The Ten Moats of the University: What LLMs Leave Behind

Phronesis AI

2026-02-17T12:36:13+01:00

An analysis following Nic Bustamante's diagnosis of the vertical software collapse

The Starting Point

In recent weeks, nearly a trillion dollars of market capitalization has been destroyed among software and data companies. FactSet fell from 20 billion to under 8 billion. Thomson Reuters lost almost half its market cap. The trigger: Anthropic released industry-specific plugins for Claude Cowork, an AI agent for knowledge workers.

Nic Bustamante, founder of Doctrine (Europe's largest legal information platform) and Fintool (AI-powered equity analysis), has described the anatomy of this collapse. His thesis: LLMs are systematically destroying the moats that made vertical software defensible. But not all of them.

Bustamante identifies ten moats. Five are being destroyed. Five hold firm. The crucial question is which are which.

I want to apply this analysis to an institution rarely conceived of as "vertical software," but which functions structurally in similar ways: the university.

What Is a Moat?

The term comes from investor language. A moat is a structural advantage that prevents competitors from entering a market. High switching costs. Network effects. Regulatory barriers. Proprietary data.

Vertical software – Bloomberg for finance, LexisNexis for law, Epic for healthcare – is characterized by particularly deep moats. Bloomberg costs $25,000 per seat per year. Retention is at 95 percent. Customers pay a lot and rarely leave.

The university functions similarly. High costs (tuition, opportunity costs). High retention (one rarely switches alma maters). Strong lock-in effects (degrees are not portable). And a rhetoric of indispensability that makes any competitor appear illegitimate.

The question is: Which of these moats are substantive, and which are interface?

The Ten Moats of the University

I apply Bustamante's framework to the university. The diagnosis is sobering.

1. Learned Interfaces → Destroyed

A Bloomberg terminal user has spent years learning keyboard shortcuts, function codes, and navigation patterns. This investment is not transferable. Those who speak the language fluently won't switch, because switching means becoming illiterate again.

The university has equivalent interfaces: How does one write an academic text? How does one cite correctly? How does one structure a DFG (Deutsche Forschungsgemeinschaft – German Research Foundation) application? How does one navigate peer review? These skills require years of practice. They are not intuitive. They are a language.

LLMs collapse all proprietary interfaces into one: chat.

What a doctoral student learns in three years – citation styles, formatting templates, the implicit rules of academic writing – an agent can apply in seconds. The question "How do I write a grant proposal?" transforms from a competency into a prompt.

Interface competency was a moat. It no longer is.

2. Codified Workflows → Evaporated

Vertical software codifies how an industry actually works. A legal research platform doesn't just store rulings. It codifies citation networks, relevance indicators, the specific way a lawyer constructs a brief.

Bustamante describes the difference between Doctrine and Fintool. At Doctrine, the team built thousands of lines of Python over years, hand-tuned relevance models, domain-specific classifiers. At Fintool, the same business logic is a Markdown file. A portfolio manager who has conducted 500 DCF analyses can codify their entire methodology without writing a single line of code.

Years of engineering versus a week of writing. That's the shift.

The university lives on codified workflows. Doctoral regulations. Habilitation guidelines. Appointment procedures. DFG formatting templates. These workflows are complex, historically grown, and their mastery requires insider knowledge.

LLMs turn this insider knowledge into a file.

The DFG application that previously required a Fleißliese (diligent workhorse) – someone who knows the formatting templates, understands the implicit expectations, keeps the deadlines in mind – is now a skill you describe in Markdown. The Fleißliese was the human version of what an agent system with access to DFG documentation accomplishes in a fraction of the time.

3. Access to Public Data → Commodified

A large part of the value proposition of vertical software consisted in making hard-to-access data easily searchable. FactSet makes SEC filings searchable. LexisNexis makes case law searchable. These are genuine services. SEC filings are technically public, but try reading a 200-page 10-K in raw HTML.

Before LLMs, accessing this data required specialized software and significant technical infrastructure. Companies like FactSet built thousands of parsers, one for each document type.

LLMs make this trivial. Frontier models already know from their training data how to parse SEC filings. They understand the structure of a 10-K. You don't have to build a parser. The model is the parser.

The university had a similar monopoly: access to knowledge. The library. Journal subscriptions. Databases. Those not at a university had no access to JSTOR, no access to primary sources, no way to participate in scholarly discourse.

This monopoly has already largely eroded (Open Access, Sci-Hub, preprint servers). LLMs complete the erosion. They have internalized the library. You no longer need access to JSTOR if the model already knows the content.

