UNLMTD.GROWTH · WORKING PAPER NO. 2

The UNLMTD AGV Density Framework

A framework for measuring Agent Generated Value — the ratio of operational value produced by autonomous agents versus human labor — and the emerging economic dimension that will define which companies thrive in the AI-native era.

Alexander Borodich
General Partner, UNLMTD.Capital

Dubai · May 2026
Version 1.2

ABSTRACT

This paper introduces Agent Generated Value (AGV) as a measurable economic dimension and the UNLMTD AGV Density Framework as the operational methodology for assessing it in early-stage companies. AGV is defined as the ratio of operational value produced by autonomous agents versus human labor within an organization. The Framework specifies four sub-dimensions: the AGV ratio itself, the rate of growth of the company's reusable agent skill library, the autonomous decision rate, and the quality of human-agent collaboration at handoff points.

The central claim of the paper is that AGV Density is becoming a structural predictor of capital efficiency and revenue-per-employee for companies built in the AI-native paradigm — and that this dimension was not measurable, and largely did not exist as a coherent category, prior to approximately 2023. The paper proposes the Framework as the first published methodology for treating AGV Density as a first-class metric in venture analysis, formalizes the AI-Native Classification that distinguishes AI-Native, AI-Augmented, and AI-Layered companies, and develops scoring rubrics, observable measurement signals, case study applications, and limitations.

Because AGV Density is a new and rapidly evolving phenomenon, the framework is presented with explicit epistemic discipline. The retrospective evidence base is small — perhaps three years of observable AGV-native company operations — and the prospective predictive power of the framework has not yet been validated against multi-year outcomes. The paper distinguishes the conceptual contribution (introducing AGV Density as a measurable category) from the predictive claim (high AGV Density correlates with capital efficiency advantages), and treats the latter as a testable hypothesis rather than an established finding.

The Framework is the second published artifact of the UNLMTD.Growth Methodology, following the Borodich Trend Velocity Model. Together, the two papers operationalize the universal-law layer of the broader methodology for two of the most underrecognized dimensions in venture analysis: market timing and human-agent operational composition.

> AGV Density is the first new venture metric of the AI-native era. The paper time-stamps its authorship within the canonical record of the era, while explicitly acknowledging that the empirical work to validate the framework's predictive claims is years rather than months away.

1 · WHY A NEW METRIC

The Metric Venture Did Not Have

For the first time in venture history, the operational composition of a company — how much value is created by humans versus by autonomous agents — is itself a structural predictor of outcomes. Venture has no settled vocabulary for this. The paper proposes one.

Until approximately 2023, every venture-backed software company operated on a single underlying assumption: that operational value was produced by human labor augmented by tools. The tools could be sophisticated. The leverage could be substantial. But the production function had a constant ratio in which one human was approximately equivalent to one unit of decision-making capacity, with software multiplying that capacity but not replacing it. Within this assumption, capital efficiency was a function of how many humans a company needed per unit of revenue, and how much each human could produce.

This assumption is no longer universally true. A growing class of companies — Sierra, Decagon, Cognition, the productized AI assistant cohort emerging from 2023 forward — operates with a different production function. In these companies, a substantial fraction of operational decisions is made by autonomous agents rather than by humans. The agents are not tools in the conventional sense. They are coordinating actors with reasoning capacity, memory, and the ability to compose into multi-step workflows. The companies that build around this production function exhibit revenue-per-employee, capital efficiency, and scaling dynamics qualitatively different from companies that do not.

Venture analysis has no settled vocabulary for this phenomenon. The existing language — "AI company," "agentic startup," "AI-first business" — collapses several distinct underlying realities into a single ambiguous category. It also fails to provide a measurable metric that would allow investors to distinguish companies whose AGV ratio is structurally high from companies whose AGV ratio is incidentally low even while their marketing emphasizes AI capability.

This paper proposes a vocabulary and a measurement methodology. AGV — Agent Generated Value — is the share of operational value within a company that is produced by autonomous agents rather than by humans. AGV Density is the framework for measuring it. The framework is not the only possible methodology for this measurement, and it is not the final one. It is an opening proposal, advanced with the conviction that the underlying phenomenon is too economically important to remain unmeasured.

Why This Matters Economically

If AGV Density is a structural rather than cosmetic feature of a company, it has economic consequences that ripple through every other dimension of venture analysis.

Capital efficiency

A company with AGV ratio of 0.7 — seventy percent of operational value produced by agents — requires approximately a third of the human headcount to produce the same revenue as a comparable AGV-0.1 company at scale. This is not a marginal efficiency gain. It is a discontinuity in the cost structure of running a business. Companies operating at high AGV Density should achieve gross margins, revenue per employee, and burn multiples that are unreachable for traditional production functions, and these advantages should compound rather than degrade with scale.

Scaling dynamics

Traditional software companies face a hiring problem at scale: each new tier of customer count requires non-linear additions of support, success, sales, and operations headcount, with substantial coordination cost. AGV-native companies face a different problem: each new tier of customer count requires expanding the agent skill library and validating new autonomous decision categories, but with substantially lower coordination cost because agents do not require management overhead in the traditional sense. The result is that AGV-native companies scale through different bottlenecks — primarily evaluation infrastructure and skill library completeness — rather than through hiring.

Defensibility composition

Traditional defensibility comes from data network effects, brand, distribution, or switching costs. AGV-native companies acquire an additional defensibility layer from their compounding skill library: the accumulated agent capabilities, calibrated to the company's specific customer base and workflows, that cannot be replicated by competitors without traversing the same operational learning curve. This is a new kind of moat, and the paper argues it should be measured separately from other defensibility forms.

> AGV Density is not a marketing claim. It is a measurable property of the company's production function, with structural consequences for capital efficiency, scaling dynamics, and defensibility that ripple through every other dimension of venture analysis.

What This Paper Claims and Does Not Claim

The clarity of the central concept requires equal clarity on its limits.

The paper claims: that Agent Generated Value is a measurable economic dimension; that the four sub-dimensions specified in the framework capture the underlying phenomenon with sufficient operational precision to be applied in investment decisions and founder self-assessment; that high AGV Density correlates with capital efficiency advantages in companies built within the AI-native paradigm; and that the AI-Native Classification distinguishing AI-Native, AI-Augmented, and AI-Layered companies captures structurally distinct categories with measurably different operating dynamics.

The paper does not claim: that AGV Density predicts outcomes deterministically; that high AGV companies will necessarily outperform low AGV companies on any specific timeline; that the framework as specified is the final form of AGV measurement; that the empirical base supporting the framework's predictive claims is mature (it is not, the cohort of AGV-native companies operating at scale is small and recent); or that AGV Density is the only structural dimension that matters for AI-era venture analysis. Other dimensions matter equally; this paper develops one.

With those constraints stated, the rest of the paper develops the framework in operational detail.

2 · THE FRAMEWORK

The Four Sub-Dimensions

AGV Density is one metric, but the underlying phenomenon has four distinct components that must be measured separately and aggregated coherently.

