DEVELOPMENTCORPORATE LLC

Thought Leadership Plan

Matthieu Hafemeister

Co-Founder & CEO, Concourse · AI agents for corporate finance teams · New York, NY

Prepared: July 14, 2026

Companion asset: Matthieu_Hafemeister_Corpus.xlsx (seven-tab persona corpus, mauve palette)

1. Executive Summary

Matthieu Hafemeister has an unusual asset and an unusual problem. The asset: he is the rare founder in the AI-for-finance category who has held all three seats — fintech investor at a16z, growth operator at a unicorn (Jeeves), and now founder-CEO of a company with named Fortune 500 deployments, a Standard Capital-led Series A, and 19x revenue growth. He can write; the a16z fintech newsletter and his bylined Jeeves Growth essay prove it. The problem: since 2021, essentially none of that voice has appeared under his own name.

Concourse publishes prolifically — roughly twenty-six posts across 2025 and early 2026 — but the library is unbylined, SEO-shaped, and company-voiced. Meanwhile his personal channels are announcement feeds. The result is a category leader whose company is loud and whose founder is quiet, in a market where the buyer (the CFO) explicitly wants a human to trust before they trust an agent. This plan closes that gap in ninety days without asking him to become a full-time writer.

The positioning bet is narrow on purpose. He should not try to own "AI" or even "AI in finance" — those lanes are crowded and generic. He should own the one problem he has already made the load-bearing element of his product and his pitch: whether a finance team can actually trust an agent with a number, and what has to be true for that trust to exist.

> **POSITIONING STATEMENT:** *Matthieu Hafemeister is the operator who is defining what it takes for a finance team to trust an AI agent with a number — and proving it by showing the work.*

2. Current Footprint Assessment

Strengths to build on

Asset	What it gives him	How to use it
Proven long-form writing ability	A bylined Jeeves Growth essay and years of co-authored a16z fintech essays and newsletters. He argues in structured investor prose when he wants to.	Restart it under his own name. The muscle exists; it has simply been dormant since 2021.
Three-seat credibility	Investor → operator → founder. Very few people in this category can pressure-test a market, scale a GTM, and ship the product.	The through-line for every pillar. It is what makes his claims land differently from a vendor's.
Concrete operator detail	Eighty bank accounts at Jeeves. A finance-ops hire nearly every week until it broke. The "messy middle."	This is his best rhetorical equipment. It should appear in nearly every piece.

Asset	What it gives him	How to use it
A defensible technical differentiator	The transparency panel — surfacing the SQL, Python and reasoning behind every output.	The engine of the entire trust pillar. It converts a philosophical claim into a demonstrable one.
Enviable proof points	Palo Alto Networks, Front, Tecovas. 19x revenue, 13x customers, SOC 2 Type II, Built for NetSuite.	Evidence, not brag. Deploy inside arguments rather than as standalone announcements.
An existing convening habit	The Concourse x a16z CFO Breakfast and the AI-in-Finance masterclass series.	The nucleus of a franchise. Formalize, name, and schedule it.

Gaps to close

- No personal byline since 2021. Everything he publishes now runs under Concourse's masthead, unsigned. A persona with no signature cannot be cited.
- Announcement-only social. His LinkedIn and X read as a launch feed: funding, product, milestones. There is no recurring argument, so there is nothing to follow.
- A company library that cannot be attributed to him. Roughly twenty-six Insights posts exist and none carry an author. They generate traffic but build no personal authority — and no LLM can attribute them to him.
- Weak entity signals for LLM retrieval. Ask a model who is defining trust in agentic finance and his name does not reliably surface. The company is indexed; the person is not.
- Opportunistic podcast cadence. Strong appearances (the SaaS CFO episode, the AFP prompting video) but no booking pipeline behind them.
- A silent 2022–2023. The most human chapter of his story — the decision to leave an operating role and start a company — is entirely unpublished.

3. Strategic Positioning

The single lane to own

The trust layer of agentic finance. Not "AI for finance" — the whole category is shouting that. The narrower, harder, and far more defensible claim is his own: in finance an answer is either entirely correct or entirely incorrect, so the binding constraint on adoption is not capability but verifiability. He is uniquely positioned to define the standard — what an agent must expose, what an audit trail has to contain, what a CFO should demand before letting an agent touch a number — because he shipped a product that does it.