4. Talent Scarcity → Inverted

Building vertical software requires people who understand both the domain and the technology. Finding an engineer who can write production-ready code and understand how credit derivatives are structured is extremely rare. This scarcity limited the number of serious competitors.

At Doctrine, Bustamante reports, hiring was brutal. Every week, lawyers gave internal lectures to explain the legal system to engineers. It took months before a new engineer became productive.

At Fintool, this doesn't exist. Domain experts write their methodology directly in Markdown skill files. They don't have to learn Python. They write in plain text what a good DCF analysis looks like, and the LLM executes it.

The university is based on the same scarcity: people who have mastered a field. "You have to know the field." "You have to have read the literature." "You have to understand the debates."

LLMs invert this scarcity. The model knows the field. It has read the literature. It understands the debates – at least well enough to cover 80 percent of cases.

Talent scarcity was a moat. Engineering is now trivially accessible. The scarce resource (domain expertise) can now become software directly, without the engineering bottleneck. The barrier to entry collapses.

5. Bundling → Weakened

Vertical software companies expand by bundling adjacent capabilities. Bloomberg started with market data, then added messaging, news, analytics, trading, and compliance. Each new module increases switching costs.

The university is the ultimate bundle: teaching, research, certification, socialization, network, career preparation, life phase. You don't buy "a course." You buy "studying." And

because everything is intertwined, you can't simply take the network part and leave out the lecture part.

LLM agents break bundling because the agent itself is the bundle.

At Fintool, Bustamante describes, alerts are a prompt. Watchlists are a prompt. Portfolio screening is a prompt. There's no separate module for each. The agent orchestrates across ten different specialized tools in a single workflow. The user doesn't even know which five services were queried.

What does this mean for the university? The agent can take teaching from one provider (online courses), research access from another (Open Access), network from a third (Twitter/X, Discord communities), certification from a fourth (alternative credentials). The incentive to buy the entire bundle evaporates.

This doesn't mean bundling is dead overnight. The operational complexity of managing ten vendor relationships instead of one is real. But the directional pressure is clear.

6. Proprietary Data → Stronger (but the university has hardly any)

Some vertical software companies own or license data that exists nowhere else. Bloomberg collects real-time price data from trading desks worldwide. This data was collected over decades, often through exclusive relationships. You can't just scrape it.

If the data is truly non-replicable, LLMs make it more valuable, not less.

The university has hardly any proprietary data. Research results are published. Teaching materials are largely standardized. What the university "owns" is not data, but reputation – and reputation is not a dataset.

The only exception: unpublished research, lab notebooks, negative results. But these are systematically not shared, not because they are valuable, but because the incentive system devalues them.

7. Regulatory Lock-in → Structurally Intact

In Bustamante's analysis: In healthcare, Epic's dominance isn't just product quality. It's HIPAA compliance, FDA certification, and 18-month implementation cycles. Switching EHR providers is a multi-year project that literally risks patient safety.

HIPAA doesn't care about LLMs. FDA certification doesn't get easier because GPT-5 exists.

The university has strong regulatory lock-in:

- State recognition of degrees
- Accreditation procedures
- Professional regulations (doctors, lawyers, engineers must have university degrees)
- Examination regulations with force of law
- BAföG (federal student aid) tied to enrolled students

As long as the state only recognizes accredited degrees, the university is indispensable as a certification body. This is not an interface moat. This is a regulatory moat. It holds.

The question is: For how long? Regulatory lock-in is politically changeable. If alternative credentials gain societal acceptance, if companies stop requiring degrees, if the state recognizes new certification pathways – then this moat erodes too.

But that's a question of decades, not quarters.

8. Network Effects → "Marriage Market & Bromances"

Some vertical software becomes more valuable the more industry participants use it. Bloomberg's messaging function (IB Chat) is the de facto communication layer of Wall Street. If every counterparty uses Bloomberg, you have to use Bloomberg. Not because of the data. Because of the network.

LLMs don't break network effects. If anything, they make communication networks more valuable.

The university has strong network effects, which a friend of mine summarized succinctly:

"The moats are really just networking in place aka marriage market & bromances."

That's brutally precise. The real careers are made: - In the cafeteria - At conferences - In office conversations - In appointment committees - In doctoral colloquia

These networks are not digitizable. They require physical co-presence, shared time, the slow accumulation of trust and mutual obligation.

An agent cannot simulate a "bromance." It cannot accelerate a career through a coffee in the right cafeteria. It cannot send and receive the subtle signals that determine who makes the shortlist.

This moat holds – for now.

9. Transaction Embedding → Partially Intact

When software sits directly in the flow of money – payment processing, lending, claims processing – switching means interrupting revenue. No one does that voluntarily.