A single ratio — say, sixty percent of operational value produced by agents — collapses information that should be kept analytically separate. Two companies with identical AGV ratios may have very different compounding dynamics, defensibility profiles, and scaling trajectories. The framework therefore decomposes AGV Density into four sub-dimensions, each measuring a distinct aspect of the underlying phenomenon, with the composite AGV Density score computed from their conjunction.

On What "Operational Value" Means

Before specifying the four sub-dimensions, the framework requires clarity on a term it uses repeatedly: operational value. The phrase is doing substantial conceptual work, and conflating its different meanings produces measurement errors that compound across all four sub-dimensions.

Operational value is not a single dimension. It includes at least five qualitatively distinct categories, each of which may be produced by humans or by agents in different proportions, and each of which carries different weight in the company's economics.

- **Execution value:** ticket handling, scheduling, document processing, content generation, classification. High-volume, individually-modest decisions that aggregate into substantial operational output. This is the category AGV measurement captures most cleanly.
- **Coordination value:** routing decisions, synchronization across workflows, orchestration of multi-step processes, handoff specifications. The category that determines whether execution value compounds or fragments.
- **Strategic value:** product direction, resource allocation, organizational architecture, acquisition decisions, capital deployment. Low-volume but individually consequential decisions. Often the territory where human judgment retains primacy regardless of AGV trajectory.
- **Relational value:** trust formation, negotiation, persuasion, conflict resolution, customer intimacy. The category where agent participation faces nonlinear acceptance thresholds — customers may tolerate agent execution but reject agent-led negotiation in the same workflow.
- **Symbolic value:** founder presence at critical moments, institutional legitimacy, accountability signaling. The category that operates partly outside efficiency calculations

because its purpose is to communicate that humans remain meaningfully accountable for outcomes.

The AGV Ratio specified in Sub-Dimension 1 below captures Execution value cleanly, Coordination value with reasonable fidelity, Strategic value partially (because strategic decisions are infrequent enough that volume-weighting underweights them), Relational value with substantial difficulty, and Symbolic value almost not at all. Analysts applying the framework should hold this asymmetric coverage in mind: a company with high AGV Ratio is a company with high Execution and Coordination autonomy, not necessarily a company that has displaced humans across all five value categories.

This asymmetry has a practical implication. Companies that attempt to maximize AGV across all five categories — including Strategic and Relational — typically encounter trust ceilings, customer resistance, or accountability failures that erode the operational advantages they sought. The framework's prescriptive force is appropriately bounded: high AGV is valuable where the value category permits autonomous execution, and counterproductive where it does not. Subsequent versions of the framework may develop category-specific scoring; the current paper treats the asymmetry as an acknowledged limitation rather than fully resolves it.

Sub-Dimension 1 · AGV Ratio

Measures: The share of operational value produced by autonomous agents versus human labor.

Range: 0.0 (entirely human-produced value) to 1.0 (entirely agent-produced value).

The AGV Ratio is the foundational measurement. It captures the fraction of operational value — output that the customer pays for or that supports the value creation pipeline — that is generated by autonomous agents acting without explicit human approval on each decision.

Critically, AGV Ratio is not a measurement of how much AI the company uses. A company can run extensive AI tooling throughout its operations while its AGV Ratio remains low — if humans remain in the decision loop for every customer-facing or operationally-significant action. AGV is measured at the unit of value-creating decisions, not at the unit of AI tool deployment. The question is not "does AI participate in the workflow" but "is the final decision made by an agent operating autonomously, or by a human reviewing and approving an agent recommendation."

Measurement is conducted through workflow audit: enumerate the company's value-creating decisions over a representative time window, classify each by whether the actor with final authority was a human or an autonomous agent, and weight by economic value of the decision. The result is a defensible AGV Ratio estimate with confidence intervals that depend on workflow visibility.

Sub-Dimension 2 · Skill Library Growth Rate

Measures: The velocity at which the organization codifies new agent capabilities into reusable, composable skills.

Range: 0.0 (no codification, ad hoc agent usage) to 1.0 (systematic skill library with measurable accumulation).

AGV Ratio measures the current state of agent autonomy. Skill Library Growth Rate measures the trajectory. A company with AGV Ratio of 0.4 today and a high skill library growth rate is on a different trajectory than a company with the same ratio and a low growth rate. The first is compounding; the second is plateaued.

A skill is a discrete, named, reusable agent capability that the organization has validated for deployment in production decisions. "Reschedule customer meetings when conflicts emerge" is a skill. "Triage incoming support tickets to the appropriate workflow" is a skill. "Detect unusual patterns in financial reconciliation" is a skill. Skills accumulate in a library that is itself a corporate asset — the more skills a company has, the more workflows it can automate, and the more compounding occurs across automation initiatives.

Skill Library Growth Rate is measured by the rate of validated skill additions per month, weighted by the breadth of workflow coverage each new skill enables. A company adding one or two skills per quarter, each addressing narrow workflows, scores low. A company adding five to ten skills per month, each composable with existing skills, scores high. The metric explicitly rewards composability because composable skills produce non-linear capability growth.

Sub-Dimension 3 · Autonomous Decision Rate

Measures: The share of agent decisions made without human approval per individual decision instance.

Range: 0.0 (every agent decision requires human approval) to 1.0 (agents act fully autonomously within their domain).

AGV Ratio captures the share of value produced by agents. Autonomous Decision Rate captures how much of that production happens without human checkpoint. The distinction matters because a company can produce most of its value through agent assistance — agents drafting documents, agents suggesting actions, agents preparing analyses — while requiring human approval on every individual instance. Such a company has high AGV Ratio but low Autonomous Decision Rate, and it does not scale through agent leverage in the way a higher Autonomous Decision Rate company does.

Autonomous decisions are those where the agent executes without a human reviewing and approving the specific output. Human-approved decisions are those where each agent recommendation passes through a human checkpoint. The distinction is granular: a company may have ten agent workflows, of which three are autonomous and seven require approval, producing an Autonomous Decision Rate of 0.30 weighted by decision volume.

Measurement requires workflow inspection rather than self-report. Companies frequently overestimate their Autonomous Decision Rate because they confuse "AI-assisted" with "AI-decided." The framework's scoring methodology explicitly probes the location of the human checkpoint for each major workflow.

Sub-Dimension 4 · Handoff Quality

Measures: The quality of collaboration at points where work transitions between agents and humans, or between agents and other agents.

Range: 0.0 (handoffs lose context, create errors, require rework) to 1.0 (handoffs preserve context, maintain consistency, minimize error introduction).

The fourth sub-dimension captures something that AGV Ratio and Autonomous Decision Rate together miss. A company can have high AGV Ratio and high Autonomous Decision Rate while losing substantial operational efficiency at the boundaries between agents and humans, or between different agent systems. Handoff failure is the most common operational pathology in AI-native companies and is also the least visible from external assessment.

Handoff Quality is measured along several axes: context preservation across handoffs, error introduction rates at boundaries, rework volume due to handoff failure, time spent in handoff coordination versus value-producing work, and the existence of structured handoff protocols rather than ad hoc transitions. Companies with high Handoff Quality typically exhibit explicit handoff specifications, well-defined inputs and outputs for each agent role, and minimal context drift across multi-agent workflows.