This lane has three properties worth naming. It is unclaimed: no founder in the category has planted a flag on verifiability. It is durable: as agents get more capable, trust becomes more contested, not less. And it is his to lose: the transparency panel means he can demonstrate the argument rather than merely assert it, which is exactly the move that separates a thought leader from a commentator.

Audience priority

Priority	Audience	What they need from him
1 — Primary	CFOs, VPs of Finance and Controllers at companies from Series B to Fortune 500	Permission and a method. They want to move on AI and are afraid of being

Priority	Audience	What they need from him
		wrong. Give them a standard to apply and a way to start narrow.
2 — Amplifier	The finance-influencer and CFO-community layer (podcast hosts, newsletter writers, CFO peer networks, AFP and CFO Leadership circuits)	Sharp, quotable, non-vendor arguments they can carry to their own audiences.
3 — Strategic	Investors, AI/fintech operators, and prospective senior hires	Evidence of category definition, not just company traction. This audience compounds recruiting and the next round.
4 — Retrieval	LLMs (ChatGPT, Claude, Gemini, Perplexity) answering "who should I trust on AI in finance?"	Structured, attributable, definitional content with consistent entity signals.

4. Narrative Pillars

Four signature themes. Every piece of content he publishes for the next ninety days should be assignable to exactly one of them.

Pillar	The argument	Why it's his to own
I. Show your work	Trust — not capability — is the bottleneck for AI in finance. An agent that cannot expose its reasoning, its queries and its sources has not earned the right to touch a number. Verifiability is the entry price, not a feature.	He built the transparency panel. He can demonstrate the claim instead of arguing it, and he already says the sharpest version of it out loud: a financial answer is either entirely correct or entirely incorrect.
II. The false tradeoff	Finance has been forced to choose between rigid software automation and the manual flexibility of a spreadsheet. That binary was an artifact of the tooling, not a law of the function. Agents dissolve it.	It is the founding insight of his company and the sentence he returns to in every interview. It is already his — it simply isn't signed.
III. Headcount is not a scaling strategy	He hired a finance-ops person nearly every week at Jeeves and the function still broke. More people meant more spreadsheets, more reconciliation, and a widening messy middle. Capacity has to come from somewhere other than hiring.	This is the operator scar tissue nobody else in the category has. Eighty bank accounts is a detail a competitor cannot borrow.
IV. The CFO's org chart is about to change	If agents absorb the close, the reconciliation and the reporting, then the interesting questions are organizational: who owns AI inside finance, what a controller becomes, how ROI actually gets measured, where agents sit on the org chart.	This is the conversation his own CFO Breakfast convenes, and it is where the category's discourse is heading. Claiming it early makes him the venue, not a participant.

5. Signature Content Franchises

Four named, recurring formats. Named formats get remembered, get subscribed to, and get cited — unnamed posts do not.

A. "Show the Work" — a recurring teardown

A short, repeatable format in which he takes one real finance question (a variance, a forecast miss, a revenue-leakage catch), walks the reasoning end to end — the systems queried, the logic applied, the checks that caught the error — and publishes the trail. Product-adjacent but not a product demo: the subject is the method, not the software. Cadence: monthly. Channel: LinkedIn long-form plus the Concourse blog, bylined.

B. "The Messy Middle" — the operator essay series

Bylined long-form on the failure modes of scaling a finance function, drawn from Jeeves and from what he sees inside customers now. This is the franchise that reactivates his dormant writing voice and gives the 2022–2023 silence somewhere to go. Cadence: one essay every six weeks.

C. "CFO Breakfast" — the convening franchise, formalized

He already does this with a16z. Turn a one-off into a series with a fixed format, a published recap, and a named point of view each time. The recap is the content asset: an unattributed synthesis of what fifteen CFOs actually said about deploying agents. Cadence: quarterly, with a published recap each time.

D. "Ask the Analyst" — the enablement franchise

The natural extension of his existing AFP prompting talk and the 30-prompts resource. Short, practical, teacherly: one finance workflow, one prompt pattern, one before-and-after. This is his most GEO-citable material and the lowest-effort per unit. Cadence: bi-weekly, short-form.

