

CASE STUDY

Strella

Turning qualitative research
into an AI superpower

Discover:

- How Strella achieved 150% NDR on their first cohort of renewals
- Why the ideal pricing model makes usage frictionless before optimizing for revenue
- Building a durable moat not based in model selection

Bessemer Venture Partners

Copyright @ 2026



Strella: Transforming qualitative research from a bottleneck into an AI superpower

Strella's path to an AI platform that powers fully automated interviews 3

Key founder lessons 4

The challenge 5

The 'tremendous number of use cases' that became Strella's TAM

When voice AI caught up to the problem

The solution 6

From a co-founder posing as AI to a full-scale AI moderator

Strella's seat-based pricing strategy: removing friction to expand

The result 8

How Strella created a repeatable sales motion to go to market

The proof is in the renewals (plus Strella's vision for the future)



Organizations often run the same interviews over and over, but it doesn't have to be that way. Whether user research, customer discovery, market mapping, or expert calls—each study is expensive, slow, and doesn't scale. The limiting factor is people's time, not their lack of curiosity. Surveys are faster but fraught with fraud and rushing respondents. The result is that the decisions driving product roadmaps, go-to-market strategies, and market entry calls often rest on thinner customer insight than most want to admit.

Qualitative research has always been the most valuable and least scalable form of business intelligence. Consider a product manager who needs 50 customer interviews to validate a roadmap decision, or a consultant running the same expert call over 10 times on a single engagement. This is how Strella, the AI-moderated customer research platform, came to bet on voice AI early on, believing it would finally change the equation of these scenarios—not by replacing human insight, but by removing the human bottleneck from collecting it. Today, Strella is rapidly automating research projects for some of the most customer-obsessed companies, including Amazon, Duolingo, and Chobani.

As part of our case study series, *Launching AI products that win*, we sat down with the CEO and Co-founder of Strella, Lydia Hylton, to unpack how she and Co-founder Priya Krishnan successfully built and commercialized an AI-powered customer research platform that helps teams run customer research at scale in a fraction of the time. We explore how their firsthand experience with slow, manual qualitative research — across consulting, investing, UX research, and product management — inspired them to build Strella and the early product decisions that paved the way.

Strella's path to an AI platform that powers fully automated interviews

Situation	Lydia Hylton and Priya Krishnan met in high school and spent their early careers seeing the same problem play out across their experiences in management consulting, UX research, product management, and investing: understanding customers is mission-critical yet painfully slow. The duo founded Strella in July 2023, spending the first year navigating the idea maze—including Co-founder Priya pretending to be an AI moderator in interview sessions to gather intel—before landing on their core product in mid-2024.
Challenge	Existing tools weren't solving the right problem. Survey platforms were fast but produced questionable data quality. Tools marketed as "AI research" were essentially surveys with a voice recording layer. What was missing was something that could replicate the experience of a skilled interviewer (adaptive, conversational, and responsive) at scale.
Solution	Strella built an AI-moderated interview platform anchored by a fully agentic voice moderator that conducts real-time, voice-to-voice interviews—asking dynamic follow-up questions, processing verbal and nonverbal signals simultaneously, and sustaining conversations for up to 90 minutes. The platform also automates the surrounding research workflow: discussion guide setup, participant recruitment, and synthesis into actionable insights.
Result	All 12 of Strella's original design partners converted to paid at launch. Duolingo used Strella for video concept testing and completed a project in two days that had previously taken over six weeks, with insights reaching the C-suite. The company reached \$1.6M ARR in its first year of monetization, and <u>Bessemer led the company's \$14M Series A in October 2025.</u>



Key founder lessons from Strella

1. Bet on the technology curve, not the technology ceiling.

Strella launched mid-2024 specifically because voice AI had just crossed a quality threshold that made true voice-to-voice conversation possible. Timing the technology wave—not just recognizing the problem—was a deliberate strategic decision that gave them a meaningful head start on product quality.

2. If you feel ready to ship, you've waited too long.

Strella used a rigorous design partner program, instead of internal testing, to validate the product under real conditions. The standard wasn't perfection; it was whether strangers on the internet would give their time to an early product. When they did, and when everyone in the cohort converted to paid, that was the signal.

3. Remove every barrier to usage before optimizing for revenue.

Strella deliberately chose seat-based pricing to eliminate per-project budget approval within organizations. The bet: making it frictionless to run a Strella interview compounds adoption faster than maximizing revenue per project.