The Composite AGV Density Score

With all four sub-dimension scores collected, the composite is computed as the geometric mean — the same aggregation function used in the Trend Velocity Model and for the same conceptual reason. The sub-dimensions are conjunctive rather than additive: weakness in any single dimension substantially reduces the value of strength in the others.

> AGV Density Score = (AGV_Ratio × Skill_Library_Growth × Autonomous_Decision_Rate × Handoff_Quality)^(1/4) Where each sub-dimension is scored 0.0 to 1.0, and the composite is the geometric mean to preserve the conjunction property: a company with high AGV Ratio but poor Handoff Quality cannot achieve a high AGV Density score, because the handoff failures will erode the value created by autonomous decisions in practice.

A company with sub-dimensions 0.9 / 0.9 / 0.9 / 0.2 has an AGV Density of approximately 0.61 — the low Handoff Quality dominates the composite. The intuition is correct: a company that produces excellent agent work but fails at the boundaries is, in practice, less efficient than a company with more modest agent capability but excellent handoffs. The geometric mean captures this empirical reality.

> A high AGV Ratio without high Handoff Quality is a company that has built a beautiful machine and is operating it through a faulty switchboard. The customer experiences the switchboard.

3 · CLASSIFICATION

The AI-Native Classification

Three structurally distinct categories of company hide behind a single fashionable label. Distinguishing them is the first methodological prerequisite for everything else.

The current investment landscape includes thousands of companies that describe themselves as AI companies, AI-first companies, AI-powered platforms, or any of several adjacent terms. Without a classification scheme, the category is analytically useless: it includes companies that operate on fundamentally different production functions and exhibit fundamentally different economics, all collapsed into a single ambiguous label.

The UNLMTD AI-Native Classification specifies three categories distinguished by their position on the AGV Density spectrum.

Category 1 · AI-Native

Definition: AGV Density above 0.5, with an agentic core architecture in which autonomous agents make a substantial fraction of value-producing decisions.

AI-Native companies are organized around an agentic core. The fundamental unit of operational work is the agent decision, with humans serving in roles of strategic direction, edge-case escalation, and skill library development rather than as the primary actors in each value-creating workflow. AI-Native companies exhibit measurably different revenue-per-employee, gross margin, and burn multiple profiles compared to their AI-Augmented and AI-Layered peers, with the differences becoming more pronounced as the company scales.

Examples (where publicly observable): Sierra, Decagon, and the cohort of agentic customer experience companies emerging from 2023 forward; Cognition's Devin and the agentic engineering assistant cohort; the cohort of AI-native vertical workflow companies that have explicitly built around agent operations rather than human operations augmented by AI. The category is small in absolute numbers but growing rapidly.

AI-Native companies typically score high across all four sub-dimensions of AGV Density. The scoring profile is recognizable: not just high AGV Ratio, but explicit skill library architecture, deliberate autonomous decision boundaries, and engineered handoff protocols. Companies claiming AI-Native status without exhibiting this profile across all four sub-dimensions are typically AI-Augmented or AI-Layered companies using the AI-Native label for positioning purposes.

Category 2 · AI-Augmented

Definition: AGV Density between approximately 0.2 and 0.5, where humans remain the primary actors but are systematically augmented by AI tooling that increases their per-capita output.

AI-Augmented companies represent the majority of so-called AI companies in 2026. They use AI substantially — often pervasively — but the AI serves as a productivity multiplier for human workers rather than as the primary decision-making agent. A salesperson uses AI to draft outreach; the salesperson decides which prospects to contact and approves the outreach before sending. A customer support agent uses AI to draft responses; the agent reviews and approves before sending. A developer uses AI to suggest code; the developer accepts, modifies, or rejects each suggestion.

AI-Augmented companies achieve real productivity gains compared to non-AI peers, often substantial ones. Per-capita output may increase by factors of two to five. But the production function remains human-centered, and the company's capital efficiency profile, while improved, does not exhibit the discontinuous jump that AI-Native companies exhibit. AI-Augmented companies scale by hiring more humans who use AI productively, with all the linear scaling dynamics that hiring implies.

The category is methodologically important because many companies that describe themselves as AI-Native are in fact AI-Augmented, sometimes intentionally — the AI-Native label commands a valuation premium in current markets — and sometimes through self-misunderstanding. The framework's scoring methodology is designed in part to distinguish these cases through workflow audit rather than self-report.

Category 3 · AI-Layered

Definition: AGV Density below approximately 0.2, where AI features have been added to an existing product or workflow architecture that was not built around AI.

AI-Layered companies have legacy products that have been enhanced with AI features. The underlying product architecture, customer workflows, and economic model predate the AI era and have been retrofitted. The retrofit may produce genuine customer value, but the company's underlying economics remain those of the pre-AI era. The AI features are marketing surface and incremental value-add rather than structural reorganization of how the company operates.

Examples include the substantial cohort of enterprise SaaS companies that have added AI assistant features to existing products, traditional consumer products with AI-powered recommendations bolted on, and legacy services businesses that have layered AI tooling on top of human-delivered services. These companies remain valuable, in many cases produce real returns, and may grow profitably. But they are structurally different from AI-Native companies on every dimension the AGV framework measures, and the venture economics applicable to them should not be confused with the venture economics applicable to AI-Native companies.

On Classification Boundaries

The three categories above are presented as discrete classifications, but real companies frequently exhibit hybrid profiles. A company may operate at AI-Native density within its support function, at AI-Augmented density within its sales function, and at AI-Layered density within its strategic operations — all simultaneously, within the same legal entity. The linear classification scale captures the average AGV across functional areas; it does not capture this internal heterogeneity.

This is a methodological simplification with practical consequences. A company averaging AGV Density of 0.45 might be a fairly uniform AI-Augmented operation, or it might be a company with AI-Native support (AGV 0.7) coexisting with AI-Layered sales (AGV 0.2). These two profiles look identical on the linear scale but represent very different organizational realities, with very different migration trajectories and very different economic implications.

Where classification matters operationally — investment decisions, founder self-assessment, organizational planning — the framework recommends functional decomposition rather than single-score classification. Score AGV Density within each major functional area separately, then assess the company's overall profile from the distribution rather than from the average. Future versions of the framework may develop multidimensional classification topology to formalize this; the current paper treats classification as primary while acknowledging that the discrete categories are conceptual anchors rather than strict boundaries.

Migration Dynamics

Classification is not permanent. A company may begin as AI-Layered and migrate to AI-Augmented and ultimately to AI-Native through deliberate operational restructuring. The migration is difficult — it typically requires substantial workflow re-engineering — but it is observable and measurable through changes in the four AGV sub-dimensions over time.

Migration velocity, like the underlying AGV scores, is itself a signal. A company actively migrating from AI-Augmented toward AI-Native — measurable as rising AGV Ratio and Autonomous Decision Rate over multiple quarters — exhibits different economics than a company stable at AI-Augmented level. The framework's prospective application includes migration tracking for portfolio companies and for prospective investments where classification is in flux.

