APR 22, 2025

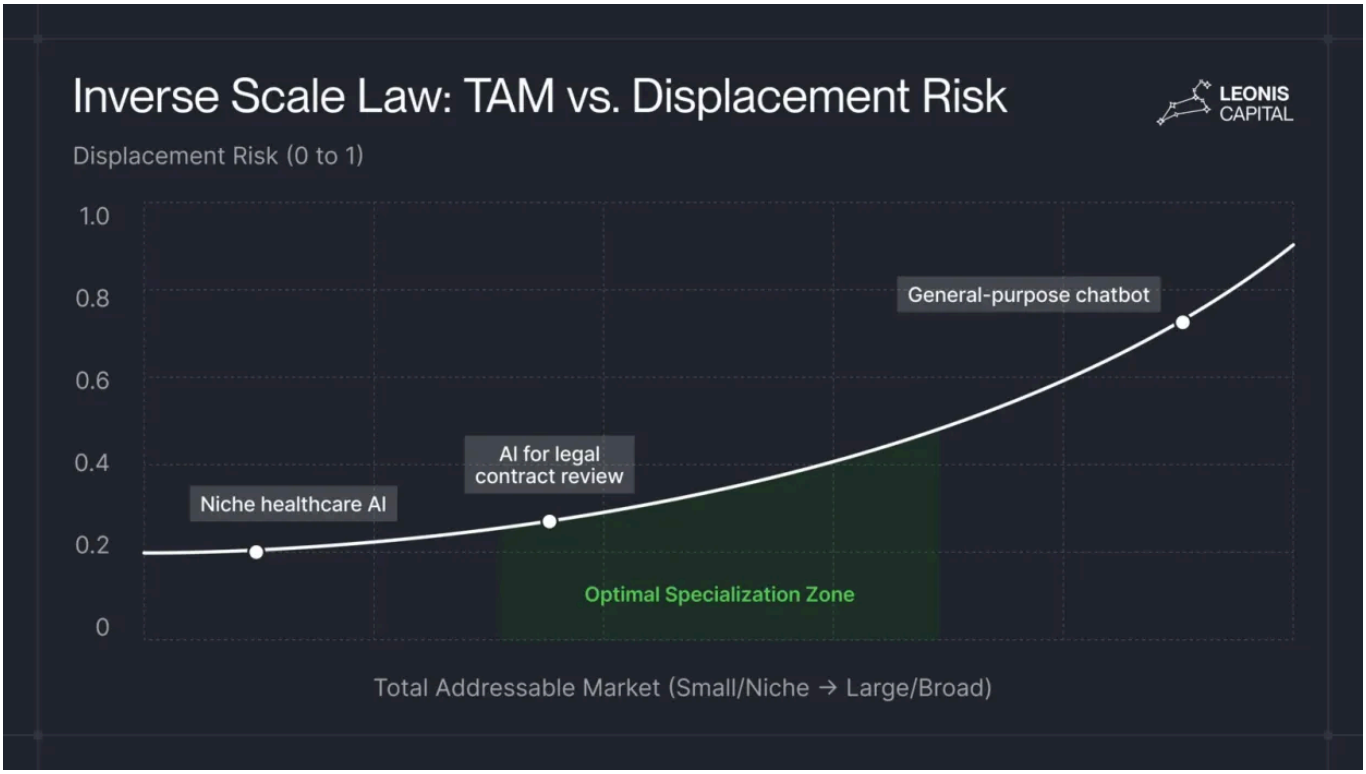
Zero or Hero: A Technical Framework for Valuing AI Companies (Part II: AI Applications)

Why Vertical AI Apps Will Survive (But Horizontal Ones Might Collapse to Zero)



By Jenny Xiao, Liang Wu, and Jay Zhao

Leonis Capital · leoniscap.com/research/zero-or-hero-a-technical-framework-for-valuing-ai-companies-(part-ii-ai-applications)/



In March 2025, OpenAI's upgraded 4o model took the internet by storm with its enhanced image generation capabilities. Users worldwide flocked to create remixes in various artistic styles, including those inspired by Studio Ghibli, the renowned Japanese animation studio. In total, users generated over 700 million images in just the first week. This viral phenomenon instantly placed leading image generation companies like Midjourney in a fight for relevance. And as history has shown, competitors are likely not far behind OpenAI in releasing similar capabilities. Because foundation model companies are locked in a hyper-competitive race, they likely won't stop at image generation either. OpenAI has already released audio generation features through its API, and it's only a matter of time before it makes these capabilities available to users through

ChatGPT. All of this begs the question: **Is there a need for specialized AI applications when foundation models offer similar features and much more?**

In Part I, we introduced the "zero-value threshold" principle, where if crossed, the value of a foundation model or application collapses to zero. This principle is especially critical for AI applications, which face the unique risk of being cannibalized by the very platforms they're built upon. We've repeatedly witnessed foundation models wipe out entire categories of AI applications simply by releasing features that users can access natively in a foundational model's interface.

Application-level AI startups face dual threats. They face risks from foundation models expanding vertically, integrating new features like image generation or domain-specific reasoning that were once the selling point of standalone applications. They also face horizontal threats from other application startups that expand into their industry. As the cost of code goes to zero, it is much easier for companies to reach feature parity and much harder to build a moat around single features.

In Part II of this valuation series, we analyze what types of features are most vulnerable to absorption and why certain positioning strategies prove more resilient than others. We highlight how, for an AI application to build a sustainable moat, there is an inverse relationship between the total addressable market (TAM) and the sustainability of the value they generate, and explain why targeting seemingly niche vertical markets (at least to start) often fares better against displacement risk. We also explore how moats evolve across different stages of a startup's lifecycle and how this impacts displacement risk in the broader context of our valuation framework.