4. A strong moat is use-case-specific IP, not model selection.

Strella switches model providers regularly. Their durable advantage lives in the problems they've solved that off-the-shelf models don't address: like building a moderator that won't interrupt a participant mid-thought—a failure mode that's acceptable in customer support, but catastrophic in research contexts.

5. AI interviewers can unlock just as much data, if not more, as humans can.

Based on Strella's experience, participants can be measurably more open and candid with AI than with a human. Rather than sacrificing depth, Strella's approach has, in many cases, increased it—transforming an AI research conversation from "good enough" to "genuinely better."



The ‘tremendous number of use cases’ that became Strella’s TAM

For anyone who has run qualitative interviews professionally, you know the process is laborious: 1) schedule the interviews, 2) brief the moderator, 3) run the sessions, 4) transcribe the recordings, 5) synthesize the findings, 6) repeat. On a single consulting engagement during her time at Bain & Company, Lydia estimates that she was spending more than half her time on just the mechanics of research—not drawing conclusions from it, but executing the logistics around it. "I would run surveys, I would run expert interviews until the cows come home," she recalls. And this was at a firm where some of the world's largest companies were paying millions of dollars a month for exactly that kind of work.

The problem wasn't unique to consulting. As an investor, Lydia ran the same diligence calls repeatedly, asking the same questions to different people across a single deal. Her co-founder, Priya, had lived a similar experience from the product and UX side—a constant cycle of interview scheduling, note-taking, and synthesis that consumed weeks before a team could act on what they'd learned. Between the two of them, they had seen the same bottleneck play out across nearly every function that does qualitative research: product teams, strategy teams, consultants, and investors. "Think about how many instances people run basically the same interview 10 times, five times, however many," Lydia says. "That's our TAM, and that's a tremendous number of use cases."

The tools that existed before Strella didn't solve the problem; they reframed it. What passed for "AI-powered research" before was closer to an asynchronous survey with a voice recording layer. A question would appear on screen, a participant would record a voice note to answer, and then it would proceed to the next question. This was faster than a traditional survey, but it wasn't really an interview. There was no follow-up, no probing, and no ability to chase an interesting point from an answer. "Anyone who came before us has a souped-up survey experience," Lydia explains. "Sure, that's better than taking an hour-long one with all open-ended questions, but you don't really forget that you're taking a survey."

The deeper consequence Strella's co-founders witnessed was that customer insights—the foundation of good product decisions, sound strategy, and effective go-to-market—were being systematically deprioritized across organizations. And not because people didn't believe in the value of this kind of business intelligence, but because the infrastructure around it couldn't keep pace with what teams needed to do their jobs.

When voice AI caught up to the problem

Lydia and Priya didn't stumble into Strella; they spent a full year deliberately navigating the idea maze, testing directions, and waiting for the right convergence of problem and technology before committing. The unlock came when voice AI crossed a quality threshold that made something genuinely new possible. The tools that came before Strella weren't bad products—they were just constrained by the current technology of the time. By mid-2024, that constraint had lifted. Real-time voice-to-voice conversation at acceptable latency was finally viable, and Strella made an early, decisive bet that this would be the superior format. They leaned hard into the qualitative side of research rather than the quantitative, at a time when most of the market was still building variations on the survey model. What made the bet feel right wasn't just the technology; it was what they found when they tested it.



In those earliest experiments, before they had a real product or before any AI moderator had been built, they discovered something that would become one of Strella's most important product insights: "People are shockingly willing to talk to AI," Lydia says. The social friction that might have made an AI interview feel awkward or off-putting simply didn't materialize. Participants were engaging, opening up, and staying in the conversation. It also reduced the response bias in human-led interviews (e.g., wanting to please the interviewer or concern about judgment).

From a co-founder posing as AI to a full-scale AI moderator

Most founders wait until their product is ready—or at least minimally viable—before they find out if anyone wants it. Strella started before building any AI: its first "product" was Priya on a Zoom call with her video off. What followed was a product evolution that happened in three phases.