> The label "AI company" is increasingly meaningless. The classification — AI-Native versus AI-Augmented versus AI-Layered — is the analytically useful distinction. The framework provides the methodology for making it rigorously.

4 · MEASUREMENT

The Measurement Protocol

Self-reported AGV scores are unreliable. The framework requires workflow-level audit. The protocol specifies how that audit is conducted.

Measurement of AGV Density cannot rely on company self-report because companies systematically overstate the autonomy of their AI systems. The overstatement is sometimes intentional (the AI-Native label commands valuation premiums) and sometimes structural (founders often genuinely misunderstand how often humans remain in their workflows once they have built sophisticated AI tooling). The framework therefore specifies a workflow-level audit protocol that produces defensible AGV scores even in the face of overstatement.

AGV Ratio Measurement

AGV Ratio measurement requires enumeration of the company's value-creating decisions over a representative time window, typically one to three months. For each decision category, the analyst classifies the actor with final authority and weights by economic value.

Step 1 — Workflow enumeration

List the company's primary value-creating workflows. For each workflow, identify the decisions that produce or affect customer-visible value. A customer support workflow may produce a single decision per ticket (the resolution). A sales workflow may produce multiple decisions per opportunity (qualification, outreach content, pricing, contract terms).

Step 2 — Decision classification

For each decision category, identify the actor with final authority. The question is: who has the last word before the decision affects the customer or the company's external state? A human reviewer who reads an agent's output and approves it before send is the actor with final authority, even if the content was agent-generated. The agent is the actor with final authority only if it executes without per-instance human approval.

Step 3 — Economic weighting

Weight each decision category by its economic significance. A customer support response is weighted differently from a contract negotiation. The weighting may be based on customer-visible value, internal economic impact, or proxy measures such as revenue attribution or cost avoidance.

Step 4 — Composite calculation

The weighted average of autonomous decisions divided by total decisions produces the AGV Ratio. A confidence interval is computed based on workflow visibility — companies that resist workflow audit, or that have substantial workflows the analyst cannot inspect, should be scored with substantially wider confidence intervals.

Skill Library Growth Rate Measurement

Skill Library Growth Rate requires inspection of the company's actual skill library — the codified, named, validated agent capabilities — over time. Three observable indicators.

- Documented skill count and addition velocity, measured monthly
- Skill composition graphs — visualizations of which skills compose into which workflows, indicating composability versus isolation
- Skill quality validation — what processes the company uses to validate that new skills perform reliably before deployment

Companies without documented skill libraries score low on this sub-dimension regardless of how much AI capability they appear to have, because undocumented capability does not compound. The codification itself is the locus of the compounding mechanism.

Autonomous Decision Rate Measurement

Autonomous Decision Rate measurement requires probing the specific locations of human checkpoints in each workflow. The protocol asks, for each workflow identified in Step 1 above, whether the human reviews each individual instance or whether the human establishes boundaries within which the agent acts autonomously.

Per-instance human review reduces Autonomous Decision Rate to near zero regardless of how much agent capability the workflow contains. Boundary-setting human oversight — where the human establishes thresholds, escalation criteria, or policy guardrails but does not review individual decisions — preserves Autonomous Decision Rate. The distinction is operationally critical and is often opaque from external observation, requiring direct workflow inspection.

Handoff Quality Measurement

Handoff Quality is the most observable sub-dimension externally, because handoff failures produce visible operational pathologies — customer complaints, support escalations, internal coordination overhead. The framework measures Handoff Quality through three signals.

- Existence of explicit handoff specifications — documented input and output contracts for each agent role and human role
- Error rates at handoff boundaries — what fraction of multi-step workflows fail due to context loss or specification mismatch at transitions
- Rework volume — how much work has to be redone because of handoff failures, as a fraction of total work output

Companies with high Handoff Quality typically have explicit Handoff Quality measurement as a first-class operational metric. Companies with low Handoff Quality typically do not measure it explicitly, which is itself one of the signals.

Founder Interview Protocol

Workflow audit is the primary measurement methodology, but it is supplemented by a structured founder interview that probes the founder's articulation of the company's AGV architecture. The interview is approximately twenty-five minutes and includes four question categories.

Architecture probe (8 minutes)

"Walk me through your primary customer-facing workflow from start to end, identifying each decision point and who has final authority on each." The probe distinguishes founders who understand their own AGV architecture from those who do not.

Skill library probe (5 minutes)

"Show me your skill library — the named, codified, reusable agent capabilities you have built. What was added in the last three months? What is composable with what?" Founders without explicit skill libraries fail this probe immediately.

Autonomous boundary probe (7 minutes)

"For each workflow, describe the policy or boundary within which the agent operates autonomously. What triggers escalation to a human?" Probes the location and explicitness of autonomy boundaries.

Handoff probe (5 minutes)

"Where do handoffs occur — between agents, between agents and humans, between humans and agents — and what specifications exist at each boundary? What is your error rate at the most failure-prone handoff?" Founders with high Handoff Quality typically know their handoff error rates by heart. Founders with low Handoff Quality typically struggle to answer.

5 · CASE STUDIES

Three Worked Cases

The AGV framework applied to three publicly observable companies operating across the three categories of the AI-Native Classification. The cases illustrate the framework in operation, with explicit acknowledgment that scoring at this stage relies on partially observable evidence.

Each case study below applies the AGV Density Framework to a publicly observable company based on available evidence from product documentation, customer reports, press coverage, and founder interviews. Because none of these companies have undergone the full workflow audit specified in Chapter 4, the scores are estimates with wider confidence intervals than would result from privileged access. The cases are intended to illustrate the framework's application rather than to produce definitive scores.

Case 1 · Sierra (AI-Native)

Public profile: Conversational AI platform for customer experience, founded by Bret Taylor and Clay Bavor; unicorn valuation in mid-2024.

AGV Ratio — estimated 0.75 ± 0.10

Sierra's customer-facing product is, by architecture, an agentic platform. The agent handles substantial fractions of customer interactions end-to-end, with humans serving in escalation and configuration roles rather than as the primary actors in each conversation. Public materials and customer reports suggest that for well-configured deployments, agent handling fractions exceed seventy percent of conversations.

Skill Library Growth Rate — estimated 0.80 ± 0.10

The company's product architecture explicitly centers on a deployment of named, reusable skills configured per customer. Each new customer deployment expands the canonical skill library by adding capabilities that generalize across customers. Founder communications emphasize this compounding effect.

Autonomous Decision Rate — estimated 0.70 ± 0.15

Within configured customer deployments, agent decisions execute autonomously across a substantial fraction of conversation paths, with human escalation triggered by policy boundaries rather than per-instance review. The estimate reflects observed customer deployment patterns; actual rates vary significantly by customer configuration.

Handoff Quality — estimated 0.65 ± 0.15

The agent-to-human handoff is a critical operational dimension for the product category, and Sierra has invested substantially in handoff specifications. Public customer reports include both successful handoff examples and instances of context loss at agent-human transitions. The estimate reflects active engineering investment with imperfect outcomes.

Composite AGV Density: approximately 0.72 — AI-Native classification with strong sub-dimension profile across the four metrics. Conjunction strength is high; no single sub-dimension dominates as a weakness.

