing Task-Specific Teams, Handling Disagreements, and Ending Collaboration

After positions are selected for individual work units, they must form a Task-Specific Team. The team exists only for the current Industry Compass objective and does not create a permanent department. Members retain their existing Agent Identifiers, capability and performance profiles, and human accountability assignments and receive the context, Governance Hats, permissions, knowledge, skills, tools, and delivery standards needed for this task.

The Scheduling Agent maintains continuity across the task. It holds the overall objective, work-unit dependencies, and current status, confirms whether the preceding output sufficiently supports the next step, and handles handoffs when positions or models are replaced. Coordination connects the stages; it does not make every judgment along the way. Business results are still produced by the relevant positions, and validation cannot be handed back to the result provider to perform on its own.

Disagreements are normal in industry research. Different perspectives may reach different judgments about operating relevance, data reliability, or scoring direction. The Scheduling Agent first locates the disagreement in a specific work unit, organizes the parties to explain their evidence, and clarifies concepts. It returns work to the relevant stage only after confirming that content requires modification. Insufficient evidence sends the work back to material preparation; a substantive change in definition, purpose, or scoring requires a new version and a repeated Full Six-Perspective Review. When existing rules cannot cover a matter, or permissions or scoring principles change, it is escalated to a named human for discretion. Coordination cannot change rules on its own to accelerate progress.

After the task is completed, the Task-Specific Team disbands and positions return to the resource pool for their next assignment. Ending collaboration must not erase identities or history. The Run Identifier, task, inputs, outputs, disagreement handling, revision records, and final deliverable for each instance run remain, and human

accountability assignments continue. Validated experience may inform capability calibration and rule revision, but a single collaboration does not automatically create a permanent organizational assignment.

13.5 The Seven-Stage Task Workflow for AI Employees: Task Intake, Rule Matching, Team Formation, Execution, Validation, Deliverable Consolidation, and Archiving

Work units and Task-Specific Teams need a common task sequence. An Industry Compass proceeds through seven stages: task intake, Rule Matching, team formation, execution, validation, Deliverable Consolidation, and archiving. The purpose is not to introduce seven additional approvals, but to specify what each step receives and produces, the conditions required to continue, and where to return when problems arise.

Stage	Main Work	Result Produced	Exception Handling
Task Intake	Confirm industry scope, task objectives, input materials, and expected deliverables	A task record with clear boundaries	Return for supplementary information if scope or materials are unclear
Rule Matching	Confirm applicable rules, versions, permissions, and escalation conditions	Rules and constraints adopted for this execution	Refer gaps in rules or permission conflicts for handling
Team Formation	Match positions to work units, configure context, and launch Agent Instances	A Task-Specific Team, assigned responsibilities, and Run Identifiers	Reconfigure if capabilities, permissions, or availability do not meet requirements
Execution	Complete factor identification, measurement specifications, Scoring Rules, and other work according to dependencies	Component outputs accompanied by "sources" and "versions"	Return to the corresponding preceding unit if inputs conflict or premises change
Validation	Check component results against inspection conditions	Acceptance opinions, objections, or return requirements	Return nonconforming results to the relevant stage for rework
Deliverable Consolidation	Combine results that have passed validation and address unresolved matters	Structurally complete deliverables	Do not proceed to archiving while disagreements remain unresolved
Archiving	Retain task status, inputs, outputs, run records, and responsibility links	Task archives available for subsequent review	Complete missing records before completing archiving

Table 13-1: The Seven-Stage Task Workflow for Industry Compass Tasks

The seven stages make status visible. Before a work unit begins execution, it should be clear whether dependencies are satisfied; before a result enters Deliverable Consolidation, it should be clear whether validation has passed. When a problem is found, its location is first clarified, and work is returned to the specific stage only when revision is necessary, rather than issuing a blanket requirement to "redo everything." Different units in the same task may work in parallel according to dependencies, but each output must be traceable to its own inputs, rules, and Run Identifier.

The validation stage checks whether results meet the inspection conditions preset for the task; it does not replace the Six Review Perspectives and Eligibility and Composition Review. Those roles examine only the design, eligibility, and composition of this round's candidate factors; enterprise data collection, actual scoring, calibration, deployment, and publication belong to subsequent implementation. Deliverable Consolidation does not eliminate all differences. It distinguishes established content, clarified matters, and parts that clearly require revision. A substantive change in definition, purpose, or scoring should produce a new version and a repeated Full Six-

Perspective Review. Archiving retains the whole task chain so that, after a task ends, each output can still be traced to its instance, handling process, and responsible party.

Figure 13-1 brings the seven-stage sequence, candidate design review, and process records together in one task chain.




Figure 13-1: The Seven-Stage Workflow for Industry Compass Tasks

13.6 Example: A Chain Convenience Store Industry Compass: From a Task Sheet to a Candidate

This chapter explains, in sequence, how a task is defined, decomposed, allocated to positions, organized into a team, and advanced through the seven stages. The following example connects these steps using "SIM-01" as the common task identifier to show how an Industry Compass enters task intake, handles disagreements in collaboration, and undergoes inspection. SIM denotes simulation only and is not an actual backend task number; V0.1 and V0.2 are illustrative versions for this example, not approved industry deliverable versions.

Task Intake: Narrowing the Meaning of "Chain Convenience Store"

A chain convenience store may mean one store, a regional network, a brand headquarters, or a franchise system. Different observation units mean that average daily sales, gross margin, and inventory turnover days are measuring different things. The task intake stage therefore narrows the scope first, then records inputs, deliverables, and constraints on the same task sheet. Rule Matching confirms the Scoring Rules version for this round and lists uncovered matters, such as a uniform standard for delineating trade areas, as escalation conditions.

Item	Assumptions in This Example
Task	SIM-01: Generate candidate factors, quantitative specifications, and Scoring Rules on a 1-to-3 scale for chain convenience stores

Item	Assumptions in This Example
Observation Object	Stores within a chain system and their franchise network, focusing on monthly operating facts; supply chain financing and headquarters capital structure are outside the scope of this task
Inputs	Industry background materials, publicly available operating knowledge, and existing research; this example does not use measured enterprise data
Deliverables	Candidate factors for the Return Dimension and Risk Dimension, each with a definition, purpose, statistical object, unit, period, data path, Scoring Rules, and review records
Constraints	Rule design only; enterprise data collection, actual scoring, calibration, and publication belong to subsequent implementation
Escalation Conditions	Refer to a named human when a new statistical standard not covered by existing rules is required, or scoring principles must be changed
Candidate Tracked in This Example	The candidate "store density," starting with V0.1

Table 13-2: The Simulated Task Sheet for Chain Convenience Stores

Team Formation and Execution: Multiple Perspectives Begin Work Together

At team formation, positions are matched to work units to form a temporary team: A-Coordination manages task status and dependencies; B-Industry Analysis, B-Measurement Specification, and B-Scoring undertake the three types of work units; C-Six Review Perspectives and Eligibility and Composition Review perform design review, eligibility determination, and composition selection; D-Support maintains material permissions, versions, and record links. When positions enter the task, the system launches Agent Instances for them, and every instance run generates a separate Run Identifier.

At execution, the Three Industry Perspectives and three models work simultaneously and offer judgments. Investor perspectives then examine business, price, and risk, and candidate factors for the return and risk sides emerge in succession, forming this task's candidate pool. This example does not elaborate each factor; it selects just one from the pool to track: "store density."

The starting point for "store density" is intuitive: convenience stores serve nearby customers, so the number of stores in the same trade area affects the customer traffic available to each. V0.1 therefore contains only a name and one statement of purpose: "Reflects competition and coverage in the store's trade area."

That statement leaves three questions unanswered: Are the stores counted those of the same brand, competitors, or all convenience stores? How is the trade area delineated, by whom, and when? Is having more stores good or bad for this store?

At this point, the seven-stage workflow has reached its fourth stage: the task has boundaries, the work is decomposed, positions have identifiers, and the candidate has a version. V0.1 arrives at B-Scoring with these three questions (see Section 14.6 of Chapter 14).

Chapter 14 How Continuing Positions Support Dynamic Collaboration [Operating Procedures]

Chapter Overview | This chapter answers: How do continuing positions support dynamic collaboration?

Key rules:

- The departmental affiliation line maintains positions' continuing identities and responsibilities; the task workflow line forms teams across departments as delivery requires. Temporary collaboration does not change a position's affiliation or named accountability assignment;
- A team must cover four functions: coordination, execution, control, and support. Coordination handles handoffs and overall delivery coordination, but cannot replace professional judgments or bypass independent Acceptance Review; execution positions produce component outputs but cannot approve their own results; control positions independently check permissions, evidence, and rule boundaries and cannot withdraw objections or lower inspection conditions to meet delivery schedules; support positions ensure operations but cannot treat technical availability as business acceptance;
- Named humans undertake rules, permissions, and escalated matters according to their authorization;
- Overall coordination and component responsibilities are recorded separately and remain traceable after the team disbands.

**General readers may stop here. The detailed text is intended for professionals such as organizers, practitioners, technical due diligence teams, and AI auditors.*

Part III proposes organizing task collaboration separately from rule supervision. Through the two lines of departmental affiliation and task-based team formation, this chapter explains how continuing responsibilities remain continuous and how current collaboration can proceed flexibly (see Section 9.1 of Chapter 9).

After an Industry Compass has been decomposed into executable, inspectable work units, the distributed results must be organized back into a complete deliverable. Materials, industry analysis, scoring judgments, control checks, and technical support may come from different departments, and participating positions may change with the task. Looking only at a temporary team can obscure continuing identities and responsibility records; looking only at fixed departments makes it difficult to bring capabilities across departments together promptly around the same objective.

Collaboration along two lines connects these needs. The departmental affiliation line maintains AI Organizational Positions' continuing identities, capability and performance profiles, and responsibility assignments; the task workflow line calls on positions and launches Agent Instances around the current objective. Agent Identifiers, the AI Position Register, and Run Identifiers connect the lines: each instance run generates a Run Identifier linked to its task, enabling inputs, judgments, inspections, and deliverables to be traced back to the corresponding position and named accountable person. Teams can change flexibly without responsibility ending when they disband.

14.1 The Departmental Affiliation Line and the Task Workflow Line

The departmental affiliation line manages relatively stable position relationships. It continuously records a position's department, capabilities, completed work, and named human accountable owner. Departmental affiliation does not restrict a position to tasks within that department. It provides a continuing identity and management foundation, enabling the organization to maintain rules and permissions, update capability and performance profiles, and connect multiple tasks into a continuous history.

The task workflow line organizes work solely around one specific delivery. Once an Industry Compass enters execution, the system calls on positions from different departments according to objectives, inputs, delivery standards, and constraints, launching Agent Instances as needed. Every instance run receives a separate Run Identifier linked to its task, current context, permissions, knowledge, skills, tools, and inspection requirements. When the task ends, the temporary collaboration ends, positions return to their original affiliations, and Agent Identifiers and historical records remain.

The Run Identifier bridges the two lines. The Agent Identifier identifies the continuing position; the Run Identifier describes what a particular instance run of that position did and links it to its task. One Industry Compass task may call on multiple positions and contain multiple runs by the same position. A task number identifies the entire task and cannot replace the Run Identifier for each instance. The workflow line shows what happened at each stage, and the AI Position Register also shows which tasks an Agent participated in and which roles it held at the time.

This distinction is especially important for 5,000 AI employees. Department names alone do not explain who completed a delivery; a task list alone does not identify who manages participating positions over time. Collaboration along two lines separately answers "Who is this position?" and "What did it do in this task?" Agent Identifiers and Run Identifiers then connect the answers accurately.

14.2 Temporary Teams Cover Four Functions, and Continuing Responsibility Records Remain Intact

A Task-Specific Team is a combination of responsibilities configured around a task; it cannot list only business execution positions. An Industry Compass requires content production, organization of dependencies, boundary checks, and maintenance of operating conditions. Every team must therefore cover the four functions of Administration, Business, Control, and Deployment according to the actual task, ensuring that results are produced, processes connected, boundaries checked, and operating conditions maintained.

Function	Work within the Task	Limits on Each Function's Authority
Administration	Organize team formation and maintain dependencies, delivery handoffs, and return arrangements	Does not form professional judgments in place of Business or perform independent Acceptance Review
Business	Organize materials, conduct industry analysis, make scoring judgments, and deliver component outputs	Cannot perform final validation of its own results alone
Control	Check permissions, evidence, rules, and risk boundaries, and retain inspection opinions	Does not withdraw objections or lower inspection conditions to meet delivery schedules
Deployment	Provide models, data, tools, operating environments, and technical support	System permissions do not confer authority to make business judgments or modify rules

Table 14-1: The Responsibilities of Four Functions in Dynamic Collaboration

The four functions may proceed in parallel within the same stage. As Business develops analysis and Scoring Rule designs, Control checks evidence and rule boundaries, Deployment handles data and tool conditions, and Administration maintains inputs, outputs, and handoffs. Function coverage is a responsibility requirement; it does not prescribe a fixed number of positions. Whether roles may be combined and when they must be separated remain subject to formal rules. Two fundamental requirements must hold: result providers do not replace independent Acceptance Review, and technical feasibility does not replace Business and Control judgments.

Temporary team formation changes current participation relationships, but not positions' continuing identities. Positions retain their Agent Identifiers, departmental affiliations, and accountability assignments after joining a team. Each Agent Instance run generates a Run Identifier linked to its task, current role, permissions, inputs, outputs, and inspection results. If a task involves rework, reexamination, or version changes, each run is recorded separately, rather than collapsed into a result that obscures the sequence.

A complete delivery must also have a clearly assigned overall accountable owner. A team cannot obscure the coordinating party by saying it was "completed jointly," and that party cannot absorb all component responsibilities into its own name. Who organized materials, formed industry judgments, submitted scoring

results, and provided inspection opinions must be recorded separately. Overall accountability identifies who organized delivery; component records explain how each judgment was formed. Together, they form a complete chain of responsibility.

Figure 14-1 summarizes the responsibility relationships among the four functions, named accountability assignments, and task progression.



Figure 14-1: Organization and Responsibilities within HAT over MASS

14.3 The Scheduling Agent Handles Routine Team Formation, Handoffs, and Overall Delivery Coordination

The Scheduling Agent keeps a task running continuously under established objectives and rules. At the start, it forms the team according to industry scope, delivery boundaries, applicable versions, and required functions, confirming each member's role for the current task. Admission to the team depends on department and responsibilities, as well as capability and performance profiles, availability, required permissions, and suitability for the current task.

Once the team is formed, the Scheduling Agent specifies what each member receives and delivers, which permissions it may use, and which checks its results must pass. For work with dependencies, the preceding output must become usable before the next activity can begin. Whenever the Rule Version, input materials, or delivery standards change, coordination must identify affected stages to prevent results from different versions being mixed during Deliverable Consolidation.