The Obsolescence Cycle: When Your Product Becomes a Feature

A pattern of obsolescence emerges where AI startups lose ground to foundation model improvements. This "obsolescence cycle" looks something like this: an AI product initially captures attention with a novel use case like advanced summarization or a specialized image style. Then, a major LLM release arrives with upgrades that match or surpass the specialized offering. Almost overnight, that once-exciting product becomes just another "feature" of the new base model. Users begin to question the value of a standalone solution, churn accelerates, and the application's valuation ultimately plummets. This isn't a new phenomenon; we've seen it happen to companies like Jasper, and this phenomenon is accelerating at breakneck speed.

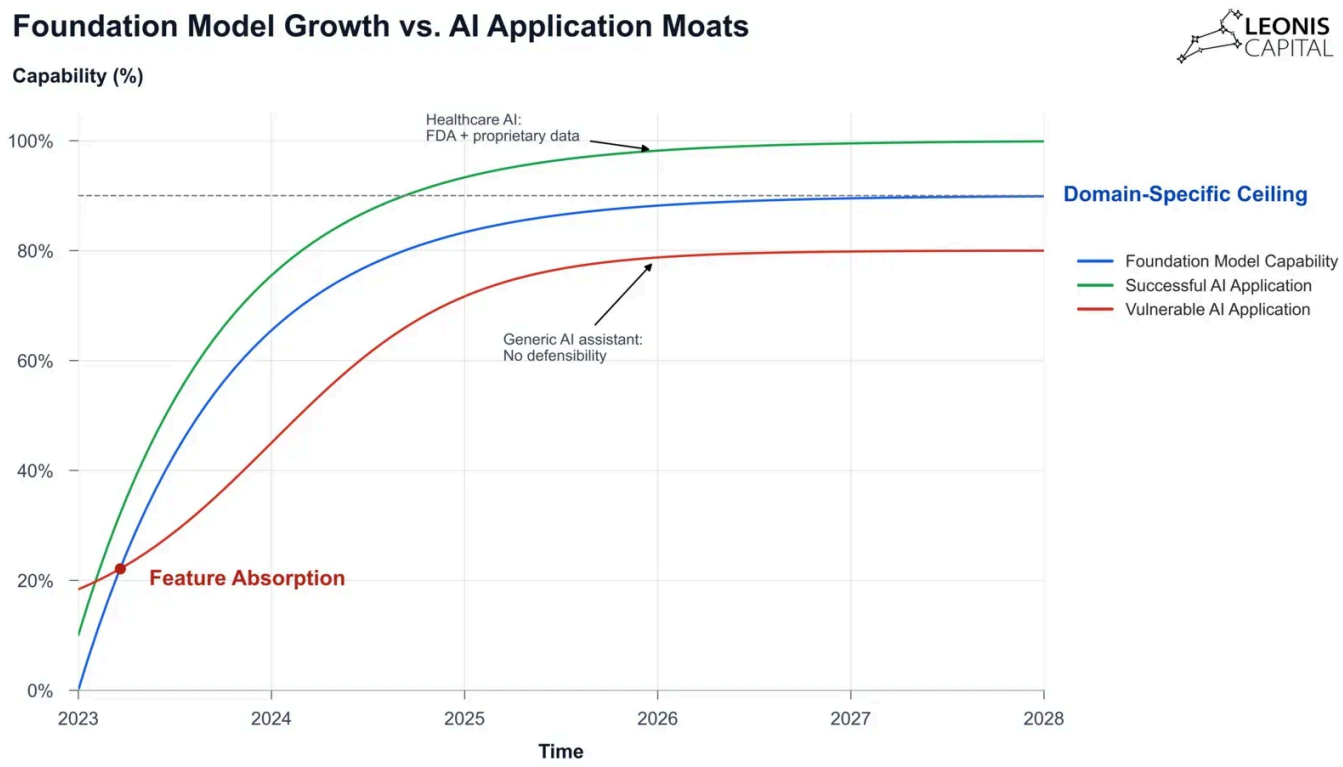


Figure 1: AI applications face the threat of absorption by foundation models, but foundation models also face a domain-specific performance ceiling.

Today, this pattern is playing out in multiple industries. In March 2025, Anthropic released Claude Code, its coding agent that developers can delegate tasks to in natural language within the terminal. Claude Code has quickly become one of the go-to tools for non-technical "vibe coders." Although Claude Code today doesn't challenge Cursor's grip in the developer market, its agentic capabilities put traditional AI coding assistants at risk.

Similarly, in AI-powered search, multiple foundation model companies like OpenAI, Grok, and Google have rolled out search functionality. All of this increases pressure on AI-powered search applications like Perplexity and renders AI search a commodity feature. Each new foundation model feature set pushes another class of single-purpose apps toward the zero-value threshold, underscoring how little time application-layer startups have to convert a clever feature into a durable moat.

Absorption Risk: Which Features Get Consumed First?

Not all AI application features face the same absorption risk by foundation models. The speed at which foundation models integrate specialized features varies based primarily on two factors: **how verticalized a feature is and how technically complex it is to build it**.