The university is partially transaction-embedded: - Degrees are prerequisites for professions (doctor, lawyer) - Academic titles are prerequisites for academic careers - Publications are prerequisites for appointments

But the embedding is less deep than with Stripe or Bloomberg. You can get a job without a university degree – it's just harder. The transaction (career) doesn't flow through the university; it's only influenced by it.

10. System of Record → Long-term Threatened

When software is the canonical source of truth for critical business data, switching is not just inconvenient. It's existentially risky. What if data gets corrupted during migration?

The university is the system of record for: - Educational biography (CV) - Publication list - Academic reputation - Qualification credentials

But Bustamante warns: Agents are quietly building their own systems of record.

Agents don't just query existing systems. They read SharePoint, Outlook, Slack. They collect data about the user. They write detailed memory files that persist across sessions. The agent accumulates over time a richer, more complete picture of a user's work than any single system of record.

Agent memory becomes the new source of truth. Not because anyone planned it, but because the agent is the one layer that sees everything.

What does this mean for the university? If the agent documents my learning history, my projects, my skills better than any transcript – why do I still need the transcript?

The First Assessment

Moat	University	Status
Learned Interfaces	How to write papers, proposals, citations	☒ Dying
Codified Workflows	DFG formats, doctoral regulations	☒ Becoming Markdown
Data Access	Library, journals, databases	☒ Commodified
Talent Scarcity	"You have to know the field"	☒ Inverted
Bundling	Teaching + Research + Degree + Network	☒ Agent unbundles
Proprietary Data	Hardly any	☒ No protection
Regulatory Lock-in	Accreditation, examination regulations, titles	☒ Holds (for now)
Network Effects	"Marriage market & bromances"	☒ Holds
Transaction Embedding	Degrees for professions	☒ Partial
System of Record	CV, publication list	☒ Long-term threatened

Five moats destroyed or dying. Three wobbly. Two hold.

But the analysis is not yet complete. The moats that hold have deep structure.

The Hidden Dimensions: Intergenerationality and Place

"Marriage market & bromances" is brutally precise as a description, but analytically too shallow. There are two more dimensions buried within it that make the network moat more substantial than it appears at first glance.

Intergenerationality

The university is one of the few institutions where generations systematically encounter each other.

The 60-year-old professor sits with the 35-year-old postdoc, the 28-year-old doctoral student, and the 22-year-old undergraduate in one room. This is not trivial. In almost all other areas of life, we segregate by age:

- School: Peers
- Work: Similar career stage
- Leisure: Similar life phase
- Social media: Algorithmically filtered peer groups

The university forces vertical encounters. The 22-year-old sits in the 60-year-old's seminar. The 60-year-old reads the 22-year-old's work. Knowledge, attitudes, networks are passed down across generations.

What this intergenerational structure transports:

Tacit Knowledge. That which isn't in books. How to read a reviewer. Which journals count and which don't. When to disagree and when to nod. These things aren't learned from documents. They're learned by watching someone who has mastered them.

Style Formation. How does one think? How does one argue? How does one write? Every academic tradition has a style, and this style is passed down through imitation. The doctoral student learns not only *what* the professor thinks, but *how* they think.

Network Inheritance. The professor introduces the doctoral student to the colleague. The colleague becomes a reviewer. The reviewer becomes a mentor. Networks aren't built; they're inherited.

Historical Consciousness. What has already been tried? What has failed? Which questions are exhausted, which fruitful? This knowledge exists nowhere in writing. It exists in the minds of those who were there.

Books are the other intergenerational transmission technology. I read Kant, though Kant is dead. The book bridges generations asynchronously.

LLMs have internalized the books. They can tell me what Kant wrote. They can even tell me how to argue "in a Kantian way." What they cannot do: connect me with someone who lives the field. Who knows the gossip. Who knows who is feuding with whom, who is rising, who is falling, which topics make careers and which end them.

The intergenerationality of the university is not primarily knowledge transfer. It is network transfer. And networks cannot be prompted.

Place

The university is a real estate institution. This is chronically underestimated.

Oxford is not just a university. Oxford is a city built around the university. The colleges are real estate. The libraries are real estate. The cafeteria is real estate. The courtyard where you "accidentally" meet is real estate.

What the physical place accomplishes:

Co-presence Compulsion. You have to be there. Physically. With your body. This sounds

trivial, but it is not. Co-presence forces attention that is not enforceable digitally. Someone sitting in a seminar cannot simultaneously watch Netflix (at least not unnoticed).