During execution, the Scheduling Agent continually maintains task relationships: tracking the status and destination of component work, filling handoff gaps, and organizing clarification and reexamination within the rules. When questions arise, it first identifies causes, then clarifies the evidence. Only after confirming that content requires revision does it return work to the corresponding position, accompanied by the reason for failure, version used, and scope of impact. Revision produces new run records, after which the work reenters the process and continues to undergo the appropriate checks.

Coordination authority cannot expand into judgment authority. The Scheduling Agent may request supplementary evidence, organize reexamination of differing opinions, and suspend work whose prerequisites remain unmet. It has no authority to change Governance Hats, expand position permissions, or modify Scoring Specifications on its own. It organizes component results into a complete delivery but cannot rewrite specialist judgments, bypass independent Acceptance Review to meet a deadline, or waive Control checks.

14.4 Named Humans Handle Rules, Authorization, and Escalated Matters Outside the Rules

Every Agent is assigned to a named human who assumes accountability. The assignment places responsibility for rules, authorization, and matters outside the rules with an identifiable party; it does not require the human to perform each operation in the Agent's place. The accountable owner must know the rules under which the Agent works, the resources it may call on, the circumstances requiring it to stop execution and escalate, and the tasks affected by rule changes.

Routine tasks already covered by rules and sufficient permissions are entrusted to Agents for completion under those rules. The Scheduling Agent organizes teams and handoffs, Business positions produce component results, and Control and Acceptance Review positions perform inspections. If the human accountable owner were to approve every material-organizing, analysis, and scoring activity anew, large-scale collaboration would revert to human direction, erasing the Agent's operational boundaries within the rules.

Scoring here means calculation and scoring proposals within the rules. Final enterprise scores that take effect externally, compliance approvals, investment judgments, and suitability judgments are still made or confirmed by humans with the appropriate authorization and qualifications; internal inspection cannot replace final approval.

Humans primarily handle matters outside the rules. When a new business situation exceeds existing rules, position permissions are insufficient, or multiple compliance requirements require discretion, the Agent should stop extending its actions on its own and submit the facts, existing rules, disagreements, and potential effects together. The named human then decides whether to continue, revise, or terminate, and records the decision in the task and relevant run records.

Accountability assignment also requires the human to remain informed about rule execution, records, and backtesting. Being "informed" does not mean taking over each item; it means assessing whether existing boundaries remain effective and whether repeated failures, rule circumvention, or environmental changes have arisen. Humans maintain boundaries that ensure every Agent has an accountable owner, rules can be explained, and exceptions can be escalated.

Stopping and resuming actions with significant consequences should follow the handling sequence in Section 7.2 of Chapter 7. The organization's formal authorization table should identify the persons responsible for handling the matter and granting approval. Accountability assignments, professional judgment authority, stop authority, and resumption approval authority should be specified separately.

14.5 Overall Delivery and Component Responsibilities Coexist, with Acceptance Review and Control Remaining Independent

An Industry Compass deliverable contains many component judgments. The Scheduling Agent's organization of overall delivery does not eliminate the responsibilities of materials, analysis, scoring, and inspection positions. Each position remains responsible for its inputs, judgments, results, and applicable rules; coordination records explain how these results were connected, clarified or revised, and returned to the deliverable.

Named human responsibilities are also separated by level. A human is responsible for the rules, authorization boundaries, and escalated matters of the Agent assigned to them, but one accountability-assignment record cannot replace the current task's run records. A delivery must identify the coordinating party, participating positions, each instance run's Run Identifier, its task, applicable versions, inspection opinions, and matters escalated to humans. Routine execution, overall coordination, and final discretion can then each be assigned appropriately.

The relative independence of Control is a necessary limit on coordination authority. The Scheduling Agent may require inspection opinions to identify specific rules, evidence, or risks and may organize supplementary materials and reexecution, but cannot remove unresolved objections because of schedule pressure. Execution positions may explain their judgments but cannot declare acceptance on their own; Deployment may demonstrate that systems or tools operated as required but cannot replace business or compliance conclusions.

The responsibilities and boundaries of AI Control are set out in Section 11.3 of Chapter 11. Independent quality inspection, compliance or legal judgments, and approvals that take effect externally must each retain their own evidence and responsible parties.

Independent Acceptance Review also cannot be performed entirely by the result provider. Acceptance Review uses preset standards to determine whether a result may enter the next stage; questions are clarified first, and work returns to the corresponding position only when revision is clearly needed. The next chapter develops the Six Review Perspectives and Eligibility and Composition Review, whose scope is limited to this round's candidate factor design review, eligibility determination, and composition selection. Enterprise data collection, actual scoring, calibration, deployment, and publication belong to subsequent implementation. Eligibility categories are Excluded, Pending Clarification, Pending Revision, and Eligible. Eligible does not mean finally selected or enterprise-validated. The authority boundaries must first be maintained: coordination is not approval, execution is not Acceptance Review, and accountability assignment does not erase component responsibilities.

14.6 Example: A Chain Convenience Store Industry Compass: Who Clarifies, Revises, and Decides When Disagreements Arise

While B-Measurement Specification is still defining the statistical object, B-Scoring already faces an obstacle: assigning a score from 1 to 3 requires knowing whether high density is good or bad, but the Three Industry Perspectives disagree:

The Ordinary Operator Perspective sees more convenience stores in the same trade area as diverting more customer traffic, so density should be interpreted as pressure.

The Industry Veteran Perspective notes that chain brands often open additional stores in established trade areas. More stores of the same brand may indicate precisely that the trade area has been proven and distribution is more consolidated.

The Adversarial Manipulation Perspective does not discuss direction, asking only this: if a trade area can be delineated after the data are seen, the same store can draw a slightly smaller boundary to exclude a few competitors. Operations remain unchanged, but the reading can be adjusted at will. Is that factor still dependable?

Three models make separate judgments for each perspective, and the results are then summarized by perspective. Within the same perspective, the three models' outputs do not converge either: one favors a competition explanation, another a scale explanation, and the third finds the definition insufficient for judgment. Disagreement itself is not the problem; what matters is where it occurs.

Locate First, Then Clarify, and Return Work Only When Revision Is Clearly Needed

A-Coordination does not vote or rule for either side. It first assigns the disagreements to specific work units: the dispute over direction belongs to factor identification, and the dispute over boundaries to measurement specification. It then asks B-Industry Analysis to describe all three interpretations as verifiable operating mechanisms: the evidence for competition, the evidence for scale, and the field targeted by the manipulation concern. Once the evidence is laid out, clarification establishes that V0.1 put two things into one name: diversion by competitors and coverage from additional stores of the same brand. Their directions may be opposite, so the combined reading cannot be assigned unambiguously to a scoring band. The work is therefore returned to B-Measurement Specification.

V0.2 limits the statistical object to the number of convenience stores not belonging to the brand within the same trade area. The brand's own store count is observed separately and is not developed in this example. It also requires the trade-area boundary to be fixed before data collection and its version retained, the store list to

identify its source and version, and the Observation Period to be consistent. No scoring direction is presumed. B-Scoring records "discriminative evidence pending" and does not invent thresholds. D-Support archives V0.1 and V0.2 separately, retaining all opinions and the Run Identifier for each instance run. The record notes that the data path for the competitor store list has not yet been connected and is an actual data gap to be filled.

One further matter is beyond the team's authority: existing rules do not cover a uniform standard for delineating trade areas. B-Measurement Specification cannot establish a new industry-wide standard, nor can A-Coordination decide on its behalf to accelerate progress. This gap is registered as an escalated matter and referred to a named human. This example does not specify an approver or approval outcome. Before human discretion is exercised, V0.2 records only the control requirement that "boundaries must be fixed in advance and recorded," without specifying a radius.

After revision is completed, V0.2 is submitted to C-Six Review Perspectives and Eligibility and Composition Review. Eligibility and Composition Review defers its eligibility determination until the findings of the Six Design Reviews are available (see Section 15.6 of Chapter 15).

Stage	Simulated Role	Work Performed in This Example	Authority Boundaries (This Example Specifies Neither Approver nor Outcome)
Location	A-Coordination	Assigns disagreements to factor identification and measurement-specification units and organizes clarification	Cannot rule for either side or reach a conclusion for the inspection role
Revision	B-Industry Analysis B-Measurement Specification B-Scoring	B-Industry Analysis documents the evidence for three interpretations; B-Measurement Specification produces V0.2, limits the statistical object, requires boundaries to be fixed in advance, and retains versions of the list and period; B-Scoring records "discriminative evidence pending"	Cannot approve its own revisions or invent thresholds
Record Retention	D-Support	Retains V0.1, V0.2, all opinions, and each instance's Run Identifier	Cannot describe "saved by the system" as "accepted by the business"
Escalation	Named Human	Handles the gap in rules concerning a "uniform standard for delineating trade areas"	This example does not invent an approver or approval outcome
Reexamination	C-Six Review Perspectives and Eligibility and Composition Review	Receives V0.2 and prepares a Full Six-Perspective Review	Cannot reexamine only the item in which the problem was first found

Table 14-2: Position Responsibilities in the Simulated Chain Convenience Store Task

The team disbands with the task, but its records remain. Every run by each position in the table has a Run Identifier, and escalated matters are assigned to named humans. Who raised the problem, revised the definition, and undertook escalation can therefore be identified separately afterward. This identification relies on the actual

record of each run; accountability assignment to a named human identifies who is responsible for the position and cannot replace those records.

Chapter 15 How a Deliverable Is Inspected and Revised [Operating Procedures]

Chapter Overview | This chapter answers: How is a deliverable inspected and revised?

Key rules:

- Acceptance Review must use standards agreed in advance and specify the object inspected, conclusion, and next stage permitted;
- The Six Review Perspectives separately examine candidate designs; Eligibility and Composition Review determines eligibility and composition from their findings. Completing the Six Design Reviews does not mean passing, Eligible does not mean selected, and selected does not mean enterprise validation is complete;
- Unresolved matters must identify the gaps and arrangements for handling them and be archived with versions and review opinions;
- Substantive changes in definitions, purposes, or Scoring Rules require a new version and a repeated Full Six-Perspective Review;
- When subsequent facts reveal discrepancies, first distinguish causes in materials, capabilities, and rules, then make the corresponding revisions and retain historical evidence.

**General readers may stop here. The detailed text is intended for professionals such as organizers, practitioners, technical due diligence teams, and AI auditors.*

Part III makes independent inspection and End-to-End Traceability conditions for trusted delivery. This chapter further explains how those requirements are implemented in Acceptance Review, run records, and accountability tracing (see Sections 10.1 and 10.3 of Chapter 10).

Completing a candidate factor table is only half the work. It must also be checked for a clear explanation of the operating problem, calculability, data paths, and the information it adds to a composition. Only when these questions are answered can an output enter the next stage. Even after it passes the current review, new facts arising later may require us to reassess it.

This chapter first defines the scope of this round's work: the Six Review Perspectives and Eligibility and Composition Review are responsible for candidate factor design review, eligibility determination, and composition selection. Enterprise data collection, actual scoring, calibration, deployment, and publication belong to subsequent implementation. Later discussions of backtesting and actual score checks explain how deliverables undergo inspection in subsequent use; they do not assign all implementation responsibilities to the Six Review Perspectives and Eligibility and Composition Review.

15.1 Acceptance Review Is the Admission Gate for Task Progression

Acceptance Review first answers a basic question: has the submitted work completed what was originally agreed? Task objectives, industry scope, input materials, Rule Versions, permissions, and delivery standards all provide the basis for judgment. These conditions should be clarified before work begins, rather than adjusting standards to fit results after they appear.

Factor identification, measurement specification, data-path assessment, and Scoring Rules produce different deliverables, but share inspection requirements: they must respond to the original task, use applicable rules, provide clear definitions and verifiable evidence, and allow them to be reviewed. If upstream materials or rules have changed but downstream work continues using an old version, even a complete table cannot pass Acceptance Review under the new conditions.

Acceptance Review must also examine how a result was produced. Records must identify the inputs, knowledge, models, and tools used, whether work remained within authorization, and whether prescribed checks were completed. Executors may explain their own results, but independent inspection must still be performed by positions carrying the corresponding responsibility. The capability to generate results does not confer approval authority.

Acceptance Review conclusions must be retained with the task, output version, applicable rules, and inspection records, and specify which stage the output may enter. Entry into candidate composition, formation of deliverables, and archiving each have their own conditions. When conditions are unmet, the problem and evidence are explained first, followed by a determination of clarification, revision, or exclusion. Where work must be returned, the modification scope and resubmission requirements are listed. Rule changes, permission exceptions, and discretion are referred to named humans.

Subsequent stages therefore receive more than an "accepted" marker: they receive an output with a clear origin and status. Schedules may press work forward but cannot replace delivery standards.

15.2 The "6+1" Review Perspectives: Responsibilities of the Six Review Perspectives and Eligibility and Composition Review, Inspection Outputs, and Returns for Nonconformance

Whether a factor may enter a Compass cannot be decided simply by asking whether it "looks important." Operating relevance, calculability, data conditions, discriminative evidence, relationships within the full pool, and resistance to manipulation must be examined separately. Each of the Six Review Perspectives holds one inspection responsibility; Eligibility and Composition Review determines candidate eligibility from their findings, then selects compositions. The responsibilities of these seven roles must be distinguished.

Name	Inspection Question	Corresponding Output
Business Relevance Review	Is it meaningful in operating terms?	Map of Operating Problems and Mechanisms
Calculation and Scoring Specification Review	Can it be calculated and scored unambiguously?	Table of Definitions, Parameters, and Boundaries
Data Source and Access Review	Is there a credible data path?	Table of Data Paths and Conditions
Discriminative Validity Review	Is there a basis for distinguishing states?	Table of Discriminative Hypotheses and Evidence
Incremental Information Review	What information does it add relative to the full pool?	Table of Full-Pool Relationships and Mappings
Manipulation Resistance Review	Can scores be inflated without improving the state?	Table of Attacks, Defenses, and Dependencies
Eligibility and Composition Review	Which candidates are eligible, and how should they be combined?	Eligibility Register and Final Plan

Table 15-1: Inspection Questions and Outputs of the Six Review Perspectives and Eligibility and Composition Review

The Six Review Perspectives are Business Relevance Review, Calculation and Scoring Specification Review, Data Source and Access Review, Discriminative Validity Review, Incremental Information Review, and Manipulation Resistance Review. Six Design Reviews is the collective name for these six review activities. Eligibility and Composition Review subsequently determines eligibility and composition; it is not a "seventh review perspective."

Business Relevance Review first brings a factor back to the business itself: which operating problem does it address, and through what mechanism does it affect results? However technical its name, it must explain this connection. The Map of Operating Problems and Mechanisms develops the connection for subsequent discussion of definitions, data, and scoring.

Calculation and Scoring Specification Review must ensure that the same materials yield only one calculated result and one scoring band. Definitions, numerators, denominators, units, Observation Periods, and outlier

treatment must align. Scoring-range endpoints must neither overlap nor leave gaps, and treatment of decimals must be specified. When denominators such as net profit are zero or negative, explicit rules are required; bands for positive values cannot be applied directly. The Table of Definitions, Parameters, and Boundaries must be sufficiently complete and unambiguous.