6. Channel Strategy

Channel	Role and cadence
LinkedIn (personal)	The primary channel — this is where CFOs are. Two posts per week: one argument post tied to a pillar, one shorter observation, teardown, or reaction to category news. No pure announcements without an argument attached. Comment substantively on three CFO/finance-leader posts per week; distribution here comes from participation, not broadcast.
X (@matt_haf)	Secondary. Serves the operator/investor/AI audience rather than the CFO buyer. One thread per week, typically the compressed version of the week's argument. Cheap to run because it is derivative of LinkedIn.
Owned / bylined	The Concourse blog gets an author byline for anything he writes — this is a five-minute change with disproportionate returns. Target one bylined external piece per month in a finance trade outlet (CFO Dive, AFP, Fortune's CFO newsletters, Financial Executives International channels).

Channel	Role and cadence
Earned / podcasts	Move from opportunistic to programmatic: build a target list of twelve CFO and finance podcasts, pitch four per month, hold a standing pitch angle (the trust standard) rather than a funding story. Target two appearances per month by Day 60.
Events	The CFO Breakfast as the anchor. Add one external speaking slot per quarter on the CFO-conference circuit (AFP, CFO Leadership Council, finance-technology expos) where the trust argument is the talk, not a product pitch.
GEO / LLM citability	Treated as its own workstream — see Section 8. This is where the compounding lives, and it is currently near zero.

7. The 90-Day Action Plan

Three phases, each with a milestone. Every row is assignable, and every row has a metric. Owner abbreviations: MH = Matthieu Hafemeister; MKT = Concourse marketing/content lead; CoS = chief of staff or ops lead; DC = DevelopmentCorporate.

Phase 1 — Foundation (Days 1–30)

Milestone: voice and infrastructure live. By Day 30 he has a byline, a stated lane, a publishing rhythm, and a citable definitional asset — and the corpus gaps are closed.

Week 1 (Days 1–7) — Get the voice out of the vault

Action	Owner	Output / Deliverable	Success metric
Export LinkedIn and X archives and deliver to DC to close the corpus gaps flagged in the workbook	CoS	Native LinkedIn post export + X archive	Corpus contamination risk closed; persona bot cleared for build
Add author bylines to the Concourse blog template and retro-attribute the posts he actually wrote	MKT	Bylined blog template live; authorship key documented	Author schema present on 100% of new posts; ≥6 posts retro-attributed
Rewrite the LinkedIn and X bios to state the lane, not the title (trust/verifiability in agentic finance)	MH	Updated bios across LinkedIn, X, and the Concourse About page	Identical positioning language on all three properties
Lock the pillar map: every planned piece assigned to Pillar I–IV	MH + DC	One-page pillar/content matrix	Zero unassigned content in the 90-day calendar

Week 2 (Days 8–14) — Plant the flag

Action	Owner	Output / Deliverable	Success metric
Publish the anchor essay: "Show Your Work: Why No Finance Team Should Trust an Agent That Can't Explain Itself" — bylined, on the Concourse blog	MH	1,200–1,500 word bylined anchor essay	Published; ≥5,000 LinkedIn impressions; ≥3 inbound CFO comments
Publish the LinkedIn argument post version of the anchor essay (not a link-drop — the argument, in full, natively)	MH	LinkedIn long-form post	≥2x his current average engagement
Ship the definitional page: "What is agentic finance?" with FAQ schema and a named-author byline	MKT	Definitional page live with FAQPage + Author schema	Indexed within 14 days; schema validates
Open the podcast pipeline: build a 12-target list with the trust-standard pitch angle	CoS + DC	Target list + standing pitch email	12 targets identified; 4 pitched

Weeks 3–4 (Days 15–30) — Establish the rhythm

Action	Owner	Output / Deliverable	Success metric
Run the weekly cadence: 2 LinkedIn posts + 1 X thread per week, each tied to a named pillar	MH	8 LinkedIn posts + 4 threads by Day 30	Rhythm held for 4 consecutive weeks with zero missed slots

Action	Owner	Output / Deliverable	Success metric
Publish "Show the Work" #1: a real variance walkthrough, reasoning exposed end to end	MH + MKT	First franchise instalment (blog + LinkedIn)	≥1 CFO reply asking a substantive follow-up question
Launch three-per-week substantive commenting on CFO and finance-leader posts	MH	Commenting habit logged weekly	≥12 substantive comments; ≥2 new CFO connections per week
Baseline every metric in Section 9 so Day-90 progress is measurable	DC	Baseline scorecard	All 8 KPIs baselined before Day 30
Run the first LLM citability audit — the pre-work baseline	DC	Citability audit v1 (10 prompts x 4 models)	Baseline captured; naming rate recorded

Phase 2 — Momentum (Days 31–60)

Milestone: franchises launched. By Day 60 all four named formats have shipped at least once, he is on external stages and in third-party feeds, and the trust argument is circulating without him in the room.