Random Encounters. The cafeteria, the hallway, the library, the copy machine – these are machines for producing chance. You meet people you weren't looking for. These unplanned encounters are the raw material from which networks emerge.

Prestige Accumulation. Old buildings are solidified time. They say: This institution has existed for centuries. It will exist tomorrow too. This permanence is itself a form of legitimacy. You don't get your doctorate at a startup.

Neutral Ground. The university creates spaces where encounters can take place that otherwise wouldn't. The professor and the student meet in the seminar, not in the professor's office (where the hierarchy would be crushing) and not in the student's apartment (which would be inappropriate). The university is a third place.

The economic dimension is real. Universities increase the value of cities. Real estate around universities is more expensive. "I studied in Heidelberg" is also a statement about milieu, about origin, about belonging to a certain class.

What LLMs change about this: Knowledge has become location-independent. Networks have not. You can access Claude from anywhere. You cannot access the courtyard of Balliol College from anywhere.

Are there digital equivalents?

The tech scene claims: Yes. YC Demo Day is a "place" (even though it occurs physically). Twitter/X is a "place" (virtual, but with random encounters). Discord servers are "places."

But these digital places are stratified. You don't randomly meet the CEO of Anthropic in a Discord server. You might randomly meet them in the café next to Anthropic's office.

Physical places democratize chance. Digital places algorithmize it.

This is a crucial difference. The algorithm shows me what it considers relevant. The hallway shows me who happens to be walking by. The hallway has no opinion about relevance. The hallway is dumb. And precisely this dumbness makes it valuable.

The Extended Moat Analysis

With the hidden dimensions, the analysis becomes more differentiated:

Sub-Moat	What It Does	LLM-Resistant?
Network (horizontal)	Finding peers	⚠ Partially replaceable (Discord, Twitter)
Network (vertical/intergenerational)	Connecting generations	✓ Hard to replace
Knowledge Transfer (explicit)	What's in books	✗ Completely replaceable

Sub-Moat	What It Does	LLM-Resistant?
Knowledge Transfer (tacit)	How do you actually do it	⚠ Partially replaceable
Network Transfer	Who knows whom	✓ Not replaceable
Place as Co-presence Compulsion	You have to be there	✓ Not replaceable
Place as Random Generator	You meet people	✓ Hard to replace
Place as Prestige Accumulator	Old buildings = legitimacy	✓ Not replaceable

The university has more deep structure than the first analysis showed. But the deep structure lies in exactly three areas: Regulatory Capture, Intergenerational Networks, and Place.

Everything else is interface. And interface is dying.

The Three Moats That Remain

After the extended analysis, three substantive moats remain:

First: Regulatory Capture. The state only recognizes accredited degrees. Professional regulations require university degrees. As long as this holds, the university is indispensable as a certification body. This moat is political, not technological. It can change, but not through LLMs – through legislation.

Second: Intergenerational Networks. The systematic encounter of generations, the inheritance of networks, the transfer of tacit knowledge and style. No agent can replicate this because it depends on bodies, time, and relationships.

Third: Place. Physical co-presence, random encounters, accumulated prestige, neutral ground. The university as a real estate institution, as a built environment in which certain things become possible that wouldn't be possible elsewhere.

These three moats are interconnected. Place enables the intergenerational network. The network generates reputation. Reputation legitimizes the regulatory privilege. An attack on one moat weakens the others.

The 200 Lines That Destroyed 200 Billion

To understand the magnitude of the shift, a concrete example helps.

The Legal plugin in Anthropic Claude Cowork is technically a skill file of about 200 lines of Markdown. These 200 lines describe how to conduct legal research: which sources to consult, how to evaluate precedents, how to follow citation chains.

These 200 lines of Markdown destroyed approximately 200 billion dollars of market value at Thomson Reuters and RELX.

Not because the file is brilliant. But because it shows that the entire "accessibility layer" – the interfaces, parsers, workflows built over years – is now a commodity capability that ships with the model.

The university should ask itself: How many of its services are "accessibility layer"? How many hours of methodology seminars, writing workshops, library introductions are in truth interface training that an agent makes obsolete?

The honest answer is: very many.

The Fleißliese Question

In the analysis of academic precarity, there is a figure who is particularly affected: the Fleißliese (diligent workhorse). Her entire value creation – writing proposals in proper format, organizing workshops, correcting footnotes, delivering on time – is exactly what Agentic AI automates.

The Fleißliese is the human version of what an agent system with access to DFG formatting templates, literature databases, and calendars accomplishes in a fraction of the time.