Phase 1: The 'Wizard of Oz' MVP

The MVP setup was deliberately lo-fi: Strella would recruit interview participants and schedule what appeared to be an AI-led session. When the call started, Priya would join camera off with the most robotic voice she could manage and conduct the interview herself, pretending to be AI. To simulate the latency of early voice AI, Priya waited a full five seconds after each participant finished speaking before responding. "AI was bad. The voices were bad at the time," Lydia recalled. The five-second pause wasn't a bug; it was a test. If participants were willing to engage with something that slow and robotic, they'd be willing to engage with the real thing. That human willingness to talk to AI was the foundational signal the team needed before investing in real infrastructure.

Phase 2: Strella's design partner program

With the core behavioral hypothesis validated, Strella moved into its design partner program in mid-2024. The program's structure was intentional and demanding. All 12 design partners had to come from cold LinkedIn outreach—no warm introductions, no existing relationships. "It's much harder to validate an idea if you need to sell a product to a stranger on the internet," Lydia said. "Ergo, if you can convince a stranger on the internet to be a design partner, you likely have something." Partners met with the team every other week, used early versions of the product, and gave structured feedback throughout. In exchange, they'd receive 50% off their first six months once Strella launched commercially. The program ended with a hard deadline: convert to paid, or don't. Every partner converted. Strella launched in early 2025 with a paying customer base built entirely from cold outreach.

Phase 3: Building the architecture for Strella's AI moderator


Building the AI itself turned out to be a different kind of challenge, one where the most important decisions weren't about which model to use but about which problems to solve that models couldn't. Strella's AI moderator is fully agentic, built on hundreds of specialized agents working in parallel—processing voice, facial expressions, and on-screen behavior simultaneously. The hard constraint that shaped this architecture? Reasoning models were off the table because latency requirements were too strict for research conversations. Instead, Strella compensates with agent breadth and use-case-specific problem solving, particularly around one of the hardest challenges in voice AI for research: building a moderator that won't cut off a participant mid-thought. In most voice AI contexts, the model estimates when a speaker has finished and responds above a probability threshold. In research, that failure mode can be catastrophic: participants often pause to think, and interrupting them can destroy the insight.

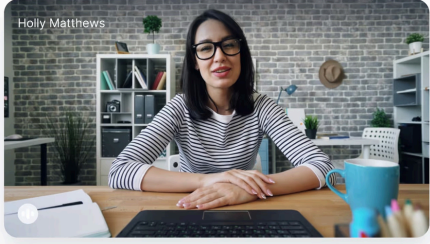


Soft Drink Preferences

AI Interview Online \$50 gift card 30 minutes

Device requirements
Computer (Mac/Windows), Webcam, Microphone

 **Watch a demo**
See how an AI-moderated session works



MacBook Pro Mic... MacBook Pro Spe... FaceTime HD Cam...

Join call

“It’s been the hardest thing we’ve done.” — Lydia Hylton on building the moderator.

"Making sure you marry all of those [signals] and actually have a cohesive response is pretty challenging," said Lydia. "It's been the hardest thing we've done."

Strella’s seat-based pricing strategy: removing friction to expand

Pricing an AI product is never just a monetization decision (hint: it's about what you believe drives adoption). For Strella, the choice to charge on a per-seat basis rather than per interview or per project was a deliberate bet on organizational behavior: make it frictionless to run research, and team use will expand.

"We wanted to remove every barrier to running a Strella interview that you possibly could have," Lydia explains. "One barrier would be if you need budget approval every time you run a project, we thought people wouldn't adopt us as quickly within an organization." A seat-based model sidesteps that entirely. Teams negotiate their spend once, upfront, and can run research as often as the work demands without going back to finance.

The pricing model also reflects something specific about Strella's cost structure that not every AI company can say. Many AI products are under real margin pressure from inference costs, which is why usage-based pricing often makes sense for them—it's how you protect your economics as consumption scales. Strella's situation is different. "Our cost per interview pales in comparison to the cost of paying the person for their time who's participating in the interview," Lydia notes. Because most research participants are compensated, and because that compensation far exceeds Strella's AI costs, the company doesn't face the same pressure to pass off inference costs to the customer. That structural advantage gave them pricing flexibility to prioritize adoption over per-unit revenue optimization.



The enterprise procurement reality reinforced their decision even further. Upper-market buyers often don't want to manage usage approvals on an ongoing basis; they want to negotiate annual spend, set expectations, and then simply use the product. "Some of our buyers just don't want to have to go through budget approval regularly," Lydia says. Seat-based pricing maps naturally to how those organizations budget, reducing friction at the point of sale as well as during day-to-day use. Within the seat model, Strella uses feature gating by tier rather than usage caps, which mirrors how research organizations actually function. Some users run research; others consume it. The tiers reflect that division rather than trying to optimize for maximum billing. As Lydia describes it, the tiers align with "the types of stakeholders you have in an organization — some people are responsible for the execution of research, and others want to do analysis."