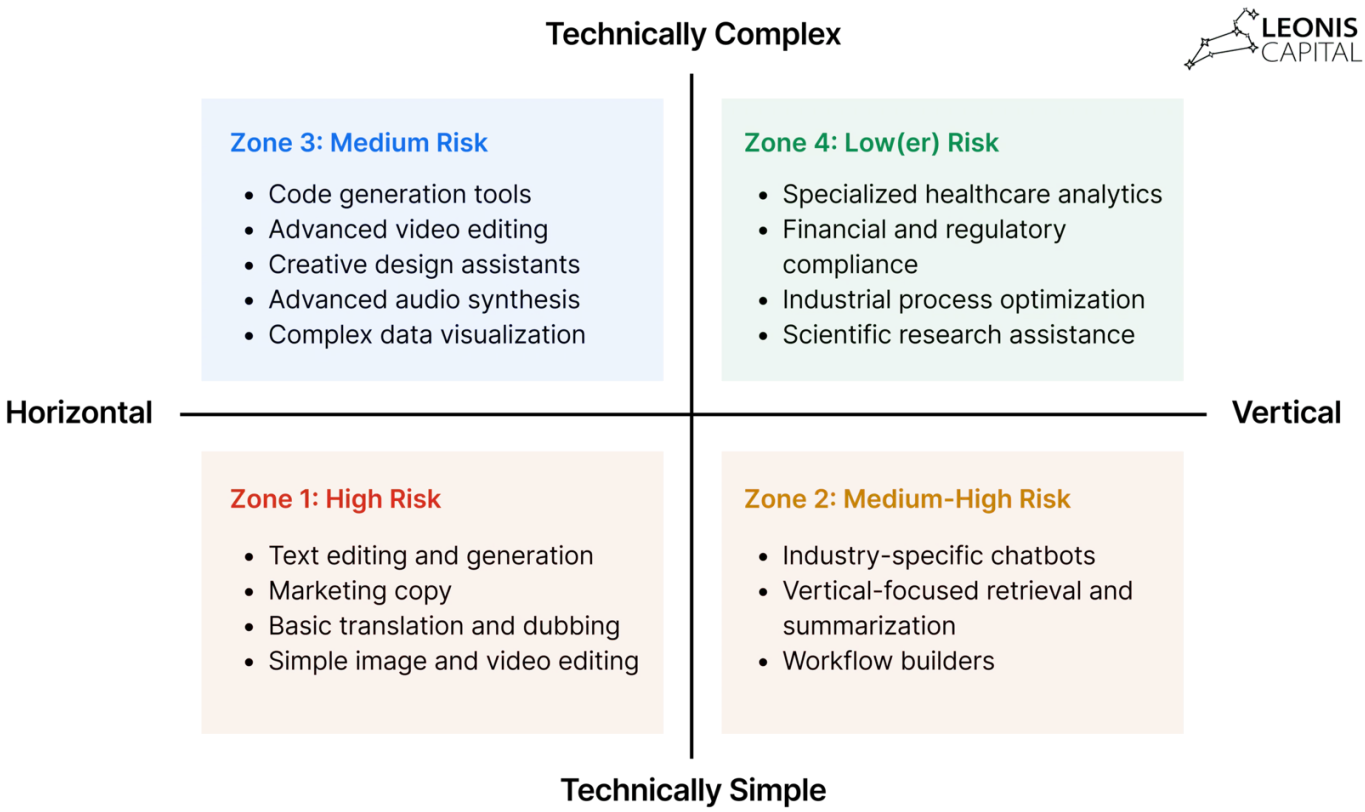


Figure 2: A 2x2 map of the absorption risk of different AI applications.

When we map these dimensions, four distinct categories of absorption risk emerge:

- **Zone 1: High-Risk [Horizontal, Technically Simple]** This quadrant faces the highest and fastest absorption risk. Foundation models prioritize these features because they're in high demand and relatively straightforward to implement. The absorption cycle is rapid: foundation models first match the feature's core functionality, then match its quality, and finally exceed it while offering additional capabilities at a fraction of the cost. *Examples:* Basic text generation, simple image editing, content summarization, simple language translation, and generic chatbots. For example, Jasper's core marketing copy generation fell squarely in this zone, explaining why it quickly found itself outcompeted by foundation models.
- **Zone 2: Medium-High Risk [Vertical, Technically Simple]** These applications serve specialized audiences with straightforward implementations. Their survival depends on how quickly the vertical grows in importance, as foundation models tend to prioritize features with broader appeal first. However, their technical simplicity means they can be quickly absorbed once targeted. *Examples:* Domain-specific chatbots, industry-focused content retrieval tools, specialized data extraction tools, and vertical-specific simple workflow builders. These applications often fly under the radar until their market grows large enough to attract attention. Unfortunately, most of the vertical AI applications today fall into this category. They might generate ARR quickly but face a serious threat of commoditization.

- **Zone 3: Medium-Risk [Horizontal, Technically Complex]** These features are in high demand but require specialized expertise or significant computing resources to execute well. While foundation models actively target these capabilities due to market demand, the technical barriers provide AI applications a longer runway to establish differentiation. *Examples:* Two clear cases are Cursor, which tackles advanced code generation inside the IDE, and Perplexity, which combines retrieval with live LLM summarization for AI-powered search. Both generated traction initially because foundation models could not immediately match their depth of features, but their lead is quickly shrinking. When comparing Zone 2 to Zone 3, technical complexity is more defensible than a vertical application that is technically simple.
- **Zone 4: Lower-Risk [Vertical, Technically Complex]** Applications combining specialized domain knowledge with technical sophistication have the strongest defenses against foundation model absorption risk. These typically involve unique data assets, specialized workflows, or human-in-the-loop processes that can't be easily replicated. *Examples:* Two standout examples are Harvey [1], which embeds legal reasoning and firm-specific knowledge into case preparation workflows, and Hippocratic AI, whose clinically validated, voice-based agents must satisfy rigorous healthcare-safety and compliance standards. Both products operate in domains where accuracy, liability, and regulation create barriers that a general-purpose LLM cannot quickly clear. When comparing Zone 3 and Zone 4, Zone 3's end outcome is likely being acquired by a foundation model company (e.g., OpenAI's buying Windsurf), whereas for Zone 4, they lose if there is another more specialized application that can do the same as the existing app but adds something incremental.