Data Source and Access Review asks where data come from, how they can be obtained, and what scope they cover. However important a factor, it lacks a calculation basis if sources are unclear or collection conditions are unavailable. The Table of Data Paths and Conditions must bring sources, access conditions, and limitations together. Listing a web address does not establish a data path.

Discriminative Validity Review checks whether changes in values correspond to meaningful operating differences. Fluctuating data show only that something is changing, not necessarily that the data distinguish good from bad. Therefore, the Table of Discriminative Hypotheses and Evidence must separate what we think the factor explains, how far existing facts support it, and what remains to be validated.

Incremental Information Review returns a candidate factor to the full pool and asks what it actually adds. The Table of Full-Pool Relationships and Mappings specifies, item by item, whether the candidate and existing factors have relationships of duplication, inclusion, dependency, or complementarity, helping compositions retain information that is truly additional.

Manipulation Resistance Review tries to exploit gaps in the rules: can scores be inflated merely by changing the statistical scope, recording method, or dependencies without improving operating conditions? Identifying these paths in advance reveals where defenses are needed. The Table of Attacks, Defenses, and Dependencies records score-inflation methods, blocking conditions, and the other rules on which defenses depend.

Eligibility and Composition Review receives the complete findings of the Six Review Perspectives and first classifies candidates as Excluded, Pending Clarification, Pending Revision, or Eligible. Candidates lacking eligibility are excluded; unclear materials or understanding are clarified first; designs needing adjustment are Pending Revision; and candidates meeting this round's eligibility conditions are Eligible. The Eligibility Register retains evidence and disagreements. The final plan separately explains which candidates were selected from the Eligible Pool, how they were combined, and why. Eligible does not mean finally selected, and finally selected does not mean enterprise validation is complete.

Figure 15-1 connects candidate formation, inspection by the Six Review Perspectives, and selection by Eligibility and Composition Review, making it easier to trace how different opinions enter the final plan.




Figure 15-1: Multi-Perspective Collaboration and Design Review for an Industry Compass

When disagreements arise, facts, definitions, and all opinions are checked first; work cannot be returned for rework simply because a disagreement is found. If clarification confirms that changes are needed, work is returned to the corresponding unit with the problem, evidence, original version, and resubmission conditions recorded. A substantive change in a candidate's definition, purpose, or Scoring Rules requires a new version and a repeated Full Six-Perspective Review; reexamining only the item in which the problem was first found is insufficient. Original opinions, revisions, and the new round's conclusions must all be retained, after which Eligibility and Composition Review updates eligibility and reexamines the composition.

When new industry circumstances, permission exceptions, or changes to the rules themselves arise, the Six Review Perspectives and Eligibility and Composition Review should explain the dispute and evidence and escalate them to a named human for discretion. Professional inspection must be able to raise challenges, and final discretion must also remain within established authorization.

15.3 Checking Inputs, Outputs, Rules, and Responsibility through Run Records

Acceptance Review establishes where a result may proceed; traceability explains how it reached this point. Both must refer to the same records. Each Agent Instance run generates a corresponding Run Identifier linked to its task, inputs, outputs, calls, Rule Version, inspection results, and time information. A task may contain multiple instance runs; its task number cannot be used as the traceability identifier for a single run.

To review a result, first locate the corresponding run, then examine its input materials, definitions and Scoring Rules at the time, intermediate outputs, inspection opinions, and any clarification, return, and rework. Subsequent material updates and rule changes should not overwrite that run's original conditions, or its historical conclusion will become impossible to explain.

Run records connect upward to stable AI Organizational Positions. The Agent Identifier (Agent ID) identifies the position that participated; the Agent Capability and Performance Profile explains the conditions on which it was called on at the time. Retaining continuing identities separately from individual executions makes tasks and responsibility traceable even after a team disbands.

Position records should also identify human accountability assignments and the current authorization. The authority of participants, inspectors, task leads, and ultimately responsible persons must be checked separately. Appearing on a list does not confer approval authority, and responsibility for coordination does not confer authority to change rules. Traceability must answer who made which decision under what permissions.

Status must also have corresponding evidence. "Completed" should be supported by delivery and Deliverable Consolidation records; "Returned" should show the problem and resubmission conditions; "Pending" should explain the obstacle and who will handle it. Current status may be updated, but earlier objections, returns, and version changes must still be retained.

A Run Identifier is only a lookup point. If input, call, or inspection records have not been retained, even a complete set of identifiers cannot reconstruct the process. Claims should reach only as far as the records support, and missing elements must be identified explicitly. An "Acceptance Review passed" marker cannot fill an evidence gap.

15.4 Backtesting against Facts Drives Capability Calibration and Rule Revision

Passing candidate design review establishes only that the candidate met the eligibility and composition conditions at that time. Subsequent enterprise implementation brings new facts through data collection, actual scoring, and calibration, and earlier hypotheses may encounter counterexamples. At this stage, backtesting against facts examines whether delivered judgments still explain actual operations.

Backtesting first fixes the object and version, then matches factual sources, sample scope, Observation Period, original rules, and original conclusions. If conditions differ, differences between old and new results cannot immediately be treated as support or refutation. Backtesting frequency, samples, and thresholds should have appropriate professional justification; one set of numbers cannot cover all industries.

When discrepancies are found, their causes are investigated first. Changes in the original factual conditions require input updates; available information that was omitted requires inspection of Agent capabilities and call conditions; discrepancies despite fully compliant execution may indicate a gap in the rules. Without distinguishing causes, changes can easily address the wrong place.

The different types of causes require separate handling. Updating materials, calibrating capabilities, and modifying rules cannot substitute for one another. Rule changes must explain their reasons, old and new versions, effective dates, and scope of impact, and be checked against outputs dependent on the old rules. A substantive change in a candidate's definition, purpose, or Scoring Rules should undergo the Six Design Reviews again under a new version before eligibility and composition are updated. Enterprise implementation then addresses affected scores accordingly.

After changes, comparable cases must be used to check whether the original problem has been blocked and whether new blind spots have arisen. Input requirements, capability configurations, rules, inspection opinions, and results before and after adjustment must be retained together to show what the changes actually resolved.

Validated experience then enters the next round of tasks to calibrate capability and performance profiles, task templates, and inspection requirements. Judgments still in dispute retain their evidence and boundaries. Historical performance may support allocation but cannot confer exemption from inspection. Every run remains subject to inspection for that execution.

15.5 Authorization Checks, Version Management, and Record Completeness

Records may contain business materials, permission information, and human information, so checks must be conducted under authorization. The white paper explains how responsibility is linked; it does not disclose sensitive tasks or personal information. Authorized inspectors should obtain the evidence needed for their checks, with public disclosure scope determined separately under the relevant rules.

Rules, capabilities, and deliverables require separate version management. Rules describe what was permitted and how to calculate and inspect at the time; capability records describe the position's call conditions at the time;

deliverables show the basis on which delivered content was formed. The three are linked, but one deliverable file's version number cannot replace all change records.

Important changes must retain their content, reasons, effective dates, applicable scope, and responsibility relationships. Reviewers can then determine whether historical results remain usable and which tasks need reexamination. Old files must be retained, and the reasons for changes must also be explained.

Record completeness is determined by examining evidence item by item: Are inputs verifiable? Are calls recorded? Have the opinions of the Six Review Perspectives and Eligibility and Composition Review been retained? Can clarification and rework be connected? Was discretion within authorization? Together, these answers determine how far traceability reaches; interface status and final files explain only part of it.

If materials were not retained, identify what is missing and which judgment is affected; if a historical version cannot be recovered, retain that limitation. Gaps also identify the retention requirements to strengthen in the next round.

Adding task experience to the continuing capability library also requires authorization and version management. Before deciding what may be reused, new knowledge, inspection methods, and risk alerts should have their sources, applicability, validation status, and sensitive-information boundaries confirmed.

15.6 Example: A Chain Convenience Store Industry Compass: Why V0.2 Is Pending Clarification after the Six Design Reviews

V0.2 changes the statistical object and narrows the purpose accordingly. This affects the factor's definition and, under this round's rules, requires submission of a new version for a Full Six-Perspective Review.

Name	Inspection Question	Opinion in This Example
Business Relevance Review	Is it meaningful in operating terms?	Clarify whether the mechanism of competitor diversion affects customer traffic, spending per customer, or product mix
Calculation and Scoring Specification Review	Can it be calculated and scored unambiguously?	The direction of the statistical object is now clear, and the requirement to fix boundaries in advance is "documented." However, the specific delineation standard awaits a decision by an authorized human following escalation, so full calculation cannot yet be reproduced
Data Source and Access Review	Is there a credible data path?	The source of the competitor store list, its update frequency, and completeness of coverage have not been established
Discriminative Validity Review	Is there a basis for distinguishing states?	There is no factual support yet for whether high or low density corresponds to good or bad operations; no direction is presumed
Incremental Information Review	What information does it add relative to the full pool?	The effects of customer diversion are already partly reflected in changes in average daily sales; the costs of overly dense store placement can be observed in the store closure rate and the time for new stores to reach operating break-even. It remains to be explained what additional information store density contributes at the "causal level"

Name	Inspection Question	Opinion in This Example
Manipulation Resistance Review	Can scores be inflated without improving the state?	Fixing boundaries in advance and recording changes reduces the scope for redrawing trade areas, but a loophole remains if list sources can be freely chosen
Eligibility and Composition Review	Which candidates are eligible, and how should they be combined?	No fundamental defect warranting exclusion has been found, but several pieces of evidence remain missing and the additional information is uncertain. The candidate does not meet the conditions for entering the Eligible Pool and is classified as Pending Clarification

Table 15-2: Six Design Review Opinions and the Eligibility and Composition Review Conclusion for V0.2 (Assumptions in This Example)

What Acceptance Review Examined

The object of this simulated Acceptance Review is the revision process itself: V0.1's problems can be located in specific fields, V0.2 specifies the changed control requirements, old and new versions correspond one-to-one, and all Six Design Review opinions are retained. It confirms that the revision process is verifiable. Classification of the candidate as Eligible by the Eligibility and Composition Review, selection into a composition, and eventual enterprise scoring are each subject to further conditions.

Deliverable Consolidation and Archiving: Record Where Work Stops

Deliverable Consolidation records three categories separately in the deliverable: established content (the statistical object and the requirement to fix boundaries), clarified matters (the dispute over direction arose from mixed definitions), and remaining gaps (the delineation standard, data path, discriminative evidence, and explanation of additional information). "Store density" leaves this task as Pending Clarification and does not enter composition discussion.

Archiving links the two versions, disagreement records, escalated matters, and Six Design Review opinions to the individual Run Identifiers under SIM-01. The team has disbanded, but the record remains. When the next task supplies the missing evidence, it can resume where the work stopped without guessing from scratch why it did not pass.

Traceability and Backtesting: Subsequent Implementation Takes Up the Next Stage of Inspection

If the candidate later obtains the missing evidence and enters enterprise implementation, backtesting continues inspection: when data and judgments differ, investigate the cause first, then revise, tracing the entire process through Run Identifiers.

The simulated sequence reaches a specific conclusion: problems must be locatable, modifications must generate versions, and Acceptance Review must state what it examined. A candidate's failure to reach the end does not mean the process failed. Precisely because its stopping point and reasons are clearly recorded, we know what should be validated, clarified, and modified next.

How 5,000 AI employees work together in their positions must ultimately be demonstrated in specific tasks and deliverables. This part explains how positions call on capabilities by task and form teams dynamically to complete industry analysis, factor definitions, and Scoring Rules, then produce deliverables through independent review and composition selection, with corresponding responsibilities assigned. The illustrative chain convenience store simulation continues across Chapters 13, 14, and 15 using the same task sheet: first decomposing the task, then locating and handling disagreements, and finally inspecting the revision and retaining records. This example deliberately follows a candidate that does not reach the end, showing where the process stops and why. This round has produced Industry Compass rules covering 500 industries; their structure and conditions of use are explained in Part II. The next part uses close readings of three categories and appendix navigation to show how these deliverables are read, allowing readers to see what was selected and preserving the boundary that enterprise data collection, actual scoring, and validation remain incomplete.

Part V Case Deliverables

This part closely examines chain convenience stores, Sichuan-Chongqing hot pot restaurants, and freshly prepared tea beverage stores to demonstrate where to begin reading an Industry Compass: first confirm the operating scope, then examine indicators, factors, and their evidence requirements along the return and risk dimensions. Convenience stores focus on frequent transactions and franchise networks; hot pot restaurants on table output and prepaid funds; tea beverage stores on cup output, channels, and cash receipts.

Even similar businesses require different observation variables.

These three cases and the 17 supplementary cases in Appendix I provide 20 windows for detailed reading, rather than a statistical sample. The cases present indicator and factor structures; they do not mean that specific enterprise data have been obtained or actual scoring completed. After understanding the structure, readers must still use the access point for deliverables covering 500 industries to check the relevant industry, Rule Version, and currently available content before assessing applicability to a specific object.

Chapter 16 Close Readings of Three Categories and How to Read the Cases

This chapter uses chain convenience stores, Sichuan-Chongqing hot pot restaurants, and freshly prepared tea beverage stores to explain how to read an Industry Compass: begin with the category introduction, examine indicators and factors under the Return Dimension and Risk Dimension, and then read a brief analysis of each dimension. Chain convenience stores continue the simulated task in Section 13.6 of Chapter 13, Section 14.6 of Chapter 14, and Section 15.6 of Chapter 15. The three categories present structures only, without developing quantitative specifications or scoring standards for each factor.

16.1 Close Readings of Three Categories

Chain Convenience Stores: Ten Indicators and Seventeen Factors

Convenience store transactions are frequent, small, and dispersed. Operating results are influenced by product mix, nighttime operations, inventory turnover, revenue reconciliation, franchise systems, and other factors, making this category well suited to illustrating how an Industry Compass is organized.