For Strella, predictable revenue was a welcome byproduct of the model rather than the primary motivation, though Lydia doesn't discount its value: "Predictable revenue is worth more than unpredictable revenue." Strella has begun experimenting with usage-based pricing for customers who prefer it, and that evolution may accelerate as the product matures. But the seat-based foundation was the right call for early market penetration. It lowered barriers, encouraged usage, and helped establish a research habit inside organizations.

How Strella created a repeatable sales motion to go to market

Founder-led sales and staying close to the customer (mid 2024 to early 2025)

By mid-2024, the team was deliberately small: a designer, three engineers, one business generalist, Lydia running product and engineering, and Priya running go-to-market. After onboarding the 12 design partners who converted to paying customers, Strella launched its product in early 2025 and moved immediately into a founder-led sales motion. Priya and Strella's first business hire, Kelsey Young Brennan, ran sales directly through roughly the first half of the year, staying close to customer conversations while keeping a tight feedback loop with Lydia and the engineering team as they continued to build.

Every customer conversation informed product, positioning, and the team's understanding of which use cases were the stickiest. The design partner discipline of learning before scaling carried directly into how Strella approached early commercial traction.

Building a repeatable motion post-Series A (fall 2025 onward)

Strella raised its Series A in the fall of 2025, with Bessemer leading the round. Since then, the focus has shifted from founder-led selling to building a more traditional sales motion. Go-to-market headcount has grown to a team of 22 people, with most of the recent growth coming from the GTM side. The product and engineering team accounts for roughly 10 of those. As Lydia noted, "We do see that we're getting a lot further with a leaner engineering team given how good the tooling is."

As Strella scales, the company isn't just competing for existing research budgets—it's expanding the market by serving adjacent roles. Teams that had never stood up a formal research practice before are starting one because Strella has made the barrier low enough. That market-creation dynamic, driven by the product's accessibility rather than any deliberate GTM play, has become one of the more unexpected signals of the platform's long-term potential.



The proof is in the renewals (plus Strella’s vision for the future)

A true sign of product-market fit isn't just the first sale—it's also what happens when the contract comes up for renewal. For Strella, that moment arrived in early 2026, roughly a year after the company began monetizing. The results were telling: over 150% net dollar retention (NDR) on the first cohort of renewals, which is considered a best-in-class benchmark.

The company's growth trajectory tells a similar story. Strella went from zero to \$1.6M ARR in its first year of monetization, which is faster, Lydia noted, than any of the comparable players in the category had done in their first year. What the metrics don't fully capture is the breadth of how customers are actually using the platform. Strella was designed for structured research, but buyers have taken it further—running Amazon shopalongs, M&A due diligence calls, competitive website comparisons, live mobile app observations, and market entry decisions for new restaurant locations. That kind of creative adoption, unprompted and unpredictable, is typically what it looks like when a product has crossed over from a useful tool to a new capability.

Strella's long-term vision is a fully agentic research workflow where AI doesn't just collect insights from humans but routes them directly to other agents making decisions. As AI systems take on more autonomous decision-making, the demand for grounded human testimony—such as what real people actually think and why—will grow. Strella's thesis is that qualitative research infrastructure becomes more valuable in that world, and not less.

CASE STUDIES

Launching AI products that win

Leaders share how they successfully built AI products. [Subscribe to get more insights to help your team win.](#)



Venture insights that matter

bvp.com/subscribe

About Stella:

Stella is an AI-powered customer research platform that helps teams run in-depth interviews and generate actionable insights in hours, not weeks. Founded in 2023 by Lydia Hylton and Priya Krishnan, the platform automates four core stages of the qualitative research process (discussion guide creation, participant recruitment, AI-moderated interviews, synthesis) and supports studies across 46+ languages with access to a panel of 8 million global participants. Enterprise customers, including Duolingo, Amazon, Square, Daily Harvest, and Apollo, use Stella to conduct research at a speed and scale that was previously impossible. Learn more at stella.io.



Venture insights that matter

bvp.com/subscribe