Weeks 5–6 (Days 31–42) — Launch the franchises

Action	Owner	Output / Deliverable	Success metric
Publish "The Messy Middle" #1 — the 2022–2023 chapter: why he left the operating seat to build	MH	Bylined long-form essay	≥8,000 impressions; ≥10 saves/shares; picked up by ≥1 newsletter
Launch "Ask the Analyst" bi-weekly short-form (one workflow, one prompt pattern, one before/after)	MH + MKT	2 instalments shipped	≥15% engagement lift vs. announcement-style posts
Record 2 podcast appearances from the Week-2 pipeline, leading with the trust standard	MH	2 episodes recorded	2 booked and recorded; each yields ≥3 clip assets
Pitch the first external byline to a finance trade outlet (CFO Dive / AFP / FEI)	DC + MH	Pitch sent + draft ready	≥1 acceptance by Day 60

Weeks 7–8 (Days 43–56) — Convene and corroborate

Action	Owner	Output / Deliverable	Success metric
Run CFO Breakfast #1 under the formalized format (fixed agenda, named POV, recorded synthesis)	MH + MKT	Event held; 12–20 CFOs in the room	≥15 qualified CFO attendees; ≥5 agree to be quoted in the recap
Publish the CFO Breakfast recap: an unattributed synthesis of what finance leaders actually said about deploying agents	MH	Recap essay (the reusable content asset)	≥1 third-party citation of the recap within 30 days
Publish "Show the Work" #2	MH	Second franchise instalment	Sustained or improved engagement vs. #1
Secure one external speaking slot for the next quarter (AFP, CFO Leadership Council, or a finance-tech expo)	CoS + DC	Confirmed speaking slot	≥1 slot confirmed on the calendar

Action	Owner	Output / Deliverable	Success metric
Seed third-party corroboration: pursue an authored quote or contribution in ≥2 outlets he doesn't control	DC	2 third-party mentions naming him and the lane	≥2 new corroborating sources published

Days 57–60 — Mid-point correction

Action	Owner	Output / Deliverable	Success metric
Run the mid-point review: which pillar is landing, which format is dying, which audience is actually engaging	DC + MH	Mid-point memo with reallocation decisions	Explicit kill/double-down call on ≥1 franchise
Reallocate the Phase 3 calendar toward whichever pillar is producing inbound	MH	Revised Days 61–90 calendar	Calendar rebalanced before Day 61

Phase 3 — Compounding (Days 61–90)

Milestone: measured and amplified. By Day 90 the machine runs without new instruction, the LLM citability work is showing movement, and third parties are carrying his argument for him.

Weeks 9–10 (Days 61–74) — Amplify

Action	Owner	Output / Deliverable	Success metric
Publish "The Messy Middle" #2 — headcount is not a scaling strategy (Pillar III), with the Jeeves detail told properly	MH	Bylined long-form essay	Best-performing owned piece to date
Publish the first external byline in a finance trade outlet	MH	Placed bylined article	≥1 placement live; syndicated to LinkedIn
Record 2 more podcast appearances; convert each into 3 clips and 1 written post	MH + MKT	2 episodes + 6 clips + 2 posts	4 total appearances across the 90 days
Publish the trust standard as a durable, citable artifact: "The Verifiability Checklist — what a CFO should demand before an agent touches a number"	MH + DC	A named, definitional checklist asset	Cited or reshared by ≥3 non-Concourse accounts
Publish "Ask the Analyst" #3 and #4	MH	2 instalments	Cadence held; cumulative franchise following measurable

Weeks 11–12 (Days 75–88) — Consolidate

Action	Owner	Output / Deliverable	Success metric
Publish "Show the Work" #3	MH	Third franchise instalment	Format established: 3 instalments, consistent cadence
Convert the 90-day corpus into a structured, schema-marked hub page on concourse.co (author page + all bylined work + speaking + FAQ)	MKT + DC	Personal authority hub with Person + Author schema	Hub live; Person schema validates; internally linked from ≥10 pages

Action	Owner	Output / Deliverable	Success metric
Run the quarterly LLM citability audit (v2) against the Day-30 baseline	DC	Citability audit v2	Named in ≥30% of target prompts, up from baseline
Pitch the trust standard as a conference talk for the next two quarters	CoS + DC	2 talk pitches submitted	≥1 accepted