This absorption pattern is not static, and features can constantly shift positions on this 2×2 as foundation models evolve and markets change. What was technically complex yesterday can become simple today, and niche markets can suddenly become popular targets.

Today, Anthropic's Model Context Protocol (MCP) represents one of the most significant developments accelerating the obsolescence cycle for many AI applications. MCP changes how foundation models interact with external data, tools, and specialized capabilities. Before MCP, AI applications could build moats around their ability to bridge between LLMs and external capabilities, whether database access, specialized API calls, or tool integration. MCP effectively standardizes this connection layer, providing foundation models with a universal interface to access various tools and data sources.

With MCP making tool integration a commodity, AI startups must shift their focus from merely *connecting* models to *building* domain-specific tools and workflows that keep their edge even when accessed through a shared protocol. Real defensibility now lies in mastering the “last mile,” the messy edge-cases, compliance rules, proprietary data, and deep workflow integration that a generic model, no matter how well wired-up, still can’t handle. And **that’s precisely why, in a world where horizontal features can be copied overnight, the strongest moat comes from going deep in a narrow vertical,** a dynamic we’ll explore next.

The Inverse Scale Law: Why Vertical Often Beats Horizontal

Traditional SaaS wisdom celebrates scale; the larger your total addressable market (TAM), the greater your growth potential and valuation. The underlying logic is that a bigger TAM typically translates into a wider pool of potential customers, accelerates revenue growth, and attracts more capital. But in the AI application landscape, this formula inverts. Broad horizontal solutions targeting general business functions across industries (like marketing, sales, or customer service) are usually the first to be subsumed by next-generation foundation models. If your application attempts to serve everyone, it becomes increasingly vulnerable to improved LLM releases that can replicate its core functionality overnight. As a result, narrowing your focus often provides a more defensible foundation for building a sustainable moat.

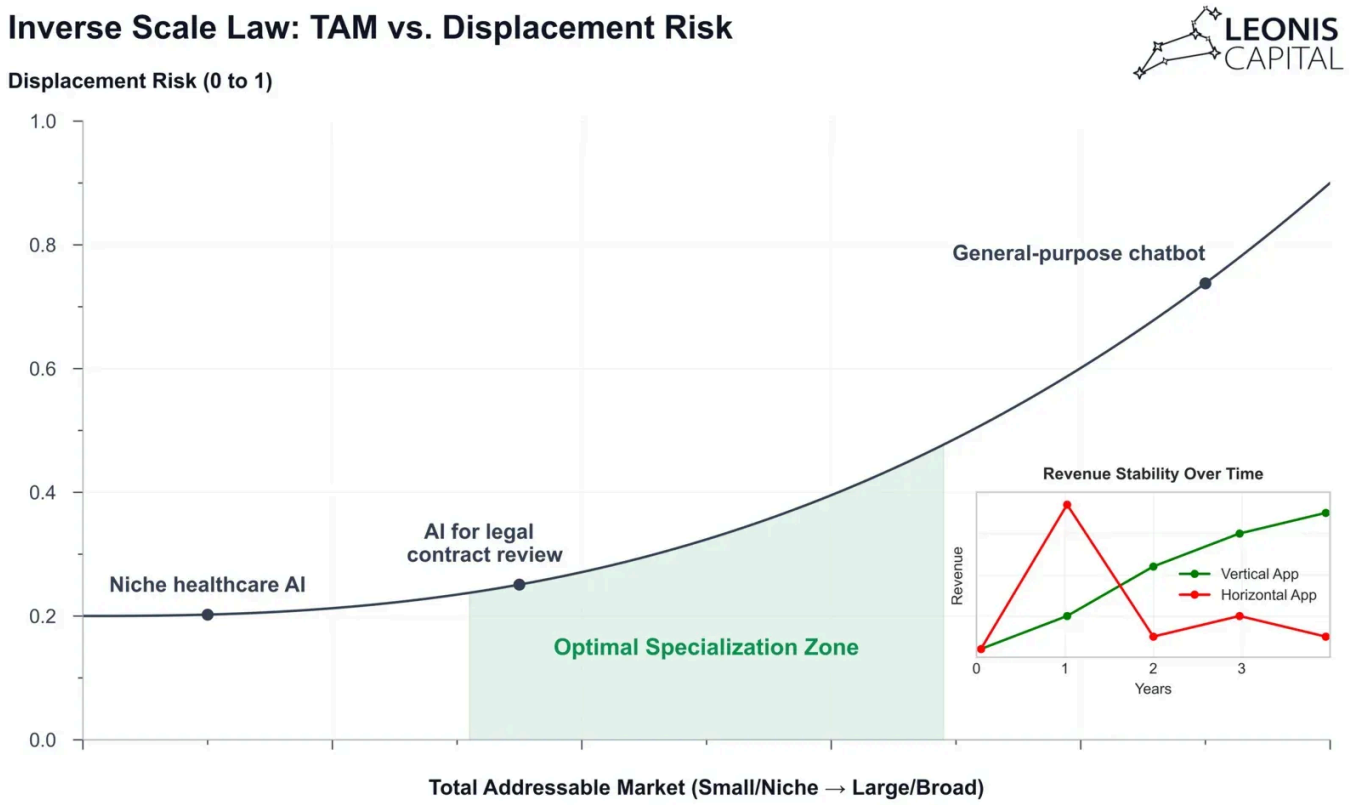


Figure 3: For AI companies, there is probably an optimal specialization zone where the product isn't too niche, cannot achieve venture-scale, and isn't too broad, and risks being absorbed by foundation models.

