Alex Johnson, co-founder & CEO of Velvet, on vertical AI for venture capital

TEAM

Jan-Erik Asplund
Co-Founder
jan@sacra.com

DISCLAIMERS

Alex Johnson, co-founder & CEO of Velvet, on vertical AI for venture capital

By **Jan-Erik Asplund**

Background

With the AI-powered financial analyst gaining steam with companies like <u>Hebbia</u> (a16z, $130M Series B) and Rogo (Thrive Capital, $50M Series B) raising big rounds, we reached out to <u>Alex Johnson, co-founder & CEO of Velvet</u>, an AI-powered financial analyst for venture capital to learn more about how venture investors get an edge in the age of AI.

Key points from our conversation via Sacra AI:

- **While public markets have Bloomberg and FactSet, <u>private markets and venture capital still lack a high engagement daily workflow anchor</u>, particularly given the heavy tool fragmentation across horizontal docs and spreadsheets like <u>Notion</u>, <u>Airtable</u> and Google Workspace, data for sourcing and evaluation like Harmonic and PitchBook, cap table like <u>Carta</u> and CRM like Affinity.** "If you're a broker or a trader, you're on Bloomberg and an exchange every day. And that's because all the research is in one place, all the data is in

one place, all the market participants are in one place. We're nowhere close to that today in the private markets, but that's the really long-term opportunity."

- **With AI transformation happening within venture firms right now, ~10% of the most data & engineering-centric VCs (Tribe Capital, SignalFire) and largest platforms (Insight) will roll their own stack—building their own sourcing & evaluation engines by combining their proprietary data with LLMs directly—while 90% will look for & buy a vertically integrated AI-for-VC solution like Velvet in VC or Rogo and Hebbia in PE.** "Horizontal agents are only competitive in the buy-and-build scenario, which most venture capital firms aren't doing. Most, 80% to 90% of them, would rather buy. And then the 10%, like the SignalFires, or the Insight Partners, the really big funds, are more likely to build their own stack because they have the resources to do that. So they might get a Glean or a Hebbia to work with enterprise knowledge search over their own dataset integrated into their own CRM that they probably built themselves. But the majority of the market's really not in that position. There's not really any VC-specific vertical AI tools that are hard in the market right now."

- **As venture deals get increasingly digitized through syndicate platforms like AngelList, secondary marketplaces like Zanbato and Augment Markets and sourcing data and workflow products like Harmonic, the center of gravity in venture will shift from sourcing & access to high velocity evaluation, value add and generating DPI.** "Historically, edges in venture have been difficult to find. It's like 'who do you know', 'what kind of deals are you sourcing?' 'What do you have access to?' But I think in another five years, it's going to be mostly company selection and value add. VC funds know this, that as the access curve flattens—VC will look more like PE. They're trying to select the best deals. They're trying to do the ones that they can add the most unfair advantages to. If you can help them identify that really quickly, that's actually the biggest edge you can offer them today."

Interview

Can you walk us through your inspiration for starting the company and a short description of what Velvet is?

So the mission of Velvet was always: how do we bring private markets more online? We wanted to actually connect the private market at scale. We originally tried to solve that with an LP-GP marketplace.

What we did basically was a digital private-placement agent model, where managers would come to us, we'd do all the homework and diligence and build some initial trust and gather data, then turn that into an Airbnb listing, and market that to LPs on a marketplace. We weeded out most of the managers and only admitted the top tier. LPs got everything for free, because in the dating app metaphor, they're the ones everybody wants to match with. So we rolled out the red carpet, did a lot of work for them.. They could use us to find the best top-decile managers.

In late November 2022, we were raising our pre-seed round, led by Outlander VC, which by the way was the first check in AngelList. We got a chance cmeeting with Sam Altman, like two weeks before ChatGPT came out. This is at the same time as the banking crisis and all of that.

We complained about our cost of diligence. Our OPEX on underwriting these funds was around like $30,000 a month. Sam's like, "Well, we can read and write the way humans do now, it might save you a little bit of time and money, it might help with that." So we got turned on to GPT right at the turning point of mainstream AI.

We went down that rabbit hole. Can you use AI to ask data rooms a question? Can you read all of the unstructured information? Can you actually do thoughtful reasoning and analysis on top of that? Could you write an investment memo?