Case 2 · A Representative AI-Augmented Enterprise SaaS Company

Profile: A widely deployed sales engagement platform that has added AI-powered features for content generation, prospect prioritization, and meeting scheduling. The category contains many such companies; this case is composite rather than identifying any single firm.

AGV Ratio — estimated 0.30 ± 0.10

The platform's AI features participate substantially in user workflows — drafting outreach content, prioritizing prospect lists, summarizing meeting notes. But the user (a human sales representative) retains final authority on every outbound message, every prioritization decision that affects pipeline allocation, and every meeting commitment. AI generates content; humans approve and send.

Skill Library Growth Rate — estimated 0.40 ± 0.15

The platform has accumulated some named AI capabilities deployed across customer instances, but the architecture is feature-based rather than skill-library-based. New AI features are released as product updates rather than as additions to a composable library. The compounding dynamic of true skill library architecture is not present.

Autonomous Decision Rate — estimated 0.15 ± 0.10

Per-instance human review is the dominant pattern. Even where AI features run continuously in the background — for example, lead scoring — the scores produce recommendations that humans use to make manual decisions. Truly autonomous decisions are rare and typically limited to non-customer-facing analytical functions.

Handoff Quality — estimated 0.55 ± 0.20

Handoffs in this category typically occur between AI features and humans within a single user's workflow, with reasonable but inconsistent context preservation. Multi-agent handoffs are rare in this architecture, simplifying the handoff problem but limiting the upside.

Composite AGV Density: approximately 0.30 — solidly within the AI-Augmented classification. The platform produces genuine productivity gains for its users but operates on a fundamentally different economic model than an AI-Native company.

Case 3 · A Representative AI-Layered Legacy Enterprise Application

Profile: A widely deployed enterprise application — for example, a procurement management system — that has added AI-powered features (intelligent search, document summarization, anomaly detection) to a product architecture predating 2020. Again the case is composite.

AGV Ratio — estimated 0.10 ± 0.05

AI features serve as additions to a workflow architecture built around traditional enterprise software conventions. The fundamental operational unit is the human user clicking through forms and reviewing reports. AI features assist with discrete tasks within these workflows but do not restructure the workflows themselves.

Skill Library Growth Rate — estimated 0.15 ± 0.10

AI capabilities are added through traditional product feature roadmaps rather than through skill library accumulation. New AI features are individually scoped, individually validated, and added through software release cycles. No composable library exists.

Autonomous Decision Rate — estimated 0.05 ± 0.05

Per-instance human action remains dominant across the application. AI features inform human decisions but do not execute autonomous actions on the company's behalf. Anomaly detection alerts a human; the human investigates and decides.

Handoff Quality — estimated 0.40 ± 0.15

The AI features integrate with existing workflows with variable smoothness. Handoff failures between AI suggestions and human action are common but generally tolerable because the AI is positioned as assistance rather than as a workflow participant.

Composite AGV Density: approximately 0.13 — clearly within the AI-Layered classification. The company remains a viable enterprise software company; it is simply not the same kind of company as an AI-Native firm and should not be analyzed using AI-Native economics.

> Across the three cases, the AGV Density scores differ by a factor of roughly five between the AI-Native and AI-Layered companies — even though all three deploy substantial AI capability and all three describe themselves as AI companies. The classification distinction is not cosmetic. It corresponds to fundamentally different economic models.

What We Do Not Yet Know

AGV Density is a new dimension. The retrospective evidence is thin. The prospective predictive power has not been validated. This chapter says so explicitly.

The Trend Velocity Model published as Working Paper No. 1 confronts a specific epistemological challenge: retrospective coherence is not predictive clarity. The AGV Density Framework confronts a related but distinct challenge: the dataset of AGV-native companies operating at scale is too small and too recent to support strong empirical claims about predictive validity. Where the Trend Velocity Model has fifteen years of portfolio outcomes against which to retrospectively fit, the AGV Density Framework has perhaps three years of observable AGV-native company operations.

This is not a defect of the framework. It is a property of the phenomenon. AGV-native companies as a category did not exist before approximately 2022. The empirical base required to validate the framework's predictive claims is being created in real time. The honest response is to publish the framework with explicit acknowledgment of this constraint, rather than to overclaim predictive power that the evidence does not yet support.

Three Specific Uncertainties

Whether high AGV correlates with multi-year capital efficiency

The framework's central predictive claim is that high AGV Density correlates with capital efficiency advantages that compound over multi-year periods. This claim is plausible by economic reasoning — agent labor scales differently from human labor, and the differential should compound — but it has not been validated against five-year or seven-year outcomes. The companies that would generate such evidence are too recent. The earliest AGV-native companies are approximately three years into their operations as of this writing, and their long-term capital efficiency trajectories are not yet observable.

Whether the AGV ceiling is bounded

The framework treats AGV Ratio as a dimension that scales from 0.0 to 1.0. But the practical ceiling for sustainable operations is unknown. Companies operating at AGV Ratios of 0.9 or above may face structural limits — escalation volumes that overwhelm human escalation capacity, edge cases that no skill library can fully cover, customer trust thresholds that cap autonomous decision acceptance. The empirical maximum sustainable AGV Ratio is unknown, and may turn out to be substantially below 1.0 across all current architectures.

Whether the framework generalizes across industries

The case studies in Chapter 5 are drawn from sales engagement, customer experience, and enterprise software. Other industries — regulated healthcare, financial services with substantial fiduciary requirements, physical-world operations — may exhibit AGV dynamics that the current framework does not capture well. The framework's scoring rubrics are calibrated against the industries where AGV-native companies have first emerged, and application to other industries should be treated as exploratory.

Calibration Agenda

The framework's response to these uncertainties parallels the response in the Trend Velocity Model: explicit commitment to ongoing calibration as the evidence base accumulates.

- **Real-time scoring of UNLMTD.Capital AI-Native deals:** every AI-Native or AI-Augmented deal entering the active pipeline receives an AGV Density score at time of investment, with the score and reasoning timestamped. Multi-year outcomes will be tracked against initial scores and the calibration gap reported.
- **Cross-sector validation:** the framework will be applied to AGV-relevant companies across healthcare, financial services, and physical operations as suitable cases emerge. Sector-specific calibration adjustments will be developed and reported as patterns become clear.
- **AGV ceiling research:** as AGV-native companies mature, the empirical ceiling on sustainable AGV Ratio will become observable. Subsequent versions of the framework will incorporate observed ceilings rather than treating the dimension as unbounded.
- **Public AGV registry:** as the methodology matures, UNLMTD will publish AGV Density scores on publicly observable companies, with both successful and failed predictions retained as part of the calibration record.

None of these calibration mechanisms will produce mature results in less than several years. The framework is published now, in advance of mature validation, because the underlying phenomenon is unfolding in real time and a vocabulary is needed before it can be measured. Publishing the framework with explicit uncertainty markings is better than publishing it with false confidence, and substantially better than waiting for empirical maturity before publishing at all.

> Some metrics must be measured before they can be validated. AGV Density is one of them. The framework is published now in the spirit of opening a research direction rather than closing one.