Horizontal AI applications often generate "vibe revenue" or hype-driven revenue when they first launch. They are shiny new toys that attract users excited by the technical innovation. But as foundation models improve, offering comparable capabilities, users begin to question why they should pay for a specialized service when the base models can perform similar functions for free or at a lower cost. A notable example is 11x.ai, which quickly gained attention for its AI sales

representatives like "Alice" and "Jordan," raising a $50 million Series B at a $350 million valuation. However, the initial excitement subsided as reports emerged that the company had falsely claimed customers. Clients like ZoomInfo discontinued their use after finding that the product underperformed compared to human sales representatives. These developments illustrate how initial hype can give way to scrutiny, especially when foundational models offer similar or superior capabilities.

Given the displacement risks explained in this essay so far, it becomes clear why vertical applications often offer stronger defenses against foundation model displacement. Verticalized AI companies enjoy moats from deep workflow integration, regulation and compliance, and proprietary data advantages. While these advantages may seem obvious, they are notoriously difficult to build (and sustain), as **most AI products achieve 80% functionality quickly, but the final 20% takes 5-10x longer and is what builds actual trust.** Even though foundation models excel at generic tasks, they often fail at edge cases. That's why general-purpose AI tools see early hype and fast churn, while vertical AI applications, though narrower in reach, often retain users and justify premium pricing through domain depth and stickiness.

The challenge for AI application founders is to find **the optimal specialization zone** and build a product that is specific enough to avoid direct competition from foundation models, yet broad enough to address a market that justifies growth. Most successful companies begin by laser-focusing on a single vertical or use case, establishing defensibility through the mechanisms described above, and then gradually broadening their capabilities after securing a solid initial foothold. In our portfolio, we've seen Describe AI expand from AI note-taking for psychiatrists and working with the largest mental health group in Florida to in-depth EHR workflow automation for multiple healthcare speciality domains. In the AI landscape, starting narrow enables companies to build sustainable moats before considering expansion beyond their core focus area.

Verticalized Multi-Agent Networks as a Moat?

If tomorrow's foundation models can copy any broad, one-off feature, defensibility must come from depth, not breadth. This is why we think there is a moat forming around **verticalized multi-agent networks**: clusters of narrowly trained agents that each handle a specific step in a regulated or domain-heavy workflow and continuously learn from the data those steps generate.

This network is bolstered by MCP, which would enable agent systems to proactively act. The result is a living, domain-specific operating system that can act in the real world and handle tasks like draft a compliant insurance quote, recalibrate a CNC machine, or triage a patient record, all the while logging every step for audit. That combination of **unique data, deep workflow integration, and evolving team-of-agents know-how** is something a generic model, no matter how advanced, cannot clone by simply reading your public docs.

Of course, nothing about multi-agent architecture is special on its own; the moat appears only if the network keeps harvesting exclusive data and embeds itself in moments where accuracy, liability, or regulation truly matter. But for founders who pick the right vertical and invest in that last-mile complexity, a specialized agent network could prove as sticky and as valuable as any software moat we have seen in the SaaS era.

The Modified ARR Framework Applied to AI Applications

The Application Layer's D Factor

In Part I, we introduced the zero-value threshold, where an application's value plummets to nothing once a foundation model catches up. Traditional valuation frameworks like ARR multiples fall short for AI applications. They don't account for the elephant in the room: **the looming threat of foundation models absorbing your product's core value**. To address this, we're adapting the standard ARR framework with a displacement risk factor, or D factor, to reflect the unique risks AI startups face.

The formula is straightforward but unforgiving:

> *Valuation = ARR x Valuation Multiple x (1-D)*

The D factor, ranging from 0 to 1, quantifies the likelihood that a foundation model will render your application obsolete. A D of 0 means you're well-defended; a D of 1 means your value is already eroding. Most AI applications are closer to 1 than their founders might hope. Below, we categorize applications based on their defensibility, from those with robust moats to those on the brink of disruption.

Category	D Factor	Moat Quality Metrics	Typical Characteristics	Examples
Reinforcing Moat System	0.1-0.2	• Core performance advantage >30% • Data + Domain + Workflow integration • Creates compounding value as system grows	• Multiple moats that strengthen each other • Each user interaction increases competitive gap • Clear unit economics improvement with scale	• Healthcare systems where clinical data improves models while workflow integration increases provider trust • Financial compliance platforms where regulatory expertise combines with model accuracy
Strong but Independent Moats	0.3-0.5	• Core performance advantage 15-30% • 2+ independent strong moats • Stable value proposition with clear user ROI	• Multiple protection layers that don't amplify each other • Good value but limited compounding effects • Defensible but static market position	• Enterprise tools with both strong integrations and unique data access • Analytics platforms with domain expertise and superior model performance
Single Major Moat	0.6-0.8	• Core performance advantage 5-15% • One significant moat with minimal secondary protection • Value proposition vulnerable to foundation model improvements	• One differentiating factor protecting from displacement • Easily replicable if primary moat is compromised • Declining advantage as foundation models improve	• Document processors with single industry integration • Content tools with basic performance advantages • Analyzers with limited proprietary data
No Sustainable Moats	0.9-1.0	• Core performance advantage <5% • No significant moats beyond basic functionality • Foundation models already matching capabilities	• Easily replicated by foundation model providers • No sustainable advantage left • Users actively migrating to foundation model solutions	• Basic summarization or generation tools • Simple workflow automations • Generic AI assistants