This was a brand new thing two years ago and it seemed applicable to work that private investors do every day. For the first time ever, I saw that you could pull all of the unstructured data and information in private markets and instantly digitize it.

I'd often make the comparison in my head to what Nasdaq and the exchanges did with the stock certificate in the 70s. Before that, your grandpa's shares were on a piece of IBM paper, and brokers would unload it for 10% to other brokers. It was all phone calls, cocktail hours, emails.

We'd already seen early digitization trends in private markets with syndicates and SPVs. The SAFE note got standardized less than 20 years ago. The cap table got digitized, finally, through Carta. These are all recent trends, but you do see this movement of private markets coming online.

The second thing we saw is that with AI we can infer across mass data sets better than any human can. Now you can derive insight and enable decision-making at scale way faster in a market that's so cramped by a lack of information, lack of information standards, opacity, and fragmented connectivity. Because of all these challenges, you still have the club-like behavior and a lot of reliance on trust and signal in private markets. These are the things that make venture a beautiful art form, but also what calls out for digitization.

At that pivot point, we moved all our resources to building AI for private markets. We built the first PDF readers and document-structuring tools before GPT could even do that. One of the first questions we asked for our managers was to see who was exposed to Silicon Valley Bank during the crisis. We did inference across all of our legal documents from these VC firms. We found out a good 30% of our managers were banking on SVB. We're like, "Oh, that's a good thing to let our LPs know."

By last summer, we came to market with a full AI co-pilot or digital analyst for private investors.

How does the AI-powered analyst work?

Broadly it's used by GPs and LPs to automate the deck-to-decision process. It follows them through the journey of the decision, structures data all the way, can read all their emails, pitch decks, entire data rooms, structure financial information from spreadsheets, pull out key terms from legal documents, and it creates this living knowledge base.

Our customers want their AI analyst to know everything about their firm. It collects information on the market, ingests their Crunchbase or PitchBook data, web-scrapes their deals, get the first-party deck and data room information from the individual companies. And then also the customer's networks—CRM data, email contacts, things like that.

Because the day to day of running a venture fund is actually very difficult. Their biggest cramp all the time is they have to make decisions at every moment as fast as possible. They have to evaluate companies as quickly as possible. You have to get down from 10,000 to 100 companies in a year. And it's a similar process on the LP side.

Where are your users seeing the biggest uplift? You mentioned investment memo drafting and reading the data room—where's the biggest time gain?

To narrow it down, there's been two really key pain points. Let's say it's a seven to 14 day process at best on evaluating a company.

First there's the organization of data, how to get everything in one place, which is to get 10,000 companies and all the information boiled down to something that's usable and actionable. The first thing you're doing is just chasing down all the missing data that's scattered across emails and decks and call transcripts. You'll see GPs and LPs still do a lot of manual data entry in spreadsheets—maybe Notion or Airtable if they're a little more evolved.

We built a tool that extracts all of the key details they want to define across the relevant documents. It'll also fill out your CRM for you. Generally CRM hygiene is weak at many firms, and analysts are spending hours every day just keeping things current.

So we save probably 80% on basic information organization, the user validates the rest. We have validation tools to do that.

But the highest ROI lift for us has actually been in the actual evaluation of companies, because it is a new lift that didn't really exist pre-AI. You didn't really have tools that could help you automate analysis, like to gain a deeper understanding of a company, or write an entire investment memo just from source material and market data.

That's really just a sunk cost. It's just I got to know and organize things.

Then the second big lift is really this deep dive. If you think about it, a due diligence process is maybe 100 hours deep, so

we can save you days of time in actual evaluation and decision support.

The naive view would be that private equity has more gravity around diligence because of lower deal frequency and operational transformation focus. But you're telling us one of the biggest uplifts is around evaluation & diligence for VCs. How does diligence fit in such a high frequency deal-making context?

Well, it's funny, that might be the initial read. In PE, yes, it's more data driven. You'll see that most of the players in AI have gone after that opportunity—investment banking, private equity.

But in venture they're also really trying to understand, as quickly as possible, is this a good company? Maybe they're answering that question at the most basic level on a first screen, but also later at the deepest level. And at that stage the question also becomes: "Can we add value to it?"