Dimension	Indicator	Factor
Return Dimension	Store Sales Scale and Growth	Average Daily Sales Items per Transaction (Average Number of Items Purchased per Transaction)
Return Dimension	Gross Profit Structure and Fresh Food Contribution	Overall Gross Margin Fresh Food Sales Share Tobacco Sales Share
Return Dimension	Nighttime Operating Contribution	Nighttime Sales Share
Return Dimension	Inventory and Working Capital Efficiency	Inventory Turnover Days
Return Dimension	Franchise System Health	Franchised Store Survival Rate Franchisee Renewal Rate
Risk Dimension	Rent and Fixed Costs	Rent-to-Sales Ratio Fixed Cost Burden
Risk Dimension	Losses and Quality Control	Fresh Food Write-Off Rate Overall Store Loss Rate
Risk Dimension	Nighttime Operating Risk	Night-Shift Attendance Rate
Risk Dimension	Revenue Authenticity and Internal Controls	POS-to-Payment Receipt Discrepancy Rate
Risk Dimension	Franchise and Expansion Risk	Store Closure Rate Time for New Stores to Reach Operating Break-Even

Table 16-1: Indicators and Factors for Chain Convenience Stores

Return Dimension: From Selling More to Sustaining Earnings

Convenience stores have small floor areas and low transaction values, accumulating returns through frequent transactions; revenue alone cannot reveal operating quality. Beyond scale, the number of items purchased in each transaction must also be considered, so average daily sales and items per transaction sit alongside one

another. Profits mainly come from high-margin categories such as fresh food, with tobacco also an important reference category; overall gross margin, fresh food sales share, and tobacco sales share together clarify the profit structure. Operating 24 hours is a feature of this format, and nighttime support for sales determines whether extended operations are justified. With many product categories and funds tied up mainly in goods, inventory turnover matters. Chain expansion mainly relies on franchising, so franchised store survival and franchisee renewal determine whether the store network can be replicated.

Risk Dimension: Costs, Losses, and Expansion

The convenience store business model is relatively uniform, and the Risk Dimension and Return Dimension correspond to one another. Rent and labor are inflexible costs, but sales fluctuate with the trade area, so the rent-to-sales ratio and fixed cost burden measure the position of the break-even line. Higher fresh food sales bring greater write-off pressure; losses and quality control are the other side of fresh food contribution. Nighttime operations depend on staffing schedules being fulfilled, and insufficient night-shift staffing increases safety, service, and cash-handling risks together. Transactions are small and dispersed, with cash and multiple payment channels coexisting. Reconciliation of checkout records with actual receipts is the most direct check for inflated or omitted revenue. The costs of franchise expansion ultimately appear in store closures, and the time for new stores to reach operating break-even shows whether the pace of expansion exceeds individual stores' cash-generating capacity. For these factors, higher risk means lower scores.

What the Table Does Not Show: Why One Candidate Is Absent

In the illustrative simulated task SIM-01 in Part IV, "store density" remained Pending Clarification because of insufficient evidence and did not enter the Eligible Pool. The simulation explains only the distinction between a candidate and a deliverable; it is not an actual generation record. The effects of trade-area competition can be observed through other factors, such as average daily sales, the store closure rate, and the time for new stores to reach operating break-even.

Sichuan-Chongqing Hot Pot Restaurants: Twelve Indicators and Sixteen Factors

Sichuan-Chongqing hot pot restaurants feature Sichuan-Chongqing spicy broth bases and freshly sliced ingredients. Returns depend on table turnover, spending per person, and gross margin on broth bases. Group buying, delivery, and stored-value programs are important means of promotion and cash-flow improvement. Core concerns include price fluctuations in main ingredients, inflexible rent and labor costs, and the authenticity of table revenue.

Dimension	Indicator	Factor
Return Dimension	Table Turnover and Table Capacity	Table Turnover Rate
Return Dimension	Transaction Value and Gross Profit Structure	Spending per Person Broth Base Gross Margin
Return Dimension	Customer Acquisition and Channels	Group-Buying Coupon Redemption Rate
Return Dimension	Repeat Purchases and Membership	Repeat Purchase Rate Member Spending Share
Return Dimension	Procurement and Cost Efficiency	Main Ingredient Procurement Price Differential Overall Store Loss Rate
Return Dimension	Growth and Cash Receipts	Sales Stability
Risk Dimension	Advance Receipts and Refunds	Stored-Value-Balance-to-Monthly-Revenue Multiple

Dimension	Indicator	Factor
Risk Dimension	Platform Dependence and Settlement Periods	Revenue Share from Group-Buying and Delivery Channels
Risk Dimension	Main Ingredient Prices and Supply	Top Three Supplier Concentration
Risk Dimension	Rent and Fixed Costs	Rent-to-Sales Ratio Inflexible Cost Ratio
Risk Dimension	Compliance and Traceability	Main Ingredient Traceability Completeness
Risk Dimension	Consumption Reconciliation and Cash Verification	Cash Flow Quality

Table 16-2: Indicators and Factors for Sichuan-Chongqing Hot Pot Restaurants

Return Dimension: Table Output and Customer Loyalty

Hot pot is a capital-intensive dine-in business with a generally fixed number of tables. Its return ceiling is mainly determined by table turnover, which can be combined with spending per person to estimate revenue per table, and with broth base gross margin to roughly infer actual returns. Group buying and stored-value programs are the main means of attracting customers, but only redemption turns them into actual revenue. Repeat purchases and member spending determine whether customer acquisition costs can be spread across more purchases. Main ingredient procurement and losses of freshly sliced ingredients directly erode gross profit, so procurement and cost efficiency belong to the Return Dimension. Sales stability identifies short-lived prosperity dependent on holidays and promotions.

Risk Dimension: Advance Receipts, Supply, Fixed Costs, and Cash Verification

Hot pot risks center on advance receipts, channel concentration, supply fluctuations, inflexible costs, compliance and traceability, and revenue authenticity. Group buying and stored-value programs generate substantial prepaid funds; a higher value of this multiple means greater pressure to fulfill obligations. An excessive revenue share from group buying and delivery weakens pricing power and subjects the business to platform settlement periods. Main ingredient prices fluctuate sharply, and supplier concentration magnifies the impact of price increases and supply interruptions. Inflexible costs such as rent and labor increase pressure on earnings when customer traffic falls. Main ingredient traceability completeness verifies ingredient sources and compliance requirements. Dine-in revenue mainly uses cash and multiple payment channels, so cash flow quality must be assessed to verify booked revenue against cash receipts and records from multiple payment channels. For these factors, higher risk means lower scores.

Freshly Prepared Tea Beverage Stores: Eleven Indicators and Seventeen Factors

Freshly prepared tea beverage stores mainly offer freshly prepared drinks such as milk tea and fruit tea. Frequent purchases, rapid cup output, and add-on sales increase individual store output. Delivery platforms and membership mini programs are the main channels. Coupon-code redemption, stored-value liabilities, ingredient losses, and the authenticity of cup output require attention.

Dimension	Indicator	Factor
Return Dimension	Single-Store Cup Volume and Output Scale	Average Daily Paid Cup Output Same-Store Net Revenue Growth Rate
Return Dimension	Transaction Value and Product Value	Average Transaction Value Received Gross Margin per Cup
Return Dimension	Member Repeat Purchases and Customer Acquisition Efficiency	90-Day Member Repeat Purchase Rate Member Spending Share

Dimension	Indicator	Factor
Return Dimension	Channel and Contract Revenue Quality	Overall Delivery Platform Receipt Rate
Return Dimension	Store Unit Economics	New Store Cash Investment Payback Period
Risk Dimension	Revenue Inflation and Output Authenticity	Share of Cup Output in Periods Exceeding Capacity Cup Output-Ingredient Reconciliation Deviation Rate
Risk Dimension	Coupon-Code and Refund Risk	Refund Rate Instant/Expedited Refund Rate
Risk Dimension	Stored-Value Risk and Repayment	Stored-Value Liability Coverage Ratio
Risk Dimension	Ingredient and Loss Risk	Ingredient Cost Ratio Overall Store Loss Rate
Risk Dimension	Settlement and Fund Authenticity	Discrepancy Rate in Reconciliation of Channel Receipt Transactions with Booked Revenue
Risk Dimension	Compliance and Operating Requirements	License and Permit Completeness Rate

Table 16-3: Indicators and Factors for Freshly Prepared Tea Beverage Stores

Return Dimension: Cup Output, Transaction Value, and Single-Store Returns

Freshly prepared tea beverages have low unit prices and mainly rely on high cup output to establish single-store revenue. Same-store net revenue growth rate can reflect stores' future growth potential. Add-ons can increase transaction value, but ingredient costs rise at the same time, so average transaction value received and gross margin per cup must be considered together. Tea beverages are purchased frequently, and 90-day repeat purchases and member spending share reflect customer loyalty. Delivery platform commissions and subsidies create a large gap between nominal revenue and actual receipts, making the overall receipt rate relevant. New store investment payback periods determine whether the single-store model can support expansion.

Risk Dimension: Cup Output Authenticity and Operating Requirements

Freshly prepared tea beverage risks focus on revenue authenticity and prepaid funds. Cup output may be overstated; output during periods exceeding capacity and its reconciliation with ingredient consumption provide physical anchors for identifying fake orders. Refunds and instant refunds are typical traces of fake orders and coupon abuse. Member stored value creates liabilities whose coverage capacity must be examined. Ingredient costs and overall store losses determine whether gross profit can be maintained. Settlement through multiple platforms can disconnect receipt transactions from booked revenue, requiring reconciliation channel by channel. Licenses and permits are fundamental operating requirements. For these factors, higher risk means lower scores.

All three cases fall within food service and retail, but have different indicator and factor compositions: convenience stores emphasize nighttime operations and franchise systems; hot pot restaurants emphasize table capacity and fulfillment of prepaid obligations; tea beverage stores emphasize cup output authenticity and coupon-code refunds.

16.2 Selecting and Reading the Seventeen Cases

After the close readings of three categories, readers can continue to the 17 formal category cases in Appendix I. These cover business formats including food service and retail, accommodation and home services, mobility and space operations, content production, and computing and technical services.

Operating Scenario	Case Number
Food Service and Retail	01 Burger Fast-Food Stores; 02 Bakery Stores; 04 Discount Snack Retail; 05 Chain Pharmacy Stores; 08 Teahouse Operations
Mobility, Accommodation, and Consumer Services	03 Car Rental; 06 Residential Elder Care Facilities; 07 Economy Business Hotels; 09 Shared Massage Chair Operations; 10 Comprehensive Luxury Goods Care; 11 In-Home Pet Care Services
Instant Retail and Content Services	12 Small Immersive Performance Venues; 13 Instant Retail Fulfillment Warehouses; 14 Esports Event Services; 15 Short-Form Drama Production
Computing and Technical Services	16 AI Computing Center Hosting; 17 AI Model Deployment and Operations

Table 16-4: Operating Scenarios of the 17 Cases in Appendix I

The 17 cases in Appendix I are illustrative samples, with the content and number of indicators and factors determined by category. For example, burger fast-food stores use monthly sales per unit of floor area and average daily orders to measure location selection and customer traffic capacity; car rental uses rental utilization and residual value realization to measure asset output and exit; residential elder care facilities use paid occupancy and the caregiver-to-resident ratio to measure monetization and care capacity; AI computing center hosting uses Graphics Processing Unit (GPU) billable utilization and accounts receivable turnover days to measure resource monetization and collections.

When reading, first confirm the category scope according to business type, then examine the indicators in the Return Dimension and Risk Dimension and the factors supporting judgments under each indicator. Appendix I provides only indicator and factor structures. Quantitative specifications, data sources, units, Observation Periods, weights, and standards on a 1-to-3 scale must be defined in each category's formal Rule Version before scoring.

Part VI Future Outlook

Chapter 17 From This Round's Deliverables to the Next Delivery

Micro Connect New Markets provides Registration and Measurement Infrastructure for the contract securitization market, and HAT over MASS is the factory that manufactures and maintains the six instruments. Beginning with business problems, this white paper explains the tools' purposes, the factory's theory, workshop operations, how to read the cases, and governance relationships subject to continuing inspection. 8·03 established admission standards; 9·03 established measurement rules; 10·03 further discloses who manufactures the instruments and how they are inspected and maintained. The next step is to validate the rules already developed against specific facts and through continuous operations.

17.1 This Round's Theory, Practice, and Deliverables

In theory, we propose HAT over MASS. HAT leaves rules, authorization, and ultimate responsibility with humans; MASS places full model capabilities into clearly bounded Microtasks. Collaboration, independent inspection, and the Twelve Sufficiency Conditions establish the conditions for scale. The resulting judgment is that expanding AI use requires more than capability and numbers; tasks, inspection, and responsibility must be able to expand together.

In practice, 5,000 AI employees are already in position. Stable positions, Dynamic Team Formation, four functions, and the Seven-Stage Task Workflow translate theoretical requirements into organizational arrangements. The illustrative chain convenience store simulation further shows that clarification, revision, and Acceptance Review should each retain evidence. Acceptance of revision records does not make a candidate Eligible, much less complete enterprise validation. Evidence of actual operations still needs to be disclosed task by task.

In terms of deliverables, Industry Compass rules now cover 500 industries. Through three cases in the main text and 17 in the appendix, this white paper presents indicator and factor structures for 20 categories in total, grounding industry judgments in operating variables that can be questioned individually. This establishes a first-round basis for discussion and application in Micro Connect New Markets. Enterprise data collection, scoring, backtesting, and Cross-Industry Calibration remain subject to the explanations supported by specific versions and evidence.

17.2 Enterprise Data Collection, Calibration, Simulation, and Progress Disclosure

The first step in the next phase is enterprise data collection and actual scoring. First confirm the operating object, applicable category, and Rule Version; check quantitative specifications, sources, periods, weights, and scoring bands; then obtain and verify evidence. Fields lacking materials remain blank. Facts cannot be filled in simply to deliver a complete score.

The second step is backtesting and calibration. Factors, thresholds, and readings with explicit versions are matched to actual operating results and results of the performance of obligations, distinguishing input changes, execution deviations, and rule problems. Each revision retains evidence, old and new versions, and scope of impact. The NUMA Adjustment Factor and cross-industry comparisons are also explained according to actual validation progress.

The third step is continuous operation in a simulated environment. Clearly labeled simulated contracts and funds connect registration, measurement, task handoffs, inspection, early warnings, and handling, with adverse scenarios added. The simulation must explain assumptions, limitations, and incomplete stages; its results are not actual trading performance.

The fourth step is continued progress disclosure. Micro Connect New Markets will report its development to the market on the 3rd of each month, explaining what has been completed, how far evidence supports the findings, and which gaps remain according to actual research and validation progress, retaining version and revision

records. Rule coverage, enterprise validation, simulated operations, and formal services each have their status explained separately. The specific scope of access is subject to the corresponding disclosures.

17.3 Continuous Operation in a Simulated Environment

With a map, Compasses, and a team in place, readers naturally ask what happens when they all operate together. As contracts continue entering, operating conditions change daily, and risks gradually emerge, can this AI army detect problems promptly, issue warnings, and handle them under the rules? These questions require inspection in an environment of continuous operation.

We plan to use simulated contracts and simulated funds to put the actual system and working discipline into operation first, connecting contract entry, measurement, simulated listing, continuous monitoring, early warnings, and handling. Listing and fund movements here are both simulated arrangements; they do not indicate that the platform already provides actual issuance, trading, or custody services. We will first check whether task handoffs work smoothly, rules can be executed, exceptions can be handled, and every reading can be traced to evidence and responsible parties. Subsequent arrangements will then be determined according to actual business functions, licenses, and conditions.