This sounds like liberation. Finally time to think! But within the system, it's a catastrophe. Because the invisibility of the Fleißliese wasn't a bug; it was her survival protection. As long as she was indispensable, she was untouchable. As soon as an agent takes over her function, she is not liberated but dispensable.

The system never valued her for her thinking. It won't suddenly start now that she has time.

What the Word Anecdote Shows

A personal experience that illustrates the argument:

Last week: spent an hour fumbling around in Word, couldn't get it done. The question: Why can't I just tell Word what I want?

The solution: Threw the Word file into Claude Code. Said: "Generate a signature for me, make it transparent, insert it where the signature should go, convert that to PDF, upload it to NextCloud, create the share for Christin, and write me the Signal message."

The agent does all of this without complaint.

Word has become pointless. Not because it doesn't work. But because the interface – clicking, formatting, exporting – has been replaced by delegation. You no longer interact with the tool. You tell the agent what you want, and the agent interacts with the tool.

The university is full of "Word." Complex systems that require interface competency: HIS, Moodle, DFG portal, examination administration, library catalogs. These systems won't be replaced. They'll be bypassed. The agent interacts with them so the human doesn't have to.

And if you consider that the text flowing through these systems is increasingly AI-generated as well: The interface is becoming obsolete from both sides. The input is AI. The processing is AI. Only the system in between is still human-made – and waiting to be bypassed.

The Scenarios

How might this develop?

Scenario 1: Slow Decline. The university gradually loses its interface moats but retains regulatory capture, intergenerational networks, and place. It becomes smaller, more expensive, more elite. An institution for those who can afford the networks and credentials. The broad middle erodes. Mass universities become certification factories. Elite universities become exclusive clubs.

Scenario 2: Disruption from Outside. Alternative credentials gain acceptance. Companies like Google, Apple, IBM accept non-university qualifications. The regulatory moat erodes because employers no longer demand it. The university loses its certification monopoly. What remains are networks and places – but these can be organized differently (see tech accelerators, artistic residencies, etc.).

Scenario 3: Reformation from Within. The university recognizes that its remaining moats are not interface but network, intergenerationality, and place. It transforms into an institution that explicitly offers these functions – less lecture, more colloquium; less examination, more project; less transmitting knowledge, more connecting people. The university as curated meeting place.

Scenario 4: Bifurcation. Top universities (with strong networks, strong reputation, and historic places) survive and become more valuable. They become the "Bloomberg Terminals" of education: expensive, exclusive, indispensable for a small elite. The rest collapses or becomes pure certification factories competing with MOOCs and AI tutors – and losing.

Scenario 5: The State Intervenes. Regulatory lock-in is strengthened, not weakened. The state protects the university through tightened requirements, professional regulations, accreditation hurdles. The university survives not because it's better, but because the state wills it. This is not a utopian scenario. It's what happens in many regulated industries: The incumbents write the rules.

The Test

Bustamante's framework ends with a test. For every vertical software company, three questions:

1. Is the data proprietary?
2. Is there regulatory lock-in?
3. Is the software embedded in the transaction?

Zero "yes" answers: high risk. One: medium risk. Two or three: probably safe.

For the university:

1. Is the data proprietary? **Hardly.** Research is published, teaching is standardized.
2. Is there regulatory lock-in? **Yes, still.** Degrees are state-recognized, professional regulations require them.
3. Is the institution embedded in the transaction? **Partially.** For some careers, the univer-

sity is mandatory; for others, it isn't.

That's one "yes" answer, one half, and one negative.

According to Bustamante's framework: medium to high risk.

But the extended test adds:

4. Are there intergenerational network effects? **Yes.**
5. Is the institution tied to a place with accumulated prestige? **Yes, for some.**

The answers vary dramatically between institutions. Oxford has all five moats. A provincial university of applied sciences may have only regulatory lock-in left.

The Question That Remains

The university doesn't survive because it is indispensable. It survives because the state protects it, because people need people, and because places accumulate prestige.

Both – all three – can change.

Regulatory lock-in is politically changeable. Networks can be built differently. Places can be designed differently.

If all three moats wobble, nothing remains.

The interface moats – how to write a proposal, how to cite correctly, how to navigate the system – are already dead or dying. They were never the real value. They were the barrier that kept competitors out.

What remains is what was never interface: the encounters, the relationships, the time, the space.

The question is whether that's enough.

This text emerged from a WhatsApp conversation, was structured by an LLM, and published via a CI/CD pipeline. It is itself an example of the shift it describes: The entire workflow – from idea to published text – touched no institutional infrastructure. No editing, no publisher, no peer review. Just a conversation, an agent, and a pipeline.

The university should ask itself: How much of what it does is this text – and how much is the conversation that preceded it?