They want to understand: do we have 50 relevant customers or other investors in our network that we can introduce the company to? Can we do a competitive analysis and understand the competitive landscape in real time? This all matters a lot at the early stage. Most of our investors on the platform right now are Seed to Series B. They're all trying to find an edge right now and be as informed as possible.

Historically, edges in venture have been difficult to find. It's like 'who do you know', 'what kind of deals are you sourcing?' 'What do you have access to?' But I think in another five years, it's going to be mostly company selection and value add. VC funds know this, that as the access curve flattens—VC will look more like PE. They're trying to select the best deals. They're trying to do the ones that they can add the most unfair advantages to.

If you can help them identify that really quickly, that's actually the biggest edge you can offer them today. That's what I've learned from selling to VCs. Historically, VCs are a pretty difficult customer, because they're all different. But decision-making is a cramp they deal with every day. They're trying to do deals as quickly as possible. But also they're worried about missing the good ones. And they're worried about learning from their mistakes. An AI agent that can learn from your data

and grow with you through the decision-making process—there's a lot of outside value there.

At least 80% of venture funds have tried to figure out how to use AI in their workflows already over the last two years. But most of them just have hacked together something with GPT models.

You mentioned most of your customers are Seed to Series B. Is that the early adopter or wedge customer? What does the early adopter look like on the LP side?

Well, to add an element to that, our initial marketplace that we had for GPs was mainly managers under a billion AUM. So it was those more emerging managers that we had distribution to, which is what I think shaped our product in the early days.

The ideal early adopter has been a lean, tech-forward investment team. Those can be big and small, but you see a lot of solo capitalists, smaller VCs, and mid-size VCs that don't have in-house research staff. And on the LP side we see family offices that don't have major infrastructure, smaller fund-of-funds as well. These are smaller teams who want to be very data-driven and be able to drive value to their portfolio and also monitor it better. These firms are resource-constrained and want to scale their diligence and their data-aggregation and decision-making without adding more headcount.

This is really more than just a little memo generator. It's basically a full OS for them. It's like an Airtable that fills out itself. It's their Google Drive. We heard a dozen times from customers, just in their VC stack, all of our stuff is in different places. They use a bunch of fragmented tools. They just wish everything was in one place.

We've sold into a few larger GPs and LPs, especially over the last six months. A lot of that's on that decision-making process and being able to add value to companies. There's also this whole liquidity value proposition that I can get into in a second.

You'll also see family offices, LPs, and even growth-equity shops and PE that basically will layer our AI diligence product on top of their existing pipeline tools. We just integrated with Affinity. So you can pull all your Affinity data into Velvet and then run it through a deeper process, a decision-making process.

What are some of the more popular specific workflows your VC customers have moved to Velvet?

For example, in our memo-generation tool, you can select the sources you want it to run across. There's all the founder backgrounds that you get from call transcripts, LinkedIn data, call notes. We pull all of that in. For the companies, it'll pull in first the business description, but then also the market trends, and the TAM that's claimed in the pitch deck, and then it'll go cross-validate on the internet.

So it'll go research articles for you and say, "Hey, this company's TAM adds up. Here's some outside research on it." Same with funding history. You can pull from a Pitchbook and Crunchbase and competitive analysis you can do by doing a similarity search across all that data. So it's mainly those pitch decks and data rooms you're pulling from to generate memos.

But we also see users bring in cap tables for doing cap table audits. There's also board minutes and investor updates. So a few of our users have built templates. We have a custom agent builder. One of them built a board-reporting workflow. So they click a button, it reads all the board decks and emails that have come in from this given company and it'll summarize the latest KPIs and everything they're working on.

There's also this whole element of network data. We have a co-investment focus tool that'll read a company's data room and it'll tell you everyone in your network who's invested in companies like this, who's interested, who's a likely co-investor. And we're shaping that for introducing customers and acquirers as well, so that it can enable the VCs' strategic value add.

The dream scenario is you run a process all the way to the decision on the platform. You have your investment memo on Velvet, you run your IC meeting on Velvet. And at the end, you already have 50 customers and investors to introduce to a company so that day one, you're giving a company an unfair advantage and you already know exactly why you invested.