Internally, this simulation is a system inspection; externally, it is a cockpit participants can test-drive themselves. As development and validation progress, participants will be able to enter conditions, adjust parameters, try different combinations, and observe how the system operates under different business scenarios. Industry and technology peers can also inspect task records to assess which rules work, which data are difficult to obtain, and which judgments still need clarification. The specific implementation scope and access arrangements will be explained once the corresponding conditions are in place.

Alongside normal operations, the simulation will include adverse scenarios such as concentrated defaults, funding stress, and fraud to check whether warnings are timely, handling follows the rules, and the system exposes shortcomings under pressure. The simulated nature will be clearly labeled to ensure that assumptions and the evidence supporting their revision are available for inspection. Trust must be built through a process that is visible, verifiable, and experienced firsthand.

17.4 Industry Feedback and Joint Inspection

Where do your business's returns come from? Where are its risks hidden?

Begin with the three cases in the main text and 17 in the appendix, and use the index of Industry Compass deliverables for 500 industries to find the Compass relevant to your operations. Feedback on indicators and factors is also welcome: which variables capture the essentials of the business, which definitions still need clarification, and which data are difficult to obtain? The 500 Industry Compasses are a starting point for this round of shared discussion. Feedback must return to specific industries, fields, and versions to continue improving the rules.

The invitation left by 9·03 will also take more concrete form through the next phase of simulated exploration:

Your Order. Your Portfolio. Your SHARP.

The map is gradually unfolding, the Compasses continue to be calibrated, and the army has entered the organization. Next, the registration and measurement instruments must collaborate within the same system, turning every reading trajectory, exception response, and revision into inspectable records. Rule coverage, enterprise validation, simulated operations, and formal services each have their status explained separately. Every step forward requires evidence and responsibility to keep pace. Micro Connect New Markets looks forward to moving toward the next delivery with you.

The Industry Compass website is live: <https://microconnect.com/industry-compass>. Click to enter, or scan the QR code after the Executive Summary.

Appendix I Seventeen Category Cases and Navigation to Research Deliverables for 500 Industries

Section 16.1 of Chapter 16 provides close readings of chain convenience stores, Sichuan-Chongqing hot pot restaurants, and freshly prepared tea beverage stores. This appendix contains the other 17 category cases, separately numbered 01–17; all indexes below refer to this appendix.

Case Scope and Reading Guide

Industry Compass research results now cover 500 industries. This white paper selects 20 categories: three are examined in detail in Section 16.1 of Chapter 16, and the other seventeen appear in this appendix. Each presents a category introduction and the indicators and factors under its Return Dimension and Risk Dimension. Case numbers support searches within this white paper and are distinct from industry identifiers in the 500-industry directory.

Each case uses the dimension–indicator–factor structure. Each dimension has 3 to 9 indicators and, in total, 3 to 9 factors within that dimension, with at least one factor corresponding to each indicator; specific selections vary with category characteristics. The stronger the return support, the higher the factor score; the higher the risk, the lower the factor score. The two sides are observed separately, and return support does not offset risk exposure.

Limits on Use

This appendix presents the indicator and factor structures of 17 categories. It does not specify quantitative definitions, data sources, units, observation periods, weights, or the specific bands for scores of 1 to 3. Before actual application, these must be defined in the formal rule version for the corresponding category. The business scope, applicable version, and evidence must then be checked before enterprise data collection, verification, scoring, and backtesting are completed.

The 17 cases in this appendix are not a representative statistical sample. They cannot replace the complete results for 500 industries or support an inference that other industries use the same indicators or factors. Cross-industry comparisons still require normalization methods, industry-side calibration, and relevant validation evidence; comparability cannot be concluded directly from the structural tables in this appendix.

Navigation to the 500-Industry Research Results

The unified results portal displays industry identifiers, result versions, access status, and materials available for online viewing. Individual industry pages present the Compass structures, rule materials, and version notes that have been made available. Case numbers in this white paper are used only for searches within it and do not serve as industry identifiers.

500-industry status fields: whether adjustments have been made, adjustment versions, validation status, and the scope of comparability must be disclosed for each industry. The complete index and scope of access are subject to the website's disclosures. This edition does not disclose verifiable adjustment versions for individual industries, so cross-industry comparability cannot be confirmed on this basis.

Case Index

Case	Category
01	Burger Fast-Food Stores
02	Bakery Stores
03	Car Rental

Case	Category
04	Discount Snack Retail
05	Chain Pharmacy Stores
06	Residential Elder Care Facilities
07	Economy Business Hotels
08	Teahouse Operations
09	Shared Massage Chair Operations
10	Comprehensive Luxury Goods Care
11	In-Home Pet Care Services
12	Small Immersive Performance Venues
13	Instant Retail Fulfillment Warehouses
14	Esports Event Services
15	Short-Form Drama Production
16	AI Computing Center Hosting
17	AI Model Deployment and Operations

Table A1-1: Index of 17 Category Cases

01 Burger Fast-Food Stores

A standardized fast-food format centered on burgers, fried chicken, fries, and beverages. It generates returns through high-traffic locations, meal combinations, and rapid food preparation. Delivery platforms and repeat purchases by members jointly support individual store revenue. Operating quality depends on preparation efficiency, ingredient costs, and the channel mix.

Return Dimension 7 Indicators 9 Factors

Indicator	Factor
I. Location and Customer Traffic Capacity	Monthly sales per unit of floor area Average daily order volume
II. Menu Mix and Average Order Value	Average order value
III. Food Preparation Efficiency and Capacity	Average food preparation time
IV. Channel Mix and Customer Acquisition Cost	Delivery order contribution ratio
V. Supply Chain and Ingredient Costs	Ingredient cost ratio
VI. Repeat Purchases and Membership Assets	Monthly repeat purchase rate
VII. Store Cash Contribution and Returns	Store operating profit margin Store payback period

Table A1-2: Burger Fast-Food Stores — Return Dimension Indicators and Factors

Risk Dimension 7 Indicators 9 Factors

Indicator	Factor
I. Inflated Order Volume and Fictitious Transaction Risk	Share of orders placed by duplicate accounts/devices
II. Discrepancies Between Food Output and Ingredient Consumption	Abnormal overall store loss rate
III. Early Recognition of Stored-Value and Group-Buying Revenue	Share of unredeemed stored value
IV. Control of Funds	Share of receipts collected outside corporate accounts
V. Cash and Refund Risk	Cash register discrepancy rate
VI. Inventory and Near-Expiry Risk	Share of ingredients nearing expiry Ingredient inventory turnover days
VII. Channel and Contract Dependence	Share of net revenue from the largest platform Share of purchases from the largest supplier

Table A1-3: Burger Fast-Food Stores — Risk Dimension Indicators and Factors

02 Bakery Stores

A chain store format focused on freshly baked bread, cakes, and pastries, with production behind the storefront or deliveries from a central bakery. Returns come from premiums on freshly made products, members' stored-value balances, and repeat purchases through private customer channels. Particular attention is required for same-day write-offs, ingredient consumption, and revenue recognition issues arising from stored value and group buying.

Return Dimension 5 Indicators 9 Factors

Indicator	Factor
I. Store Output and Sales per Unit of Floor Area	Growth rate of sales for which payment has been received at comparable stores Sales per unit of floor area Average order value
II. Product Mix and Contribution	SKU sales activity rate
III. Repeat Purchases and Member Retention	90-day paid repeat purchase rate Share of sales contributed by members
IV. Channels and Replicability	Share of revenue from private/owned channels Share of profitable stores among stores open for at least 12 months
V. Efficiency and Fulfillment	Output per employee

Table A1-4: Bakery Stores — Return Dimension Indicators and Factors

Risk Dimension 7 Indicators 9 Factors

Indicator	Factor
I. Detection of Inflated Order Volume	Share of related orders with the same delivery address
II. Material-Balance Reconciliation	Month-end ingredient inventory count discrepancy rate
III. Prepayment Violations	Share of stored-value card top-ups recognized in full as revenue on the day of top-up Share of unredeemed stored-value balances
IV. Platform Settlement Allocation	Share of overdue settlements and operating receivables
V. Production and Sales Balance	Deviation in overall store loss rate
VI. Cash Handling Chain	Discrepancy rate in reconciliation of cash, daily settlement, and deposits Share of immediately refunded orders
VII. Cash Flow Authenticity	Rate of abnormal payments and post-period reversals

Table A1-5: Bakery Stores — Risk Dimension Indicators and Factors

03 Car Rental

A short- and long-term rental service format centered on fleet assets. Returns come from vehicle rental utilization, average daily rental charges, and rental period management. Cash returns per vehicle and residual value disposal determine capital efficiency. Security of deposits and advance receipts, authenticity of vehicle status, and reconciliation of accidents and insurance are the main sources of risk.

Return Dimension 4 Indicators 8 Factors

Indicator	Factor
I. Fleet Rental Output and Utilization Efficiency	Rental utilization rate Average daily rental charge per vehicle
II. Cash Contribution per Vehicle and Capital Efficiency	Monthly gross profit per vehicle Vehicle investment payback period
III. Contract Cash Flow and Collection Authenticity	Rental cash collection rate Consistency rate across orders, contracts, and receipts Share of repeat/renewed rental orders
IV. Asset Exit and Payment Security	Residual value realization rate

Table A1-6: Car Rental — Return Dimension Indicators and Factors

Risk Dimension 5 Indicators 9 Factors

Indicator	Factor
I. Vehicle Asset Authenticity and Exclusivity	Consistency rate between vehicle status and system registers
II. Security of Deposits and Advance Receipts	Segregation of deposits and advance receipts Share of deposits and advance receipts converted to revenue
III. Revenue and Cash Flow Authenticity	Matching rate between rental revenue and bank transaction records Share of receipts collected through personal accounts Share of rental charges from related-party transactions
IV. Accident Downtime and Insurance Reconciliation	Matching rate between insurance claims and system records
V. Residual Value and Asset Exit Risk	Residual value overstatement/deviation in net exit recoveries Rate of fraudulent rentals with uncontactable renters

Table A1-7: Car Rental — Risk Dimension Indicators and Factors

04 Discount Snack Retail

A snack chain format characterized by large packages, low prices, and high turnover. Headquarters centralizes purchasing and supplies franchise stores. Returns depend on product gross margins, inventory turnover, and the pace of network expansion. Inventory pushed onto franchisees, dependence on rebates, receivable terms, and revenue authenticity require particular verification.

Return Dimension 6 Indicators 9 Factors

Indicator	Factor
I. Store Retail Output	Average daily store sales Average monthly sales per unit of floor area
II. Product Gross Margins and Supply Chain Structure	Overall store gross margin
III. Inventory Turnover and Product Efficiency	Inventory turnover days across the entire chain SKU sales activity rate
IV. Warehousing and Distribution Fulfillment Efficiency	Delivery cost ratio per store
V. Cash Cycle and Funding Efficiency	Cash conversion cycle
VI. Network Expansion and Franchisee Quality	Net store opening rate Share of new stores meeting first-year sales targets

Table A1-8: Discount Snack Retail — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Store Expansion and Concentration Dependence	Annual store closure rate
II. Ownership of Goods and Inventory Loading	Franchise store inventory depth Franchisee return rate
III. Inventory and Near-Expiry Losses	Inventory loss rate Discrepancy rate between book and physical inventory
IV. Receivables and Payment-Term Risk	Average credit term extended to franchisees
V. Advance Receipts and Funding Dependence	Franchisee advance receipts and security deposits as a share of total liabilities
VI. Procurement Concentration and Rebate Dependence	Share of purchases from top five suppliers

Table A1-9: Discount Snack Retail — Risk Dimension Indicators and Factors

05 Chain Pharmacy Stores

A chain retail format focused on prescription drugs, nonprescription drugs, and health products. Returns come from store sales, repeat purchases by members, and chronic disease management. Medical insurance settlements and online O2O channels account for substantial shares. Attention is required for medical insurance compliance, drug expiry and inventory, staffing with licensed pharmacists, and cash flow quality.

Return Dimension 5 Indicators 9 Factors

Indicator	Factor
I. Store Sales Output Capacity	Average daily net sales per store Average order value
II. Gross Margins and Category Mix	Overall gross margin Share of sales from high-margin categories
III. Member Repeat Purchases and Chronic Disease Value	Share of sales from members Member 90-day repeat purchase rate
IV. Replication of Store Network Profitability	Same-store sales growth rate Share of profitable stores
V. Contract Revenue and Cash Realization	Net sales cash realization rate

Table A1-10: Chain Pharmacy Stores — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Inventory and Expiry Risk	Inventory turnover days
II. Compliance and Professional Operations	Number of penalties for medical insurance violations Licensed pharmacist on-duty rate
III. Cash Flow and Liquidity	Ratio of net operating cash flow to net profit Medical insurance receivable turnover days
IV. Revenue Concentration and Dependence	Share of sales covered by medical insurance
V. Expansion and Store Closures	Share of loss-making stores
VI. Data and Verification Risk	Inventory count discrepancy rate

Table A1-11: Chain Pharmacy Stores — Risk Dimension Indicators and Factors

06 Residential Elder Care Facilities

A residential institution format providing long-term accommodation, daily care, and nursing services to older adults who are dependent, partially dependent, or independent. Returns depend on paid occupancy, pricing by care level, and value-added services. Security of advance receipts, care staffing, authenticity of service records, and licensing compliance are the core risks.

Return Dimension 4 Indicators 6 Factors

Indicator	Factor
I. Bed Occupancy and Monetization Efficiency	Paid occupancy rate Monthly total service revenue per bed
II. Care Levels and Pricing Structure	Share of revenue from care fees Share of residents with moderate or severe dependency
III. Contract Charges and Cash Flow Resilience	Manageable continued-residence rate
IV. Customer Acquisition Conversion and Value-Added Services	Share of revenue from value-added services

Table A1-12: Residential Elder Care Facilities — Return Dimension Indicators and Factors

Risk Dimension 5 Indicators 7 Factors

Indicator	Factor
I. Care Capacity and Scheduling Authenticity	Caregiver-to-resident ratio
II. Service Record Completeness and Retrospective Entries	Abnormal retrospective care record entry rate
III. Security of Advance Receipts and Refunds	Coverage rate of supervision over dedicated advance-receipt accounts
IV. Medical Boundaries and Compliance Permits	Completeness of operating permits Fire safety acceptance status
V. Authenticity of Temporary Absences and Abnormal Events	Ratio of funds tied up by related parties Abnormal events

Table A1-13: Residential Elder Care Facilities — Risk Dimension Indicators and Factors

07 Economy Business Hotels

A midscale and economy hotel format serving business travelers and the general accommodation market. Returns depend on occupancy, average room rates, and direct sales and member channels. OTA commissions, rent, and the staff-to-room ratio determine the room for profit. The authenticity of room nights, recognition of stored-value revenue, and collection of amounts charged to corporate client accounts require verification.