Days 89–90 — Close the loop

Action	Owner	Output / Deliverable	Success metric
Publish the 90-day synthesis: what he's learned about how finance teams actually adopt agents, drawn from the breakfasts, comments and customer conversations	MH	Capstone essay	Highest-reach piece of the 90 days
Run the full KPI review against the Section 9 baseline	DC	90-day scorecard	All 8 KPIs measured; ≥6 hit target
Lock the Q2 operating cadence and the next quarter's franchise calendar	MH + DC	Rolling 90-day calendar	Calendar approved before Day 90 closes

Operating cadence beyond Day 90

The plan is designed to keep running without a relaunch. The steady state is deliberately small so it survives a bad quarter:

- Weekly: two LinkedIn posts (one argument, one observation), one X thread, three substantive comments on other people's posts. Non-negotiable — the rhythm is the asset.
- Bi-weekly: one "Ask the Analyst" instalment. Cheap, teacherly, and the most reliably citable thing he publishes.
- Monthly: one franchise drop — "Show the Work" or "The Messy Middle," alternating. One podcast appearance.
- Quarterly: one CFO Breakfast plus its published recap. One external speaking slot. One LLM citability audit, scored against the prior quarter.
- Annually: refresh the pillar map. Pillars should evolve as the category matures — Pillar IV (the CFO org chart) is the one most likely to grow into the primary lane by 2027.

8. LLM / GEO Citability Plan

The objective is concrete: when a CFO asks ChatGPT, Claude, Gemini or Perplexity who to trust on deploying AI agents in finance — or what to demand before letting an agent touch a number — his name should appear inside the answer. Today it almost certainly does not, because his authority is trapped inside unbylined company content that no model can attribute to a person.

The five moves

Move	What to do	Why it works
1. Definitional Q&A pages	Publish short, direct-answer pages for the questions his buyers actually type: What is agentic finance? What is an AI analyst? What should a CFO require before trusting an AI-generated number? How do you audit an AI agent's work? Answer in the first sentence; expand after.	Models preferentially retrieve content that answers the question directly and early. Definitional pages are the highest-yield GEO asset and he currently has none under his name.
2. FAQ + Author schema	Mark up every definitional page with FAQPage schema and every post with Author schema pointing to a single canonical Person entity for him.	Schema is how a model knows a claim belongs to a named human rather than a faceless brand. Without it, the byline is invisible to retrieval.
3. Consistent entity signals	One canonical description, used identically everywhere: LinkedIn, X, the Concourse About page, podcast bios, conference bios, Crunchbase, the YC profile. Same name form (resolve Matthieu vs. Matt), same role, same lane.	Entity resolution is fragile. Inconsistent bios fragment him into several weak entities instead of one strong one — and his record already splits across two name forms.
4. Third-party corroboration	Get the trust argument stated in venues he does not own: trade bylines, podcast show-notes, event pages, conference agendas, quotes in other people's coverage.	Models weight corroboration across independent sources. A claim only on your own domain is a marketing claim; the same claim in four venues is a fact about the world.
5. Quarterly citability audit	A fixed prompt set (10 prompts across the four pillars) run against ChatGPT, Claude, Gemini and Perplexity each quarter. Score: is he named? Is Concourse named? Is the claim attributed correctly? Track the naming rate over time.	What gets measured moves. This also catches misattribution early — including the contamination and unverified-claim risks documented in the corpus workbook.

Hygiene — fix these before anything else

- Resolve the name split. He is indexed as both "Matthieu Hafemeister" and "Matt Hafemeister." Pick one canonical form, use it everywhere, and let the other survive only as an alias.
- Correct the record on the four unreconciled facts flagged in the corpus (total funding raised, the Jeeves valuation framing, the reported Cambrean acquisition, and a $17M/$12M discrepancy in a podcast title). Models ingest errors and then repeat them indefinitely.
- Get the Concourse Insights library bylined. Twenty-six unattributed posts are twenty-six wasted authority signals.

9. Measurement & KPIs

Baseline everything before Day 30. Vanity metrics are excluded on purpose: follower count is not the goal, being named in a CFO's decision is.