7 · A P P L I C A T I O N

Application Protocols

The framework serves three distinct audiences: investors evaluating AI-related opportunities, founders evaluating their own AGV architecture, and operators tracking AGV evolution over time.

For Investors

The framework integrates into early-stage investment evaluation at three points.

Pre-meeting classification

Before first founder meeting, conduct preliminary classification of the company against the three categories of the AI-Native Classification, using publicly available information. The classification produces a baseline expectation: AI-Native companies should be evaluated using AI-native economics, AI-Augmented companies using augmented-productivity economics, AI-Layered companies using legacy software economics. Confusion of these economic models is among the most expensive errors in current AI-era venture analysis.

In-meeting AGV interview

During the first founder meeting, conduct the four-probe AGV interview specified in Chapter 4. The interview is approximately twenty-five minutes and should be integrated into a longer conversation. The Architecture probe and the Skill Library probe are the most informative; founders who cannot walk through their own AGV architecture clearly are typically operating at lower AGV Density than they describe in marketing materials.

Workflow audit for advanced diligence

For deals proceeding to advanced diligence, the workflow audit specified in Chapter 4 should be conducted with founder cooperation. This is the only methodology that produces high-confidence AGV scores; self-report and pattern matching from publicly observable signals produce wider confidence intervals than committed capital decisions should depend on.

For Founders

The framework serves founders as both a diagnostic and a roadmap.

Self-diagnostic

Founders should be able to articulate their company's AGV Density across all four sub-dimensions, with specific evidence supporting each score. A founder who cannot identify their own AGV Ratio with reasonable precision, name the skills in their skill library, specify the autonomous decision boundaries, and describe the handoff specifications is operating without the operational visibility that AGV-native companies require.

Migration roadmap

For founders operating in the AI-Augmented category who aspire to migrate toward AI-Native, the framework provides a structured roadmap. Migration typically proceeds through deliberate workflow re-engineering: identifying decision categories where autonomous agent execution is feasible, building the skill library incrementally with explicit composability requirements, establishing autonomous decision boundaries with measurable quality validation, and engineering handoff specifications before the operational pathologies emerge. Migration is difficult — most AI-Augmented companies do not successfully migrate — but the framework makes the migration path explicit rather than mysterious.

For Operators

The framework serves operators by providing a continuously-updating metric for tracking AGV evolution over time. Quarterly AGV Density measurement on the four sub-dimensions allows operators to observe whether their AGV trajectory is rising, plateaued, or declining, and to identify which sub-dimension is the binding constraint at any given moment.

Most companies have a binding constraint that, when addressed, unlocks substantial AGV gains. For some companies, the binding constraint is Skill Library architecture — capabilities exist but are not codified into composable form. For others, it is Autonomous Decision Rate — capabilities are composable but human checkpoints remain on every instance. For others, it is Handoff Quality — autonomous decisions are made but handoffs lose context. The framework's quarterly tracking surfaces the binding constraint and focuses operational attention where it has the highest leverage.

> AGV Density is not a property that companies have or do not have. It is a trajectory that companies are on. The framework makes the trajectory measurable and the binding constraints explicit.

8 · LIMITATIONS

Limitations and Future Work

The framework's limits are as important as its claims. This chapter names them explicitly.

What the Framework Does Not Capture

AGV Density measures the operational composition of value production. It does not measure product-market fit, customer love, market timing, competitive dynamics, founder capability beyond AGV literacy, or capital structure. A company with high AGV Density and weak fundamentals on these other dimensions will fail despite its operational sophistication. AGV Density is one of sixteen Universal Laws specified in the UNLMTD.Growth Methodology and is necessary but not sufficient.

In practical application, AGV Density should be one input into a broader scoring framework, not a sole basis for decision. A composite UNLMTD.Growth score that includes Trend Velocity, AGV Density, founder cognitive complexity, plan velocity, and the other named contributions provides substantially more decision-useful signal than any single dimension alone.

Known Failure Modes

Self-report overstatement

Companies systematically overstate their AGV Density. The overstatement is the dominant source of measurement error and is the reason the framework specifies workflow audit rather than founder interview as the primary measurement methodology. Analysts relying on founder self-report should apply substantial downward adjustment to claimed AGV scores, particularly when the company has fundraising incentives.

Workflow opacity

Some companies operate workflows that are not externally inspectable, particularly companies with proprietary agent infrastructure or with customer-specific deployments that vary substantially across instances. Confidence intervals for AGV scoring should widen substantially in cases of workflow opacity, and analysts should treat such scores as preliminary rather than definitive.

Migration timing

Classification at a single point in time can mislead when applied to companies undergoing rapid AGV evolution. A company classified as AI-Augmented at time of investment may migrate to AI-Native over eighteen months, fundamentally changing its economics. Or it may fail to migrate, remaining permanently at AI-Augmented level despite founder intentions. Migration velocity is partially observable but not reliably predictable, and AGV scoring should be supplemented with explicit migration trajectory assessment.

AGV Theater and sub-dimension gaming

As the framework's vocabulary spreads, sophisticated founders may construct narratives that score well on each sub-dimension without underlying operational reality matching the narrative. This is the failure mode that becomes most acute precisely when the framework becomes most influential — what we call AGV Theater: visible AI deployment that masks the absence of meaningful agent autonomy, performative skill library construction, claimed autonomous decision rates that conceal hidden human checkpoints, and handoff specifications that exist primarily as documentation rather than as operational practice.

The interview protocol includes probes designed to surface theater — asking founders to walk through specific workflows in detail, to demonstrate the skill library in production rather than as concept, to name handoff error rates from measurement rather than estimation — but these probes face a moving target. Analysts evaluating AGV scores in years three and four of the framework's existence should apply more skepticism than analysts in year one, on the principle that gaming sophistication tends to grow with framework familiarity. Workflow audit, not founder narrative, remains the primary defense against AGV Theater, and the framework's preference for audit over interview reflects this concern.

The structural risk is not individual founders gaming individual scores; it is the broader ecosystem drift toward valuation premiums for AGV claims that exceed AGV reality. The framework's response is to commit to publishing AGV scores with explicit audit methodology, so that gamed scores can be distinguished from audited scores by external observers. This is part of the calibration agenda specified in Chapter 6.

Causal ambiguity

The framework asserts correlation between high AGV Density and capital efficiency advantages. It does not claim that high AGV Density causes capital efficiency advantages directly. Alternative causal interpretations are consistent with the same correlational pattern: companies with strong founders may be both more capital-efficient and more able to build AGV architecture; companies in particular sectors may exhibit both AGV opportunities and capital efficiency advantages from independent causes; selection bias in venture portfolios may inflate the apparent strength of any single signal. The framework is consistent with multiple causal interpretations and treats the correlation as the empirically meaningful claim.

Trust-Constrained AGV Ceilings

The framework's scoring methodology treats higher AGV Density as structurally favorable, on the empirical observation that AGV-native companies exhibit capital efficiency advantages. But this prescriptive implication requires important qualification: economically optimal AGV is not necessarily maximum AGV. Different value domains carry different trust thresholds beyond which autonomous agent decision-making produces value destruction rather than value creation, regardless of operational efficiency.