> *Table 1: How to evaluate AI applications' moats.*

The D factor measures an application's resilience against foundation model encroachment. Here are the key metrics to evaluate:

- **Foundation Model Capability Gap:** How much better is your application than a foundation model like o-3? If your performance edge is below 5%, the risk of disruption is high. A healthcare AI diagnosing rare conditions might have a 40% advantage, while a generic chatbot might barely scrape 2%.
- **Integration Depth:** How embedded is your solution in customer workflows? Applications that are deeply integrated, like a supply chain AI optimizing logistics for a manufacturer or an IDE

that developers spend all of their time in (e.g., Cursor), are harder to displace. If your tool is easily swapped out, it's at risk.
- **Data Network Effects:** Do you have proprietary datasets that improve with usage? A legal tech AI with access to decades of case law and that self-improves based on customer usage has a clear edge; a content generator relying on public data does not.
- **Regulatory/Compliance Barriers:** Does your application operate in a regulated industry? An AI with FDA clearance for medical imaging has a significant advantage over foundation models that can't easily enter these industries.

Consider an AI application with $10M in ARR and a standard SaaS multiple of 10x. In a traditional valuation, that's $100M. But if your D factor is 0.9—common for generic AI assistants—your adjusted value drops to $100M × (1-0.9) = $10M. A steep discount, but a fair one. Conversely, a healthcare AI with a D of 0.2 retains $80M of its value. The D factor forces a hard look at what your startup is really worth in the face of foundation model competition.

Many founders overestimate their defensibility. That "data moat" you're banking on? If it's just publicly available data, by definition, it's not a moat. And investors are starting to see through the hype. They're no longer paying sky-high multiples for applications with a D factor north of 0.9, regardless of the buzz. We saw this dynamic playing out across recent YC (YCombinator) batches; the best companies with an edge were getting funded at high valuations ahead of demo day, while those towards the lower end struggled to fill their round even after demo day. The market is shifting, and valuations are adjusting to reality.

Of course, by definition, no seed-stage company has a true moat. But they should have a reasonable plan to build a moat. A seed investor's decision should be based on that plan. Does the company's product offer a clear differentiation from the models themselves? Will the product be deeply integrated into a user's workflow? Is there a way to build a compounding data advantage? These are questions that early-stage investors must ask to assess the D factor. For later-stage investors, we summarized the moats to look for at each stage in the Appendix.

What Should the Valuation Multiples Be for AI Applications?

In the same vein as Part I, we'll use SaaS valuation metrics as a starting point:

```
Base Multiple = Growth Rate x NRR x 10
```

```
Valuation Multiple = Base Multiple x Gross Margin Adjustment
```

Unlike foundation models, there are four main reasons why we actually believe that *AI applications deserve a higher valuation multiple than SaaS products*.

First, AI companies have superior growth trajectories. AI applications often target transformative use cases that can disrupt entire industries, leading to exponential growth rates that outpace traditional SaaS. While some companies generated “vibe revenue,” many, if not most, AI applications unlock millions in revenue by solving real pain points that traditional SaaS cannot. Within our portfolio, we have seen companies like Sully AI grow to millions in ARR by creating agentic workflows for specific needs in the healthcare vertical.

Second, AI companies are more capital efficient than traditional SaaS, scaling revenue without scaling headcount. Traditional SaaS typically requires substantial investments in sales, marketing, and customer success to drive growth, whereas AI companies leverage their own technology to enhance engineering productivity and automate internal workflows, creating a virtuous cycle of efficiency.

Third, AI companies, especially agent companies, have higher gross margins. This margin advantage stems from several factors: reduced human intervention in service delivery, automated scaling capabilities, and the ability to continuously improve performance through data feedback loops. As these systems mature, the margin advantage over traditional SaaS becomes increasingly pronounced, creating substantial long-term value for investors.

Finally, and perhaps most importantly, a fundamental shift is also occurring in the business model itself: from Software-as-a-Service to Service-as-Software. AI companies can now charge service-level fees because their agents effectively perform jobs previously done by humans, allowing them to tap into salary pools rather than just software budgets. This expands their total addressable market significantly. For example, our portfolio company Amy.vc is targeting the AI analyst market, replacing functions traditionally performed by junior analysts while delivering consistent, scalable insights at a fraction of the cost.

These arguments are frequently made by VCs to justify higher valuations for AI companies. However, we must consider an important counterbalance: the D factor. The looming threat of foundation models means that while the growth potential, capital efficiency, and margin profiles of AI companies may justify premium multiples, their "true valuation" may or may not exceed traditional SaaS when accounting for this heightened obsolescence risk. The valuation equation thus becomes more nuanced than a simple multiples comparison.

The above analysis points to a deeper truth about AI application valuation: the D factor captures external threats from model companies and competitive pressures, while revenue quality reflects internal factors about operational efficiency and market expansion. Companies that can both minimize their D factor through defensible moats and maximize revenue quality through operational excellence will likely command premium valuations. Those that successfully replace high-value knowledge work with AI-driven solutions may eventually achieve multiples far exceeding traditional SaaS companies, but the path to proving this sustainable advantage remains long and uncertain.

In the next section, we'll offer a case study for how a higher valuation multiple might not overcome the threat of the D factor.