Return Dimension 3 Indicators 8 Factors

Indicator	Factor
I. Room Revenue Quality	Occupancy rate Net average room rate
II. Channels and Members	Share of direct sales/corporate agreement channels Share of room nights from members Ratio of deferred recognition of stored value
III. Costs and Profitability	Energy cost ratio Staff-to-room ratio Rent-to-sales ratio

Table A1-14: Economy Business Hotels — Return Dimension Indicators and Factors

Risk Dimension 4 Indicators 7 Factors

Indicator	Factor
I. Room Night and Occupancy Authenticity	Rate of room nights exceeding available inventory
II. Revenue Recognition and Window Dressing	Share of stored value recognized as revenue in full in the current period Rate of revenue not reversed after refunds
III. Channel Settlement and Commissions	OTA commission omission rate Unexplained channel reconciliation discrepancy rate
IV. Corporate Account Charges and Collections	Share of uncollected corporate account charges Share of corporate account charges overdue by 90 days or more

Table A1-15: Economy Business Hotels — Risk Dimension Indicators and Factors

08 Teahouse Operations

A space-based business format centered on tea consumption, tea retail, and private rooms for social gatherings. Returns come from table-hour utilization, average order value, and members' stored value. Reservations and exclusive venue bookings lock in part of the revenue. Attention is required for internal controls over stored-value liabilities, high-value tea inventory, and receipts collected through personal accounts.

Return Dimension 5 Indicators 9 Factors

Indicator	Factor
I. Customer Acquisition and Store Visit Conversion	Group-buying voucher redemption rate Platform order gross margin Share of revenue from direct channels
II. Table and Space Utilization	Paid table-hour utilization across all operating hours
III. Order Value and Product Mix	Average order value Overall gross margin
IV. Member Stored Value and Channel Quality	Stored-value customer top-up renewal rate Stored-value spending conversion rate
V. Operating Efficiency and Revenue Realization	Same-store net revenue growth rate

Table A1-16: Teahouse Operations — Return Dimension Indicators and Factors

Risk Dimension 5 Indicators 7 Factors

Indicator	Factor
I. Revenue Recognition and Stored-Value Risk	Rate of amounts for unfulfilled services incorrectly recorded as revenue Discrepancy rate between book and actual stored-value liabilities
II. Fulfillment Authenticity and Redemption	Consistency of ingredient consumption
III. Payment Cash and Redemption Risk	Share of receipts collected through personal accounts Cash deposit discrepancy rate
IV. Food and Beverage Preparation, Ingredients, and Inventory Risk	High-value tea inventory count discrepancy rate
V. Internal Control Permissions and Compliance	Share of days with missing data

Table A1-17: Teahouse Operations — Risk Dimension Indicators and Factors

09 Shared Massage Chair Operations

An unattended business format placing massage chairs in malls, cinemas, transport hubs, and other venues, with payment per use by scanning a code. Returns depend on location traffic, usage frequency per chair, and revenue-sharing terms. Attention is required for fabricated and retrospectively entered orders, stability of location contracts, equipment operations and maintenance, and aggregation of settlement funds.

Return Dimension 5 Indicators 9 Factors

Indicator	Factor
I. Output per Location	Daily paid order count per chair Average order value based on payments received
II. Traffic and Fulfillment	Growth rate of net cash receipts at the same location Rate of equipment online and available for paid use
III. Unit Economics	Cash payback period per chair Share of locations with a positive contribution
IV. Contractual Revenue Rights	Renewal rate of expiring location contracts Rate at which receipts are retained by the operator
V. Conversion of Revenue to Cash	Period from order placement to receipt of available funds

Table A1-18: Shared Massage Chair Operations — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Revenue Authenticity	Share of offline orders entered retrospectively
II. Contract Cost Completeness	Discrepancy rate in accruals for venue rent and revenue shares Share of operating locations without valid contracts
III. Settlement and Funds	Share of overdue settlement amounts Completeness rate of fund aggregation
IV. Location Concentration and Withdrawal	Revenue concentration in the top five locations
V. Refunds and Disputes	Refund amount ratio
VI. Equipment Operations, Maintenance, and Safety	30-day recurrence rate of the same fault

Table A1-19: Shared Massage Chair Operations — Risk Dimension Indicators and Factors

10 Comprehensive Luxury Goods Care

A professional service format providing cleaning, repair, refurbishment, and maintenance for high-value luxury goods such as bags, leather goods, watches, and footwear. Returns come from high-value skilled services and repeat purchases. Particular attention is required for custody of high-value goods and responsibility tracing, authenticity of work procedures, outsourcing controls, and compensation risk.

Return Dimension 6 Indicators 7 Factors

Indicator	Factor
I. Customer Acquisition and Sales Conversion	Conversion rate of qualified inquiries to paid orders
II. Average Order Value and Pricing Structure	Share of revenue from services beyond basic care
III. Repeat Purchases and Customer Value	12-month repeat purchase rate Customer revenue retention rate
IV. Technician Capacity and Fulfillment Efficiency	Monthly output value per technician
V. Contribution per Item and Cost Control	Contribution margin per item
VI. Contract Cash Flow	Operating cash conversion rate

Table A1-20: Comprehensive Luxury Goods Care — Return Dimension Indicators and Factors

Risk Dimension 7 Indicators 8 Factors

Indicator	Factor
I. Custody of High-Value Goods and Responsibility Tracing	High-value goods inventory count discrepancy rate
II. Authorization and Work Procedure Authenticity	Coverage rate of written authorization before work begins Matching rate of chargeable procedure records
III. Outsourcing Transparency and Control	Outsourcing provider qualification verification rate
IV. Fulfillment Quality and Rework	Rework rate
V. Compensation Exposure and Fund Security	Compensation loss rate
VI. Fund Authenticity and Fraud Prevention	Rate of receipts collected through accounts outside the operating entity
VII. Personnel and Business Continuity	Revenue concentration among core technicians

Table A1-21: Comprehensive Luxury Goods Care — Risk Dimension Indicators and Factors

11 In-Home Pet Care Services

An in-home service format providing pet owners with home feeding visits, dog walking, washing and grooming, and short-term boarding. Returns depend on regional order density, repeat purchases and retention, and value-added services. Attention is required for authenticity of service fulfillment, pet safety incidents, personnel stability, and private or diverted orders.

Return Dimension 4 Indicators 5 Factors

Indicator	Factor
I. Demand Density and Order Conversion	Order density in the service area
II. Pricing and Revenue Structure	Share of revenue from value-added services
III. Fulfillment Monetization and Completed Order Quality	Customer confirmation rate for service records
IV. Repeat Purchases and Customer Retention	Monthly repeat purchase rate 180-day retention rate after customer acquisition

Table A1-22: In-Home Pet Care Services — Return Dimension Indicators and Factors

Risk Dimension 7 Indicators 7 Factors

Indicator	Factor
I. Fulfillment Authenticity and Evidence Completeness	Share of orders with fulfillment abnormalities
II. Pet Safety and Abnormality Handling	Major safety incident rate
III. Capacity Exceedance and Scheduling Conflicts	Rate of excessive orders accepted by the same individual
IV. Personnel Stability and Private Orders	Private/diverted order rate
V. Security of Advance Receipts and Cash Flow	Months of operations covered by cash
VI. Refunds, Compensation, and Customer Complaints	Compensation rate
VII. Record Authorization and Compliant Insurance	Share of orders covered by liability insurance

Table A1-23: In-Home Pet Care Services — Risk Dimension Indicators and Factors

12 Small Immersive Performance Venues

Small and medium-sized performance spaces centered on small theaters, immersive drama, and interactive performances. Returns come from ticket sales, the lifecycles of resident productions, and nonticket revenue such as exclusive venue bookings and team-building events. Attention is required for reconciliation of performances and ticket revenue, audience capacity and safety for each performance, and compliance with performance approval requirements.

Return Dimension 5 Indicators 9 Factors

Indicator	Factor
I. Content and Production Competitiveness	Audience repeat purchase rate Average lifecycle of resident productions
II. Performance Scheduling and Seat Utilization Efficiency	Seat occupancy rate Ticket revenue per seat per performance
III. Profitability per Performance and Cost Structure	Contribution margin per performance Venue rent as a share of revenue
IV. Diversified Monetization Capacity	Share of nonticket revenue Share of performances for exclusive venue bookings/team-building events
V. Contract Cash Flow and Collection Quality	Ticket settlement collection cycle

Table A1-24: Small Immersive Performance Venues — Return Dimension Indicators and Factors

Risk Dimension 4 Indicators 6 Factors

Indicator	Factor
I. Revenue Authenticity	Discrepancy rate in three-way reconciliation of performances, ticketing, and bank records Share of complimentary/internal tickets
II. Immersive Venue Capacity and Safety	Capacity exceedance ratio per performance Validity of fire safety acceptance and safety permits
III. Cash Flow and Contract Fulfillment	Buffer against a surge in ticket refunds
IV. Qualifications and Compliance	Completeness rate of performance approvals/filings

Table A1-25: Small Immersive Performance Venues — Risk Dimension Indicators and Factors

13 Instant Retail Fulfillment Warehouses

A retail format using neighborhood fulfillment warehouses to provide instant delivery within 30 minutes to 1 hour. Returns depend on order density per warehouse, average order value, and repeat purchase frequency. Fulfillment costs are relatively inflexible. Attention is required for losses and stockouts, platform dependence, discrepancies between book and physical inventory, and cash flow pressure at individual warehouses.

Return Dimension 6 Indicators 9 Factors

Indicator	Factor
I. Order Scale and Frequency	Average daily order volume per warehouse Monthly repeat purchase rate
II. Average Order Value and Gross Margin Structure	Average order value Overall gross margin
III. Fulfillment and Delivery Capacity	On-time delivery rate In-warehouse fulfillment cost per order
IV. Contract Cash Flow and Working Capital Turnover	Cash conversion cycle
V. Settlement Authenticity and Revenue Quality	Settlement reconciliation accuracy rate
VI. Warehouse Contribution and Network Replication	Share of mature warehouses with a positive contribution

Table A1-26: Instant Retail Fulfillment Warehouses — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 9 Factors

Indicator	Factor
I. Fulfillment Cost Inflexibility	Fulfillment expense ratio
II. Losses and Stockouts	Overall store loss rate Stockout rate
III. Refunds and After-Sales Service	Refund rate
IV. Platform Dependence and Cash Flow Mismatch	Share of GMV from the largest single platform Overdue accounts payable rate
V. Internal Control Fraud and Data Consistency	Discrepancy rate between book and physical inventory
VI. Liquidity and Warehouse Profit or Loss	Months of cash buffer EBITDA margin per warehouse

Table A1-27: Instant Retail Fulfillment Warehouses — Risk Dimension Indicators and Factors

14 Esports Event Services

A service format undertaking esports event planning, production, execution, and commercial operations. Returns come from organizer contracts, sponsorship and broadcasting, and ticket revenue. Project collection cycles are relatively long. Attention is required for authenticity of attendance and traffic data, revenue recognition timing, dependence on subsidies, and return flows of funds involving related parties.

Return Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Contract Cash Flow and Collection Quality	Collection rate for receivables due
II. Event Profitability and Cost Efficiency	Gross margin per event
III. Ticket Monetization and Attendance Quality	Paid seat occupancy rate Ticket revenue per paying attendee
IV. Sponsorship and Broadcasting Commercialization	Sponsor renewal rate Acceptance pass rate for sponsorship rights
V. Customer Repeat Purchases and Event Continuity	Event customer repeat purchase rate
VI. Revenue Recognition and Delivery Quality	Realization rate of advance receipts converted to revenue

Table A1-28: Esports Event Services — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Attendance and Online Traffic Authenticity	Deviation rate in verified in-person attendance counts
II. Revenue Recognition and Cash Flow Mismatch	Share of revenue from incomplete fulfillment Deviation rate between event cash collections and recognized revenue
III. Completeness and Compliance of Costs and Liabilities	Payables for prize pools and team appearance fees
IV. Dependence on Subsidies and External Funds	Government subsidies as a share of total revenue
V. Related-Party Transactions and Return Flow of Funds	Share of collections from related parties Share of funds flowing back within 30 days
VI. Evidence Chain and Data Specifications	Matters involving conflicting measurement specifications

Table A1-29: Esports Event Services — Risk Dimension Indicators and Factors

15 Short-Form Drama Production

A content production format developing, filming, and distributing web dramas, web films, and vertical short-form dramas. Returns depend on platform revenue sharing, copyright sales, and paid traffic monetization efficiency. Attention is required for paid traffic costs and recovery, production budget overruns, revenue recognition across periods, inventory impairment, and copyright compliance.

Return Dimension 5 Indicators 8 Factors

Indicator	Factor
I. Revenue-Share Recovery and Realization	Platform settlement revenue realization rate Revenue-share collection cycle
II. Paid Traffic Monetization Efficiency	Paid traffic ROI Paid traffic payback days
III. Project Profitability	Gross margin per project Project investment recovery multiple
IV. Production and Delivery Efficiency	On-time delivery rate
V. Cash Recovery Capacity	Timely platform settlement collection rate

Table A1-30: Short-Form Drama Production — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 7 Factors

Indicator	Factor
I. Paid Traffic Risk	Paid traffic cost ratio
II. Production Cost and Fulfillment Risk	Budget overrun rate per project
III. Revenue Recognition and Cross-Period Risk	Share of recognized revenue from productions not yet released
IV. Impairment and Inventory Risk	Share of completed projects without platform release schedules
V. Related-Party Transaction Risk	Share of revenue from related-party transactions Suspected return flow of revenue-share collection funds
VI. Compliance and Copyright Risk	Authorization chain completeness rate

Table A1-31: Short-Form Drama Production — Risk Dimension Indicators and Factors

16 AI Computing Center Hosting

A data center service format providing customers with GPU computing power, racks, power, and operations and maintenance hosting. Returns come from long-term contracts, billable computing utilization, and value-added services. Capital investment is substantial and payback periods are long. Attention is required for power costs, equipment generation changes, debt service, and customer concentration.

Return Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Contract Cash Flow Quality and Lock-In	Share of the value of contracts containing minimum commitment / take-or-pay clauses Undiscounted investment payback period
II. Computing Resource Monetization and Utilization	GPU billable utilization rate Effective rack rental utilization rate
III. Pricing and Unit Economics	Unit contribution margin
IV. Collections, Cash Realization, and Capital Recovery	Accounts receivable turnover days
V. Customer Mix and Renewals	Renewal rate of expiring contracts
VI. Delivery and Service Fulfillment	Rate of deliveries meeting standards

Table A1-32: AI Computing Center Hosting — Return Dimension Indicators and Factors

Risk Dimension 6 Indicators 8 Factors

Indicator	Factor
I. Power Capacity and Energy Costs	Power costs as a share of revenue
II. Equipment Reliability and Technology Generation Changes	Share of computing power from older-generation cards GPU failure rate
III. Funding Chain and Debt Service Risk	Share of equipment financed through finance leases
IV. Customer Credit and Collection Risk	Customer revenue share
V. Minimum Compliance and Qualification Requirements	Completeness of key permits
VI. Revenue Authenticity and Manipulation Prevention	Fulfillment rate Matching rate in reconciliation of bills and collections

Table A1-33: AI Computing Center Hosting — Risk Dimension Indicators and Factors

17 AI Model Deployment and Operations

A technical service format providing enterprise customers with AI application development, domain-specific large model fine-tuning, deployment, and ongoing operations and maintenance. Returns come from project delivery and billing per call. Attention is required for authenticity of call-based billing, capacity to bear computing costs, customer concentration, delivery acceptance, and related-party transaction risk.