A fully autonomous hedge fund would, in narrow operational terms, achieve maximum AGV Density. It would also concentrate catastrophic tail risk in a way that no rational capital allocator would accept. A fully autonomous medical diagnosis system would similarly maximize AGV but face social and regulatory ceilings that no operational efficiency argument can override. A fully autonomous legal negotiation system would erode institutional accountability structures in ways that may impose costs exceeding the efficiency gains. In each case, the economically optimal AGV ceiling is meaningfully below 1.0, and the framework's prescriptive force should be modulated accordingly.

The general principle: trust thresholds are domain-specific, nonlinear, and frequently asymmetric. Customers may accept full autonomy in Execution value (an AI handling a customer service ticket) while rejecting it in Relational value (an AI conducting a contract negotiation on the customer's behalf) within the same workflow. Regulators may tolerate autonomous decisions in low-stakes domains while requiring human accountability in domains with catastrophic downside. The framework's scoring should be applied with explicit awareness of these constraints, and high AGV scores should not be interpreted as universally desirable across all sectors, value categories, or stakeholder relationships.

Subsequent versions of the framework may develop trust-constrained AGV ceiling estimates by sector and value category. The current paper treats this as an acknowledged limitation, with the practical guidance that AGV Density should be evaluated against domain-appropriate ceilings rather than against the theoretical maximum.

Migration as Political Transformation

The framework treats AGV migration as primarily an operational and architectural challenge: workflow re-engineering, skill library construction, autonomous decision boundary specification, handoff protocol design. This treatment captures the technical dimensions of migration accurately, but underweights a dimension that frequently determines migration success or failure: organizational politics.

AGV migration is not merely technical transformation. It is also political transformation. Existing organizational hierarchies are organized around human decision-making, and shifting decision authority to autonomous agents redistributes power, status, and economic claims within the organization. Managers whose role depends on coordinating human workers face displacement when those workers are replaced by agents. Departments whose budgets reflect headcount face contraction when headcount is replaced by skill library expansion. Status hierarchies built around the symbolic significance of human judgment face erosion when significant judgments are increasingly executed by agents.

These dynamics frequently produce predictable patterns: managerial displacement resistance, incentive misalignment between operational leaders and AGV migration goals, symbolic labor preservation (humans retained in checkpoint roles whose primary function is to signal continued human authority rather than to add operational value), and AI adoption theater (visible AI deployment that masks the absence of meaningful agent autonomy). Many AGV migration failures are organizationally political rather than technically inadequate.

The framework's measurement methodology partially captures this through the workflow audit approach, which surfaces gaps between claimed and actual agent autonomy. But political resistance to migration is not yet modeled as a first-class dimension. Companies attempting AGV migration should treat political alignment as parallel infrastructure to technical capability, and analysts evaluating AGV trajectories should weight observed political resistance as a leading indicator of stalled migration.

Substrate Dependencies

AGV-native companies depend on infrastructure conditions that are themselves becoming strategic constraints in 2026: compute availability and pricing, energy economics for inference at scale, model access including frontier model dependency and provider concentration, latency requirements, and jurisdictional access to AI infrastructure that varies substantially across regulatory regimes. These substrate dependencies are not currently first-class variables in the AGV Density scoring, but they meaningfully constrain the sustainable AGV ceiling for any individual company.

A company operating at AGV Density 0.7 with reliable inference economics, model provider redundancy, and favorable jurisdictional access faces a different sustainability profile than a company at the same AGV Density operating with single-provider dependency, marginal inference economics, and constrained jurisdictional access. The first is robust to substrate volatility; the second is not. Analysts should assess substrate dependencies as a complement to AGV Density scoring rather than as a separate consideration.

This dimension intersects substantially with the Capital Climate Framework anticipated as a future working paper in the broader UNLMTD.Growth Methodology. Substrate economics in the AI era — compute, energy, model access, jurisdictional positioning — operates at a scale that connects company-level AGV decisions to sovereign and geopolitical conditions. The current paper acknowledges the connection without fully developing it.

Regime Boundaries

Industry transferability

The framework is empirically grounded in software-mediated workflow companies. Application to physical-world operations (manufacturing, logistics, healthcare delivery, agriculture) requires substantial methodological adaptation. Physical operations introduce constraints — regulatory, safety, latency, embodiment — that the current framework does not explicitly address. Subsequent working papers may develop industry-specific AGV variants.

Regulatory environment

AGV Density depends partly on the regulatory environment within which the company operates. Jurisdictions with strict AI accountability requirements may structurally cap Autonomous Decision Rate for companies operating within them. Jurisdictions with looser frameworks may permit higher AGV Density at the cost of accountability tradeoffs the framework does not score. As global regulatory frameworks for AI continue to evolve, the framework's scoring may need jurisdiction-specific adjustments.

Temporal validity

The framework is calibrated against the AGV-native cohort that emerged from approximately 2022 forward. The specific scoring thresholds, sub-dimension weights, and category boundaries may need recalibration as the AGV-native paradigm itself evolves. The framework is positioned as a research program rather than a fixed methodology, and explicit revision cycles are anticipated as the empirical base matures.

Future Work

Seven extensions of the framework are anticipated in subsequent working papers, listed in approximate priority order.

1. Skill Library Economics: a full economic subsystem developing the implications of reusable agent capability as accumulated capital. Specifically: skill compounding dynamics, marginal skill reuse curves, cross-workflow transferability, capability graph density, orchestration overhead, recursive capability amplification, and recombination velocity. This may be among the strongest latent contributions of the AGV framework and warrants standalone treatment in a future working paper.
2. AGV trajectory modeling: explicit modeling of migration velocity between classification categories, allowing prospective assessment of which AI-Augmented companies are likely to successfully migrate to AI-Native.
3. AGV ceiling research: empirical investigation of sustainable AGV Ratio ceilings across industries and customer types — including the trust-constrained ceilings described above — refining the upper bounds of the framework's scoring scale.
4. Cross-sector calibration: industry-specific scoring adjustments for healthcare, financial services, physical operations, and other sectors where AGV dynamics differ structurally from the software-mediated workflow archetype.
5. AGV × Trend Velocity interaction: theoretical and empirical investigation of how AGV Density interacts with Trend Velocity Model scores — whether high AGV companies are systematically more or less sensitive to trend timing, whether AGV amplifies adaptive

capacity, and whether the two dimensions reinforce each other into a single adaptive acceleration coefficient.

6. Organizational entropy theory: the framework approaches but does not yet formalize the dynamics of coordination entropy, context loss, specification decay, and recursive drift that distinguish AGV-native organizational failure modes from traditional operational inefficiency. Handoff Quality is currently treated as operational metric; it may eventually function as the entry point into a broader organizational thermodynamics that warrants standalone development.
7. Recursive environmental effects: the framework primarily models internal organizational economics. But high-AGV firms reshape adjacent systems — labor markets, regulatory expectations, customer patience thresholds, industry competitive equilibria, capital allocation flows. AGV may eventually function as market-shaping infrastructure rather than only firm-internal architecture. This second-order effect deserves separate investigation.