Case Study: Why AI SDRs Might Be Overvalued

Nothing illustrates the danger of confusing early traction with durable value quite like the current boom in **AI-powered sales-development-rep (SDR) tools**. In less than two years, investors have poured hundreds of millions of dollars into startups that promise autonomous prospecting, email drafting, and even cold-calling bots. At first glance, the momentum looks unstoppable: In September 2024, 11x.ai closed a $50 million Series B that valued the company at roughly $350 million. The company reportedly achieved $10M in ARR. Within the next few months, Regie.ai followed with a $30 million Series B, and Artisan announced a $25 million Series A to build similar AI sales agents.

Yet when we run these businesses through our investment framework, the picture shifts dramatically. Most of the companies in this category face a D-factor in the 0.7-0.9 range, meaning as much as 90% of headline value could vanish once displacement risk is priced in. Four structural forces drive that risk.

1. **Data moats are paper-thin.** AI SDR products fine-tune on the customer's own CRM data plus public email corpora; by definition, the public data is neither exclusive nor permanent. Once an incumbent model taps the same objects via an API call, the advantage evaporates.
2. **Switching costs hover near zero.** SDR managers live and die by reply-rate dashboards. If a tool's performance dips, they churn within weeks. A recent TechCrunch piece revealed that 11x was losing 70–80 % of the companies that trialed its product after just a single three-month cycle.
3. **Platform bundling erodes pricing power.** Because the underlying technology is not difficult to build, larger software platforms can clone the feature and bundle it at marginal cost. When an incumbent CRM or marketing-automation suite folds an AI SDR into its subscription tier, stand-alone pricing power takes a significant hit.
4. **ROI decays as everyone deploys the same outbound playbook.** The more ubiquitous AI-generated outreach becomes, the faster prospects develop "inbox antibodies," pushing conversion rates down and cancelling the very cost savings that fuel the narrative.

The moral is not that AI for sales lacks merit; rather, it shows how systematic over-valuation creeps in whenever investors anchor on adoption curves and ARR run-rates without calibrating for displacement risk. For founders, the lesson is to move beyond generic email automation toward proprietary data loops, deep workflow control (think pricing negotiation, quote creation, contract red-lining), and compliance features incumbents cannot replicate quickly. For investors, the case is a vivid reminder that **high ARR multiples cannot (and should not) ignore a brutal D-factor**. Once you discount for the likelihood of foundation-model or platform absorption, many

of today's "AI SDR" darlings look far less like the next Gong and far more like transient features waiting to be bundled away.

Outrunning AGI: How Specialized Apps Can Win Even as Models Race Ahead

Looking ahead, we do not expect AI progress to slow down; if anything, it will continue to accelerate. While this acceleration will bring about disruption for many AI companies, it will also unlock new opportunities, especially as we march towards artificial general intelligence (AGI).

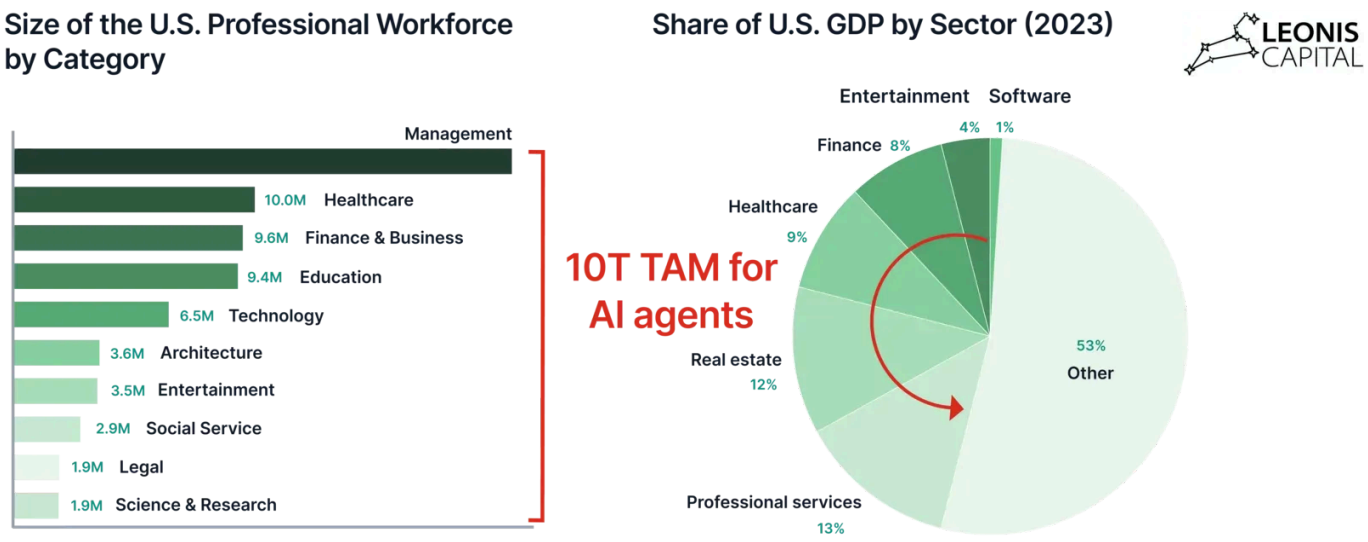
Some predict that foundation models will soon engulf the AI landscape, rendering application-layer startups obsolete as they march toward AGI.

This concern is legitimate. Compounding the absorption risk for AI applications is the accelerating speed at which new foundation models are able to do more. Research shows that the length of tasks AI can handle is doubling every 7 months, with each leap in progress more potent than the last. If these trends hold, by 2027 AI agents could handle task lengths of 1 work day, and by 2028 handle task lengths of 1 work week. Moreover, automated AI R&D could shrink innovation cycles dramatically, creating a "software intelligence explosion."