Return Dimension 4 Indicators 9 Factors

Indicator	Factor
I. Customer Acquisition Scale and Conversion Efficiency	Number of qualified paying customers Trial-to-paid conversion rate
II. Call Monetization and Unit Economics	Call success rate Revenue per thousand calls
III. Revenue Retention and Profitability Quality	Net revenue retention rate Annual renewal rate Gross margin
IV. Cash Recovery and Contract Quality	Collection rate within the contractual period Accounts receivable turnover days

Table A1-34: AI Model Deployment and Operations — Return Dimension Indicators and Factors

Risk Dimension 4 Indicators 7 Factors

Indicator	Factor
I. Fulfillment Authenticity and Data Compliance	Reconciliation rate between call logs and bills Duplicate attribution/duplicate order rate
II. Cost Capacity and Supplier Risk	Computing/resource capacity ratio Dependence on third-party model suppliers
III. Customer Concentration and Delivery Quality	Share of revenue from the top five customers Post-acceptance refund/deduction rate
IV. Related-Party Transactions and Abnormal Collections	Share of revenue from related parties

Table A1-35: AI Model Deployment and Operations — Risk Dimension Indicators and Factors

Appendix II Overview of the Twelve Sufficiency Conditions

The "Twelve Sufficiency Conditions" are a set of joint conditions that make HAT over MASS possible. They constrain how tasks are broken down, how model foundations are configured, how organizational staffing creates checks and balances, and how outputs are inspected. The four groups are connected: if tasks are unclear, subsequent staffing and inspection have no defined object; if foundations lack diversity and cross-checking, multiple results may merely repeat the same bias; if staffing has numbers without fine-grained roles and checks and balances, greater scale makes errors more likely to spread along the collaboration chain; if inspection lacks coverage, duration, or depth, short-term correctness is difficult to translate into sustained trustworthiness.

The twelve conditions therefore cannot be treated as optional strengths, nor can outstanding performance in one condition offset the absence of another. "An alert whenever any condition is missing" does not mean reducing all existing records and results to zero. It means identifying whether the conditions currently relied upon hold. Once a condition cannot be established, reliance on the part that has not been established should stop, the gap should be identified, and verification should resume after calibration.

I. Overview of the Twelve Sufficiency Conditions

Category	Condition	Core Requirement	Checks Required When the Condition Is Insufficient	Main-Text References
Tasks	Sufficiently Small	Do only one thing at a time, with defined inputs, outputs, and acceptance criteria. A small task does not reduce the capability of the model called upon.	Whether multiple judgments are mixed into one task; whether inputs and outputs are clear; whether results can undergo acceptance review against preset criteria.	Chapter 8, Sections 8.1 and 8.3
Tasks	Sufficiently Narrow	Keep task permissions and judgment boundaries limited, without relying on ad hoc discretion during execution; escalate to a human when discretion is required.	Whether permitted and prohibited scopes are clear; whether the task requires major judgments without rules; whether escalation conditions exist.	Chapter 8, Sections 8.2 and 8.3
Tasks	Sufficiently Granular	Each output can be inspected independently and each step can be replayed, allowing errors to be isolated and located.	Whether task success or failure can still only be judged as a whole; whether intermediate steps, inputs, and outputs can be reconstructed; whether errors can be located in specific stages.	Chapter 8, Sections 8.3 and 8.4

Category	Condition	Core Requirement	Checks Required When the Condition Is Insufficient	Main-Text References
Foundations	Sufficiently Diverse	Use different foundations and suppliers to reduce dependence on a single source of capability and common-source bias.	Whether multiple results actually come from the same or highly correlated foundations; whether capability sources are excessively concentrated; whether the claimed diversity consists only of different names.	Chapter 8, Section 8.5
Foundations	Sufficiently Cross-Checked	Make parallel judgments on the same task so that disagreements become visible and enter review or escalation under the rules.	Whether parallel results offer independent judgment value; whether disagreements are visible and addressed; whether consistent results may arise from common bias.	Chapter 8, Section 8.5
Foundations	Sufficiently Cost-Effective	Configure models with different capabilities and costs according to task difficulty, matching model capability, task requirements, and resource investment.	Whether every task calls on the same capability level indiscriminately; whether a lower-cost configuration can meet acceptance criteria; whether cost savings compromise inspection requirements.	Chapter 8, Section 8.5
Staffing	Sufficiently Numerous	Numbers can cover all tasks and provide capacity for parallel execution, cross-review, and necessary substitutes.	Whether any task lacks an executor or reviewer; whether growth in numbers exceeds the capacity of rules, inspection, and accountability.	Chapter 9, Section 9.6
Staffing	Sufficiently Fine-Grained	Refine the division of work enough to identify specific responsibilities, handoffs, and stages where errors occur, so that responsibility is not assigned vaguely to a single whole.	Whether the step where a problem occurred can be determined when a task fails; whether production, acceptance review, and inspection responsibilities are mixed together; whether handoff recipients are clear.	Chapter 9, Sections 9.4 and 9.6

Category	Condition	Core Requirement	Checks Required When the Condition Is Insufficient	Main-Text References
Staffing	Sufficiently Counterbalancing	Allow different positions to create checks and balances within the organization, so that questioning, review, and dissent can interrupt systemic bias. "Counterbalancing" means the necessary tension between judgment positions, rather than creating disorderly opposition.	Whether those responsible for execution also decide acceptance conclusions; whether different roles merely repeat the same judgment; whether dissent can be recorded and addressed.	Chapter 9, Sections 9.4 and 9.6
Inspection	Sufficiently Comprehensive	Every output enters an acceptance review or inspection process, with no default exemptions.	Whether any outputs are skipped because task volume is large, results appear familiar, or past performance is good; whether inspection records can be linked to tasks.	Chapter 10, Sections 10.1 and 10.4
Inspection	Sufficiently Long	Backtesting covers a sufficiently long history, different environments, and boundary cases to distinguish actual capability, short-term luck, and environmental change.	Whether the observation period includes only one phase; whether the conclusion depends on short-term performance; whether the original answer has become outdated after environmental change.	Chapter 10, Sections 10.2 and 10.4
Inspection	Sufficiently Deep	Traceability can lead from results back to inputs, model calls, rule versions, execution paths, inspection results, and accountability records.	Whether specific inputs, rule gaps, or process stages can be located when an error occurs; whether renewed verification after revision can determine if the same type of error has been blocked.	Chapter 10, Sections 10.3 and 10.4

Table A2-1: Overview of the Twelve Sufficiency Conditions

II. Limits on Use of the Overview

This table supports quick checks and navigation to the main text. It does not replace the complete explanations of concepts, organizational relationships, and failure boundaries in each chapter. Chapter 8 explains the detailed logic of the task and foundation conditions; Chapter 9 covers staffing conditions; Chapter 10 covers inspection conditions.

When using the overview, understand the four groups of conditions as part of the same evidence chain. For example, even if a task is sufficiently small, it cannot be considered sufficiently narrow if its boundaries still rely on ad hoc discretion. Even if multiple models produce results simultaneously, numbers alone cannot establish effective cross-checking if their foundations share substantially the same underlying sources. Increasing staffing numbers does not mean organizational capacity has expanded accordingly if capacity for rule maintenance, evidence checks, and abnormality handling has not also increased. Likewise, passing a single acceptance review cannot replace ongoing backtesting, and a longer observation period cannot replace deep tracing of inputs, rules, and responsibilities.

The "Twelve Sufficiency Conditions" provide conditions for judgment, rather than fixed numerical thresholds, calculation formulas, or certification labels. Tasks differ in complexity, risk level, and evidence conditions. How to achieve "sufficiency" in each case should be based on the corresponding rules, acceptance criteria, run records, and backtesting results. The overview itself cannot prove that a task, organization, or result has met all the conditions, nor can theoretical checks and balances be presented as an untested guarantee of reliability.

When a condition does not hold, first determine the scope of its impact: whether the problem arises from task design, foundation configuration, division of staffing responsibilities, or inspection depth. Then stop relying on the affected judgments, revise the relevant rules or arrangements, and verify again with new evidence. Existing records are retained because they provide the basis for identifying bias, comparing versions, and judging whether calibration is effective. Trust does not accumulate automatically from past achievements; it must be formed anew in every task, every inspection, and every revision.

Appendix III Term Definitions and Chinese–English Correspondence

This table is arranged A–Z by English name and explains the meaning of each term in this white paper. The specific rules for the functions of registration and measurement instruments, organizational operations, inspection processes, and Compass application remain those in the corresponding chapters of the main text.

Chinese Term	Full English Name	Definition
验收	Acceptance Review	A judgment that checks whether outputs and the execution process meet requirements based on predetermined task goals, applicable rules, and delivery standards, and specifies which subsequent stage they may enter.
调度管理	Administration	The A function among the four major functions, responsible for organizing goals into executable tasks and managing task decomposition, dependencies, position selection, and collaboration handoffs. It does not replace professional judgment or independent acceptance review.
骗子视角	Adversarial Manipulation Perspective	An adversarial perspective that actively seeks pathways for packaging and manipulating data, indicators, or scores. It does not identify any specific operator as fraudulent.
智能体	Agent	A technical entity capable of calling upon models, knowledge, skills, and tools to perform work toward a given goal. In HAT over MASS, its activities are constrained by the current task, Governance Hat, and scope of authorization.
智能体能力表现	Agent Capability and Performance Profile	A record of capability and actual performance used for position searches and scheduling, including capability sources, versions, applicable conditions, and task performance. Possessing capability does not confer permission for the corresponding task.
身份编码	Agent Identifier (Agent ID)	The stable organizational identifier of an AI Organizational Position, linking its responsibilities, status, accountability assignment, and historical records. The Agent Identifier remains continuous when the same position participates in different tasks; the Run Identifier identifies a particular instance execution.
智能体实例	Agent Instance	A technical execution entity started for a specific task, carrying the goals, context, Governance Hat, permissions, capability configuration, and delivery requirements of that run.

Chinese Term	Full English Name	Definition
兵对兵	Agent-to-Agent Collaboration	A collaboration channel for task assignment, execution, handoff, and review among agents within established rules. It has no authority to change governance rules or expand authorization on its own.
AI 员工	AI Employee	The organizational designation HAT over MASS uses for AI incorporated into its positions, responsibilities, and collaboration system. Internal management must further distinguish stable positions from specific run instances.
AI 岗位	AI Organizational Position	A relatively stable identity and carrier of long-term records within the organization, holding responsibilities, status, Named Human Accountability Assignments, and historical performance. It does not disappear when a single task or instance run ends.
岗位台账	AI Position Register	A management record that continuously records AI Organizational Position identities, responsibilities, statuses, accountability assignments, and historical performance, stored in a separate layer from individual instance run information.
业务执行	Business	The B function among the four major functions, responsible for producing analysis, processing, and deliverables corresponding to task goals, and retaining the inputs, supporting evidence, and process records that support results.
问业	Business Relevance Review	A review perspective that examines whether candidate factors address meaningful operating questions, together with their mechanisms of influence and the conditions under which they hold, producing a map of operating questions and mechanisms.
量尺	Calculation and Scoring Specification Review	A review perspective that examines whether definitions, parameters, units, observation periods, and boundary handling support unambiguous calculation and scoring, producing a table of definitions, parameters, and boundaries.
候选资格状态	Candidate Eligibility Status	The four statuses a candidate factor may hold after review: Excluded, Pending Clarification, Pending Revision, or Eligible. Eligible means qualified to participate in composition selection, rather than finally selected.

Chinese Term	Full English Name	Definition
候选因子	Candidate Factor	A proposed factor that still requires design review and an eligibility judgment. Becoming a candidate does not mean entry into the final plan or completion of enterprise validation.
本文案例编号	Case Number in This White Paper	The sequential numbers assigned to the 17 illustrative cases in Appendix I, such as Case 01. These numbers are used only for searches within Appendix I and are separate from the industry identifiers in the industry directory.
首席执行官 (CEO)	Chief Executive Officer	In Figure 14-1, the organizational role responsible for coordinating goals, priorities, and resource boundaries.
同源偏差	Common-Source Bias	Common bias arising across multiple judgments because they share the same or highly correlated model foundations, materials, or methods. Agreement among multiple results cannot by itself prove that the judgments are correct.
罗盘读数	Compass Reading	Factor data updated daily for registered contracts, continuously recording changes in operations, fulfillment, and evidence, with points in time, periods, and rule versions retained. Missing or delayed data must be indicated. The readings do not constitute recommendations for investment action.
罗盘终端	Compass Terminal	A working interface through which licensed institutions or qualified overseas investors with the corresponding qualifications and authorization read registration materials, rules, and readings, and form judgments using their own methodologies. It is managed separately from the public industry research portal.
合同罗盘	Contract Compass	A framework for measuring registered contracts, consisting of the Core Compass and Industry Compass, which observe contract structure and the specific business, respectively. Each has a Return Dimension and Risk Dimension. They are presented separately, are not added together, and do not offset each other.
组合登记簿	Contract Portfolio Registry	After investors complete transactions and settlement through licensed institutions, it records the relationships and facts between contracts and portfolios daily and links to original transaction and position records. It does not handle client assets.

Chinese Term	Full English Name	Definition
合同登记簿	Contract Registrar	A record retaining contract identity, terms, initial profiles, operating status, underlying evidence, and historical versions. Registration does not create assets, change rights or obligations under the original contract, or constitute endorsement.
合同终端	Contract Terminal	An interactive system that maintains dialogue with issuers or their AI agents, actively requests, verifies, and follows up on materials, and organizes contracts into structured records that meet admission rules. Admission does not constitute an investment recommendation.
管控合规	Control	"Control" in this white paper is the function through which AI collects evidence, performs checks, flags abnormalities, and conducts design reviews and quality inspections within its authorization. Compliance approvals, investment judgments, suitability judgments, final enterprise scores that take effect externally, and legal opinions are made or confirmed and recorded by people with the corresponding authorization and qualifications. AI may still calculate and form scoring recommendations under the rules. Internal calculation, independent inspection, and final approval carry different responsibilities.
通用码	Core Codes	The indicators and scoring expressions used by the Core Compass to present return support and risk load at the contract structure level. Readings for specific projects are formed from applicable rules and materials and are distinct from the Core Compass framework itself.
通用罗盘	Core Compass	A scoring framework that observes contract structure and common elements across industries, with separate Return and Risk Dimensions.
跨行业校准	Cross-Industry Calibration	The process of establishing a basis for comparing results across industries using interpretable external benchmarks, data, and version conditions. It does not mean eliminating industry differences.