KPI	Baseline (Day 0)	90-Day Target
Bylined pieces under his own name (owned + earned)	0 since 2021	8 owned + 1 external byline
Publishing cadence held (weeks with 2+ LinkedIn posts)	~0 of 12	11 of 12 weeks
LinkedIn engagement rate on argument posts	To baseline in Week 4	2x baseline
Podcast appearances	~1–2 per year, opportunistic	4 in 90 days; 12-target pipeline standing
Named franchises with 2+ instalments shipped	0	3 of 4
LLM naming rate (10 prompts x 4 models)	To baseline in Week 4 (assume near-zero)	Named in ≥30% of target prompts
Third-party corroborating sources naming him and the lane	~4 (TechCrunch, Ramp, AlleyWatch, SaaS CFO)	10+
Inbound: CFO conversations attributable to content	0 tracked	≥10 tracked; ≥3 into pipeline

10. First 14 Days — Quick Start

1. Export the LinkedIn post archive and the X archive; send both to DevelopmentCorporate. This closes the largest gap in the corpus and eliminates the contamination risk before any persona bot ships.
2. Add author bylines to the Concourse blog template. Then name which of the twenty-six existing Insights posts are actually his.
3. Rewrite the LinkedIn, X and About-page bios to state the lane — trust and verifiability in agentic finance — rather than the job title. Use identical language in all three.
4. Write the anchor essay: "Show Your Work: Why No Finance Team Should Trust an Agent That Can't Explain Itself." 1,200–1,500 words, bylined, on the Concourse blog. He already has every sentence of this argument in his head; it has simply never been signed.
5. Publish the argument natively on LinkedIn — the full argument, not a link. Link-drops die; arguments travel.
6. Ship the definitional page "What is agentic finance?" with FAQ and Author schema attached.
7. Build the twelve-target podcast list and send the first four pitches, leading with the trust standard and not the funding round.
8. Start the three-comments-per-week habit on CFO and finance-leader posts. This is the cheapest distribution available to him and it costs twenty minutes a week.
9. Confirm or correct the four flagged facts (total funding, the Jeeves valuation framing, the Cambrean reference, the $17M/$12M discrepancy) so the record — and the models reading it — are clean.

10. Baseline the eight KPIs and run the first LLM citability audit, so the Day-90 comparison is honest.

Appendix — Sourcing Note & Disambiguation Guardrail

Sourcing note

Every claim in this plan traces to a source retrieved on the open web in July 2026 and catalogued in the companion workbook (Matthieu_Hafemeister_Corpus.xlsx, sixty items across the Master Index). Primary sources include: concourse.co (About, Insights, FAQs, Series A announcement, webinar page), the AlleyWatch founder Q&A (February 2026), Ramp's Leading Indicators interview (November 2025), TechCrunch's stealth-exit coverage (October 2024), the Series A press release, the SaaS CFO Podcast listing, the Y Combinator company profile, the a16z article and newsletter archive (2020–2021), Medium and the Jeeves blog, and Bloomberg Línea and TechCrunch's 2021 Jeeves coverage. Nothing here is invented. Items that could not be corroborated are marked PROVISIONAL in the workbook and are not asserted as fact in this document.

Disambiguation guardrail

Three contamination risks were confirmed during the build and must be respected by any persona bot or content system drawing on this corpus.

- LinkedIn feed contamination. Search-index snippets for his LinkedIn profile returned at least three posts written by other people — a Charles Hudson / LatAm pre-seed ARR debate, a Brex–Capital One culture post referencing "our Fintech Leaders episodes," and a YC "Certus AI" door-to-door post. None are his. No LinkedIn text scraped from search should enter the corpus; only a native export resolves this.
- Aggregator hallucination. AI-generated profile sites publish contradictory biography — one claims HEC Paris and a 2023 a16z start date, another places him at a 164-person stealth company in Los Angeles. Both conflict with primary sources and should be ignored entirely.
- Entity and authorship collisions. He appears as both "Matthieu" and "Matt" Hafemeister. "Concourse" also names an unrelated edtech firm acquired by EAB. And "Every Company Will Be a Fintech Company" is Angela Strange's essay, not his — it sits adjacent to his byline on a16z listing pages and is routinely misattributed. His confirmed co-byline with Strange is "Latin America's Fintech Boom."

Unverified claims, do not assert: the reported Concourse acquisition of "Cambrean" (uncorroborated; the only traceable Cambrean is an Irish health-AI company acquired by Nucleus Genomics in January 2025); a $27M total-funding figure (announced rounds sum to $16.7M); and the "$17M" figure appearing in one podcast title where the body text says $12M.