According to this line of thinking, the window of opportunity for applications to differentiate and build defensibility is growing narrower. The question is no longer if foundation models will replicate your startup's functionality, but when and whether you can build sufficient value beyond the model's capabilities before that happens.

We disagree with the doomsday scenario for AI applications.

In fact, **AGI might actually help application companies unlock new markets. We think one of the first markets AGI will unlock is the vast pool of tasks still paid for out of payroll, consulting, and professional-services budgets**. The "labor TAM" is several trillion dollars and orders of magnitude larger than every corporate software line item combined. This shift will also push pricing away from seat-based SaaS toward outcome-based contracts that resemble human labor. Such performance-linked deals can accelerate revenue far faster than classic subscriptions, but they also carry less baked-in stickiness. At that point, investors will need to once again rethink how AI companies are valued.



Source: U.S. Bureau of Labor Statistics (2023).

Source: U.S. Bureau of Economic Analysis (2023).

Figure 4: AI agents can theoretically address most, if not all, of the white collar market in the U.S.

As competition in AI continues to increase, the zero-value threshold will indeed claim those who chase broad markets or rest on static moats. But for those who relentlessly innovate within their vertical, cultivate deep domain expertise, and solve genuine pain points, the future holds not just survival but access to vast, previously untapped markets worth trillions in human labor value.

Jenny Xiao is a Partner at Leonis Capital. She was previously a researcher at OpenAI and received her Ph.D. from Columbia University.

Liang Wu is a Senior Researcher at Harvard Business School, where he focuses on the evolution of business models with emerging technologies.

Jay Zhao is a Partner at Leonis Capital. He is a long-time VC with 3 IPOs and 8 unicorns in his portfolio.

Appendix: Moat Development for AI Companies

Having established that narrower, more specialized AI applications often have stronger defensive advantages, we now turn to the question of how and when these advantages truly take shape. In other words, at what point does a moat become more than just a pitch deck claim?

The timeline for developing a genuine moat is frequently longer and more complex than many AI founders anticipate. Below is an exhibit that highlights four distinct stages in moat formation, ranging from the earliest seed phases, when the threat of foundation model displacement is at its highest, to a fully proven, mature moat.

Moat Stage	Typical Funding	Moat Strength	Defining Characteristics	Key Risk(s)
Pre-Moat Stage	Seed to Early Series A	Low (mostly theoretical)	Bold ideas, early demos, minimal data or workflow integration	Very high displacement risk if foundation models advance too quickly
Moat Construction	Late Series A to Series B	Emerging (partially validated)	Initial paying customers, domain footholds, some data advantage	Still vulnerable if feature can be absorbed by next foundation model
Moat Validation	Series B to Series C	Moderate to high (tested by market)	Demonstrated resilience against new model releases, measurable retention	Quantifiable displacement risk with real data showing decreasing churn, increasing engagement
Mature Moat Stage	Series C+ and Beyond	High (proven over multiple cycles)	Deep industry integration, strong brand presence, regulatory or data barriers	Sustainable moat weathering several LLM leaps

Pre-Moat: High Ideals, High Vulnerability: A young AI startup typically begins in the Pre-Moat phase, touting future data network effects or regulatory advantages. Yet actual proof points, like paying customers or validated workflow integrations, are often scarce. In this stage, a sudden leap in foundation model capabilities can easily render an AI application obsolete, because the startup's claimed "moat" is still more promise than reality. For founders, the key imperative is to convert theoretical advantages into something real and do so faster than a competing LLM can catch up. For investors, the lesson is to discount lofty projections and focus on early signs that the startup is moving toward genuine defensibility (e.g., first customers, regulatory progress).

Moat Construction: Building Real Defenses, But Still Fragile: Once a startup acquires paying clients and embeds its solution into real workflows, it enters the Moat Construction stage. Domain-specific capabilities like partial FDA clearance or proprietary data pipelines start to form an actual defense. However, these defenses can remain fragile if a foundation model rapidly gains the same capability. For founders, this phase calls for aggressively reinforcing whatever edge they have (deeper integrations, more specialized data, or further compliance milestones) before a disruptive LLM upgrade emerges. For investors, it's time to look closely at the pace of moat-building and whether the startup can stay ahead of foundation model developments.

Moat Validation: Data-Backed Staying Power: A startup progresses to Moat Validation when it weathers at least one major LLM release without mass user churn. At this point, retention and renewal data prove that customers genuinely prefer the specialized solution over free or cheaper base models. For founders, validated retention metrics become powerful evidence in fundraising or partnership discussions, showing that the product isn't just a feature waiting to be absorbed.

For investors, it's a milestone that justifies a higher valuation multiple since there is now tangible proof that the startup's moat is more than marketing spin.

Mature Moat: Premium Valuations and Lower Displacement Risk: Finally, companies that maintain stable (or growing) revenue in spite of new foundational model versions can achieve a Mature Moat. They enjoy strong brand recognition, deep workflow integrations, and possibly exclusive regulatory advantages. For founders, this stage opens the door to premium partnerships, large-scale expansions, or lucrative exits because their defensibility has been demonstrated over time. For investors, the threat of immediate displacement is significantly reduced, justifying a valuation premium that reflects the company's proven resilience in a rapidly evolving AI market.

FOOTNOTES

1. We have heard mixed reviews of Harvey, with a few lawyers telling us that although their firm has a Harvey subscription, they still prefer to use generalist AI tools like ChatGPT or Perplexity. This goes to show that even vertical tools in difficult domains face competitive pressure from foundation models.