Chinese Term	Full English Name	Definition
模型对冲	Cross-Model Cross-Checking	An inspection arrangement in which different models form judgments separately under comparable task conditions, followed by checks of their assumptions, supporting evidence, and disagreements. It is not established automatically by increasing the number of models.
寻源	Data Source and Access Review	A review perspective that examines whether data sources, access conditions, coverage, and limitations are credible and clear, producing a table of data pathways and conditions.
收口	Deliverable Consolidation	The process of combining validated component outputs, addressing unresolved matters, and forming complete deliverables. It does not replace independent inspection or eliminate unresolved disagreements.
部署支撑	Deployment	The D function among the four major functions, responsible for preparing and maintaining models, data, skills, tools, systems, and operating environments, and ensuring controlled permission configurations and versions.
维度	Dimension	An upper-level direction of observation in a Compass, including the Return Dimension and Risk Dimension. Indicators are set under each dimension, with specific factors supporting judgments on the indicators. Both the Core Compass and Industry Compass retain these two dimensions.
分水	Discriminative Validity Review	A review perspective that examines whether factor differences support meaningful distinctions among different operating states, producing a table of differentiation hypotheses and evidence.
动态组队	Dynamic Team Formation	The process of selecting positions and starting instances to form temporary collaboration based on the work units, capability requirements, permissions, and availability of a specific task.
息税折旧摊销前利润 (EBITDA)	Earnings Before Interest, Taxes, Depreciation and Amortization	Used in this white paper's instant retail fulfillment warehouse case. The calculation scope and denominator of EBITDA margin per warehouse must be defined in the formal rules.

Chinese Term	Full English Name	Definition
执衡	Eligibility and Composition Review	The responsibility position that receives the complete findings of the Six Review Perspectives, determines candidate eligibility, and selects a composition from the Eligible Pool. Eligibility and Composition Review does not replace the six independent inspections or mean that enterprise validation has been completed.
资格台账	Eligibility Register	A record of candidate factor versions, eligibility statuses, grounds for judgments, disagreements, and updates, supporting subsequent composition selection and tracing.
全程溯源	End-to-End Traceability	The ability to trace results back to input materials, model and tool calls, rule versions, execution paths, inspection findings, and responsible parties. Its actual depth depends on whether evidence has been retained in full.
人工升级	Escalation to an Authorized Human	The mechanism for submitting relevant facts, evidence, and execution records to a named human for handling when a task encounters rule gaps, permission exceptions, or matters requiring discretion.
因子	Factor	A verifiable operating variable supporting judgments on an indicator. Its business meaning, data source, calculation specification, unit, observation period, and scoring rules must be defined. The Industry Compass sets 3 to 9 factors per dimension, with at least one factor corresponding to each indicator and each factor assigned to only one indicator.
因子全池	Full Factor Pool	The existing set of factors used to compare factor relationships, serving as the reference scope for checking candidate factor duplication, dependencies, and new information.
完整六审	Full Six-Perspective Review	The complete review in which a candidate factor undergoes inspection by all Six Review Perspectives. After a substantive change to its definition, use, or scoring rules, a new version should be created and undergo a Full Six-Perspective Review again.
全检	Full-Coverage Inspection	Every output enters an acceptance review or inspection process suited to its nature, with no default exemptions. Full-Coverage Inspection does not mean that every item is checked by a human or that results are necessarily correct.

Chinese Term	Full English Name	Definition
帽子	Governance Hat	The set of rules, responsibilities, permissions, and escalation arrangements applicable to an agent in a specific task, specifying permitted actions, prohibited actions, and how boundary breaches are handled. A Governance Hat may change with the task and is distinct from a model or administrative rank.
帽对兵	Governance-to-Agent Control	The governance channel between humans and agents centered on rules and authorization, used to determine permitted scope, prohibited actions, supervision requirements, and escalation conditions.
图形处理器 (GPU)	Graphics Processing Unit	Used in computing service scenarios in this white paper. The statistical scope of GPU billable utilization and failure rates is governed by the corresponding formal rules.
国内生产总值 (GDP)	Gross Domestic Product	This white paper refers to industry value-added data in quarterly GDP accounting. Industry value added and total GDP should be understood separately.
商品交易总额 (GMV)	Gross Merchandise Value	Reflects the scale of merchandise transactions within a specified period and scope. Refunds, cancellations, and the statistical scope must be defined in the formal rules. It is not directly equivalent to operating revenue.
帽子军团	HAT over MASS	An AI organizational method that places rules, authorization, and ultimate responsibility with humans and organizes large numbers of agents to collaborate through clear task boundaries. It consists jointly of the HAT trust mechanism and the MASS microtask collaboration approach.
HAT	Human-Commanded, AI-Operated, Trust	A human-commanded, AI-operated trust mechanism: humans establish rules, decide authorization, and bear ultimate responsibility; AI executes within boundaries. Human responsibility must be maintained, and AI must have sufficient scope to realize its potential within clear boundaries.
编号层	Identification Layer	The layer that connects stable positions, capability configurations, tasks, and execution evidence through Agent Identifiers, Run Identifiers, and their associated records.

Chinese Term	Full English Name	Definition
照影	Incremental Information Review	A review perspective that examines duplication, inclusion, dependencies, complementarity, and new information in candidate factors relative to the existing Full Factor Pool, producing a table of full-pool relationships and mappings.
指标	Indicator	An evaluation item under a dimension that specifies the question the dimension must answer, with specific factors supporting judgment and scoring. The Industry Compass sets 3 to 9 indicators within each dimension of a category. Their content is determined by category, rather than using a fixed list across categories.
行业码	Industry Codes	The indicators, factors, and scoring expressions used by the Industry Compass. Readings for specific projects must be interpreted in conjunction with industry scope, observation periods, rule versions, and applicable conditions. They are distinct from industry identifiers.
行业罗盘	Industry Compass	Category-specific scoring rules for a single asset category within the NUMA SHARP Twin Compasses, alongside the market-wide Core Compass. They observe the return sources, risk exposures, and category anchors of a specific business and use the dimension–indicator–factor structure.
行业编号	Industry Identifier	An identifier used to locate the corresponding industry and its rules in the industry directory. It does not represent a score and is distinct from Industry Codes.
老炮视角	Industry Veteran Perspective	A perspective drawing on long-standing industry experience to identify differences between rules on paper and actual operations, as well as industry cycles, practices, and implicit constraints.
破局	Manipulation Resistance Review	A review perspective that examines whether scores can be raised by changing specifications, inputs, or recording methods without improvement in the actual operating state, producing a table of attacks, defenses, and dependencies.
微任务智能体	Micro Agent	An agent undertaking a simple, short-duration task with limited boundaries in a single run. "Micro" describes the scope of the task in that run. It does not indicate lower model capability and is not a permanent position category.

Chinese Term	Full English Name	Definition
MASS	Micro Agents; Simple & Short Tasks	Micro Agents and simple, short-duration tasks. Full model capability operates within a limited task scope, achieving scale through verifiable microtask collaboration.
微任务	Microtask	A work unit broken down from a complex goal that retains complete business meaning, with clear inputs, a single delivery goal, limited permissions, identifiable outputs, and independent acceptance criteria. It can be individually located, corrected, or redone.
多视角席位	Multi-Perspective Role	A task position established for different observation, judgment, and inspection responsibilities, specifying the questions to address and responsibilities to bear. It is distinct from a model, tool, or source of capability.
具名真人认领	Named Human Accountability Assignment	An organizational arrangement that establishes a continuing link between an AI Organizational Position and an identifiable human responsible party. Within the scope of authorization, the person assuming accountability bears responsibility for rules, authorization, supervision, and abnormality handling. This does not require replacing AI execution item by item.
牛码调整系数	NUMA Adjustment Factor	An adjustment factor that supplements industry-side judgments with a reference to the operating environment, without being limited to a single data source. This edition primarily refers to relevant data from the National Bureau of Statistics of China and uses 1 as the neutral benchmark. The specific method, scope of application, and version are set out in Section 5.3 of Chapter 5. Results must undergo ongoing review and backtesting.
牛码新市帽子军团	NUMA HAT over MASS	The AI organizational practice established by Micro Connect New Markets using the HAT over MASS method, including position design, task grouping, and inspection arrangements. NUMA refers to Micro Connect New Markets. The organizational name used consistently in the main text is "NUMA HAT over MASS"; the general method remains "HAT over MASS."

Chinese Term	Full English Name	Definition
牛码 SHARP	NUMA SHARP	The central indicator and unified quantitative language used by Micro Connect New Markets to express and compare the risk-return structure of contract cash flows. Factors form indicator and dimension scores using disclosed weights. After normalization, the NUMA Adjustment Factor completes cross-category calibration to form industry SHARP, which then combines with the Core Compass to form NUMA SHARP for a single contract. Specific calculations must follow the rules of the applicable version.
观察期间	Observation Period	The time range covered by factor calculation and judgment. It should suit the industry's operating cycle and be retained with the result.
持续回测	Ongoing Backtesting	Ongoing factual inspection of judgments under identified versions to distinguish stable capability, short-term luck, and environmental change. Inspections using historical materials and validation against subsequent new facts should separately retain the versions, samples, and periods used.
线上线下融合 (O2O)	Online to Offline	A business approach connecting online channels with offline fulfillment of goods or services. Specific in-store and at-home scenarios are defined by the business scope.
在线旅行平台 (OTA)	Online Travel Agency	A platform providing reservation or agency services for travel products such as accommodation. In this white paper's hotel case, it is used to explain channels and commission costs.
订单登记簿	Order Register	The register in the 9·03 White Paper that records transaction terms, positions, distributions, and exit statuses for completed transactions, retaining links with the Contract Registrar and portfolio relationships.
正常经营者视角	Ordinary Operator Perspective	A perspective that starts from ordinary business operations and analyzes how factors affect revenue, costs, efficiency, and business continuity.
销售点系统 (POS)	Point of Sale	In this white paper, a cash register and sales record system, which is one data source for checking transaction records and payment receipts.

Chinese Term	Full English Name	Definition
登记与度量基础设施	Registration and Measurement Infrastructure	The functional system in the contract securitization market that connects contract facts, evidence, rules, readings, and records of portfolio relationships. It is used by licensed institutions and is responsible for its own data, systems, services, and governance duties.
职责层	Responsibilities Layer	The layer in the organization's three-layer design that determines who undertakes tasks, how positions divide responsibilities, and how they create checks and balances.
收益维度	Return Dimension	The Compass dimension that observes support for the generation and realization of returns. For return factors in the Industry Compass, stronger return support produces higher scores. The score is distinct from an actual rate of return or a promise of returns.
投资回报率 (ROI)	Return on Investment	In this white paper, the paid traffic return metric in the short-form drama case. Its revenue, cost, and period specifications are governed by the formal rules. It cannot be treated directly as the return available to investors.
风险维度	Risk Dimension	The Compass dimension that observes risk sources and exposures. For risk factors in the Industry Compass, higher risk produces lower scores. The score does not directly represent the probability of default or the expected loss rate.
规则匹配	Rule Matching	The process of confirming applicable rules, versions, permissions, and escalation conditions before task execution. Rule gaps or permission conflicts discovered must be submitted for handling.
规则版本	Rule Version	An identifiable marker of the content and revision status of rules applicable during a given period, used to explain the constraints, definitions, or scoring basis applied in an execution.
规则层	Rules Layer	The layer in the organization's three-layer design that determines permitted actions, prohibited actions, permission conditions, delivery standards, and escalation requirements.
运行编码	Run Identifier	The identifier for a single Agent Instance run, linking its task, inputs and outputs, permissions, calls, versions, and inspection records. Records of rework or re-execution runs should be retained separately.

Chinese Term	Full English Name	Definition
评分规则	Scoring Rules	Rules for converting observations into scores, including scoring direction, scoring bands, endpoint inclusion, and treatment of missing and abnormal values.
评分口径	Scoring Specification	Unified specifications for the objects, statistical scope, calculation methods, units, and periods on which scoring is based, allowing the same facts to be calculated consistently.
七站作业流程	Seven-Stage Task Workflow	The operating workflow in which a task passes, sequentially or according to dependencies, through task intake, Rule Matching, team formation, execution, validation, Deliverable Consolidation, and archiving. It defines each stage's inputs, outputs, conditions for continuation, and routes for handling abnormalities.
说明性模拟编码 (SIM)	SIM / SIM-01	Used in this white paper to identify illustrative simulation tasks. SIM-01 is not an actual backend Run Identifier; V0.1 and V0.2 are also simulated versions.
六审	Six Design Reviews	The collective designation for the six design reviews corresponding to the Six Review Perspectives. Six Review Perspectives names the perspectives; Six Design Reviews names the work collectively. Eligibility and Composition Review separately undertakes eligibility judgments and composition selection.
六鉴	Six Review Perspectives	Six independent review perspectives: Business Relevance Review, Calculation and Scoring Specification Review, Data Source and Access Review, Discriminative Validity Review, Incremental Information Review, and Manipulation Resistance Review. They respectively review candidate factors' business meaning, calculability, data conditions, basis for discrimination, incremental information, and resistance to manipulation.
库存单位 (SKU)	Stock Keeping Unit	An inventory management unit used to distinguish specific product items. The product scope, period, and calculation specification for the SKU sales activity rate are governed by the formal rules.
专项小队	Task-Specific Team	A temporary collaboration unit formed around a specific task goal. It may dissolve after the task is completed; members' position identities, human accountability assignments, and execution records are retained.

Chinese Term	Full English Name	Definition
牛码 SHARP 双罗盘	NUMA SHARP Twin Compasses	A parallel observation framework consisting of the Core Compass and Industry Compass, presenting return support and risk load in contract structure and specific operations, respectively.
三类行业视角	Three Industry Perspectives	Three perspectives: Ordinary Operator Perspective, Industry Veteran Perspective, and Adversarial Manipulation Perspective. They respectively focus on everyday operating mechanisms, industry experience and implicit constraints, and pathways for packaging indicators and exploiting rules.
三个投资者视角	Three Investor Perspectives	Three judgment perspectives that start from business, price, and risk, respectively, to analyze value creation and its sustainability, investment price conditions, and potential adverse outcomes.
十二个足够	Twelve Sufficiency Conditions	The four groups of conditions required for HAT over MASS: sufficiently small, narrow, and granular tasks; sufficiently diverse, cross-checked, and cost-effective foundations; sufficiently numerous, fine-grained, and counterbalancing staffing; and sufficiently comprehensive, long, and deep inspection. The specific explanations are in Appendix II. They do not represent uniform numerical thresholds.

Table A3-1: Term Definitions and Chinese–English Correspondence

Suites 2105–2108
Two Exchange Square
8 Connaught Place, Central, Hong Kong

Micro Connect Group Website: www.microconnect.com
Email: info@microconnect.com