Each of these extensions requires empirical work that will mature over coming years. The framework is published now in advance of that maturation, on the principle that providing a vocabulary for a phenomenon is the precondition for studying it rigorously. The current paper has reached the point where additional sub-dimensions or conceptual layers would increase complexity faster than operational infrastructure can support; the disciplined response is to consolidate the framework as specified, build measurement infrastructure around it, and develop further dimensions as separate working papers rather than as extensions to this one.

CONCLUSION

In Summary

Agent Generated Value is a new economic dimension. The category of company that produces substantial fractions of operational value through autonomous agents — rather than through humans augmented by tools — did not exist as a coherent class before approximately 2022. By 2026 the class is small but rapidly growing, and the venture economics applicable to AGV-native companies differ structurally from the venture economics applicable to AI-Augmented or AI-Layered companies that have used similar marketing language to describe themselves.

The UNLMTD AGV Density Framework introduces AGV as a measurable dimension with four sub-dimensions: AGV Ratio, Skill Library Growth Rate, Autonomous Decision Rate, and Handoff Quality. The composite is computed through geometric mean to preserve the conjunction property, since weakness in any single dimension substantially erodes the value of strength in the others. The AI-Native Classification, derived from AGV Density measurements, distinguishes three structurally distinct categories of AI-related company — AI-Native, AI-Augmented, and AI-Layered — that current venture vocabulary tends to collapse together.

The framework's predictive claims have not yet been validated against multi-year outcomes; the AGV-native company cohort is too recent for such validation to be possible. The paper publishes the framework now, with explicit epistemic discipline, on the principle that providing a vocabulary is the precondition for measurement, and measurement is the precondition for science.

AGV Density is the second published artifact of the UNLMTD.Growth Methodology, following the Borodich Trend Velocity Model. Subsequent working papers will develop additional named contributions of the broader framework, with the eventual UNLMTD.Growth Codex integrating all contributions into a coherent reference document.

A final note on scope. The framework presented here begins to approach territory historically occupied by organizational economics and theory of the firm — Coase, Drucker, Christensen, and the broader academic tradition of production function analysis. We acknowledge this proximity without attempting to fully claim it. AGV Density as a venture metric is one paper. AGV Density as a foundation for theory of the post-human firm is a multi-year academic project requiring engagement with the existing literature in ways this paper does not attempt. The framework is positioned as a venture methodology with implications for firm theory, not as firm theory in its own right. Whether the implications develop into independent academic contribution depends on empirical work that neither this paper nor any single working paper can complete.

> Some metrics must be invented before they can be validated. AGV Density is one of them. The honest

response is to publish the framework, name its limits explicitly, and commit to revising it as the empirical base matures.

Alexander Borodich

General Partner, UNLMTD.Capital

Dubai · May 2026

APPENDIX A

AGV Density Scoring Rubric

This appendix consolidates the full scoring rubric in single-page reference form for operational use during investment decisions or founder self-assessment.

A.1 · AGV Ratio

Score	Criteria
0.8–1.0	Workflow audit confirms 80%+ of value-creating decisions executed by autonomous agents; humans serve in escalation and configuration roles only.
0.5–0.7	Workflow audit confirms 50-79% of decisions executed by autonomous agents; majority autonomous but humans remain primary actors in significant minority of workflows.
0.3–0.4	30-49% autonomous decisions; AI participates substantially but humans hold final authority on majority of value-creating decisions.
0.1–0.2	AI assistance pervasive but humans approve essentially every customer-facing or operationally-significant action.
0.0	AI features present but do not participate in value-creating decisions; AI as tooling, not as actor.

A.2 · Skill Library Growth Rate

Score	Criteria
0.8–1.0	Explicit skill library with documented composition; 5+ new skills added monthly; composability designed into architecture.
0.5–0.7	Skill library exists; new additions occur regularly; composition is partial; some workflows recombine skills.
0.3–0.4	Skills exist but library is implicit; additions occur but are not systematically codified; minimal composability.
0.1–0.2	AI capabilities are released as product features rather than as library additions; no codification of reusable skills.
0.0	No skill library exists in any meaningful sense; AI deployments are ad hoc.

A.3 · Autonomous Decision Rate

Score	Criteria
0.8–1.0	Most agent decisions execute without per-instance human approval; human oversight occurs via boundary-setting and escalation triggers.
0.5–0.7	Mixed pattern; some workflows operate fully autonomously with boundary setting, others require per-instance approval.
0.3–0.4	Minority of decisions are autonomous; most agent outputs reviewed by humans before action.
0.1–0.2	Per-instance human review is the dominant pattern; autonomous decisions limited to internal or low-stakes contexts.
0.0	All decisions reviewed by humans; AI provides only recommendations or analysis.

A.4 · Handoff Quality

Score	Criteria
0.8–1.0	Explicit handoff specifications for each transition; measured error rates; rework volume tracked as first-class metric; context preservation engineered.
0.5–0.7	Handoff specifications partially documented; error rates known but not systematically tracked; reasonable but inconsistent context preservation.
0.3–0.4	Handoffs occur but specifications are ad hoc; error rates not systematically measured; visible operational pathologies at boundaries.
0.1–0.2	Handoff failures common; significant rework volume; context loss at agent-human transitions is a recognized but unaddressed problem.
0.0	No handoff specifications; transitions are ad hoc with high failure rates; rework dominates output.

A.5 · Composite Interpretation

AGV Density Score is computed as the geometric mean of the four sub-dimensions. Thresholds and corresponding classification:

AGV Density	Classification	Economic Model
≥ 0.5	**AI-Native**	Agentic core architecture; capital efficiency and revenue-per-employee dynamics qualitatively different from human-centered companies; scaling through skill library rather than hiring.
0.2–0.49	**AI-Augmented**	Human-centered production function with AI as productivity multiplier; 2-5x per-capita output improvement but linear hiring-based scaling dynamics.
< 0.2	**AI-Layered**	Legacy product architecture with AI features added; underlying economics of the pre-AI era; AI as marketing surface rather than structural reorganization.

ABOUT THIS PAPER

WORKING PAPER	UNLMTD.Growth Working Paper No. 2
TITLE	The UNLMTD AGV Density Framework
VERSION	1.2 · May 2026
AUTHOR	Alexander Borodich, General Partner, UNLMTD.Capital
PARENT METHODOLOGY	UNLMTD.Growth Methodology — Universal Law U15 and Archetype 9 Classification
PRECEDED BY	Working Paper No. 1 — The Borodich Trend Velocity Model
NEXT IN SERIES	Working Paper No. 3 — The Borodich Founder Spiral Diagnostic
LICENSE	Free for public reading, citation, and academic use with attribution. Commercial reproduction with attribution to UNLMTD.Capital.
CONTACT	UNLMTDgrowth.org · UNLMTD.Capital

Citation:

Borodich, A. (2026). The UNLMTD AGV Density Framework. UNLMTD.Growth Working Paper No. 2. UNLMTD.Capital. Version 1.2.

U N L M T D . C A P I T A L

D u b a i · A b u D h a b i · G l o b a l