You talked about being the youngest team to ever acquire a private-markets broker dealer. Tell us about how that part of the business model came together. Why does it

make sense to combine an SaaS-style AI product with a transactional business?

We built it from scratch, actually. We didn't buy it. We built it from scratch.

I think part of it was organic because it's overhang from really a prior business model of raising money for GPs. Late 2022 was when I got my first secondary transaction done, and we started to realize all the LPs in our network were really interested in direct opportunities, primaries, secondaries, and SPVs. People kept asking us for SpaceX shares and Neuralink shares and growth stage AI companies like Anthropic and OpenAI.

And there's this massive buzz and rush into that. There's still real excitement around the secondaries market today. Over the last 14 months, we've executed about 220 million dollars in private placements into these. We've done that sometimes as a broker directly taking people to companies, getting them on the cap table. Other times we've done our own SPVs that we've led. We've done GP-led secondaries, as well as primaries that we've collaborated on, and same on the LP side. So we've built really a sort of bespoke boutique investment bank in the secondary space, while we were growing and taking this AI product to market.

And the interactivity between them is at least to me pretty obvious. Because actually, the biggest need in private markets is always liquidity. The recycling of capital, the being able to book returns for your LPs.

If you zoom out—70% of DPI from venture, I think, over the last year has come from the secondaries market, not from acquisitions or IPOs. So we're sitting on really strong tailwinds with a platform that's already capturing data across the investment lifecycle.

That's why I think the AI product is so interesting as it focuses on decision making. Because it's all about making better decisions faster, getting into better deals, getting returns faster. Why couldn't I use that same AI analyst to help investors find liquidity? Either by doing syndicates more quickly, like what AngelList did really well. Or maybe through getting secondary liquidity across a portfolio. That's how I see those two interacting.

One's a SaaS product. The AI analyst is a SaaS product with annual contracts that range depending on the size of the customer and the volume and queries they're making. And then we have a transactional business that is generally taking 2% to 5% placement fees or SPV fees on raising capital for specific opportunities or placing capital.

Right, got it. You touched on liquidity there, let's dig into that. Why would Velvet be well positioned to step in as a liquidity solution?

Yeah, I mean, private markets are generally illiquid. You have a lot of pressure in the venture space, especially over the last two or three years, because the IPO window has been so tight, and it's so difficult to deliver returns to LPs, so that managers can then go on to raise more funds. So these Seed through B investors have a bunch of Series C and D companies that they're actually pretty interested in booking returns on right now, just from a market-pressure standpoint.

But imagine in five to 10 years, you have a private market that's more liquid, more transparent. There used to be mutual fund brokers 40 years ago, now you can buy a mutual fund in a couple clicks on your broker. You can buy 10 million of Apple shares instantly, really without thinking twice. We're so far away from that in the private market, but it's mainly because the right regulation and information standards haven't been put in place and nobody's centralized the marketplace of private capital yet.

I think the best way to do that is through a daily-use tool. If you're a broker or a trader, you're on Bloomberg and an exchange every day. And that's because all the research is in one place, all the data is in one place, all the market participants are in one place. We're nowhere close to that today in the private markets, but that's the really long-term opportunity. Is there an exchange that can exist in the private markets? It's probably somewhere between the Bloomberg Terminal and NASDAQ or the New York Stock Exchange. A lot of people have tried in this space, including Carta, and Carta CEO Henry Ward has spoken out a lot on the opportunity and the challenges.

His position, correct me if I'm oversimplifying it, is basically that with respect to secondary market platforms,

the problem is that the startups themselves don't want them.
The companies don't want them, right, because it changes the status quo with regard to price discovery. He was referring more to the price discovery issue.

To be fair, I think we're still far from price discovery. People are trying to solve the liquidity discovery problem first. That's why there's Forge. That's why there's Zanbato. That's why there's Hiive, and they're killing it right now. But it's run entirely by brokers. It's not run by the first-party investors. I think that first step, the discovery of where liquidity exists, matching buyers and sellers, *that* still hasn't been solved. Price discovery is the next step after that.

Do you worry at all about what happened with Carta's Carta X exchange, where conflicts of interest arose between the back-office business and the secondary marketplace? How do you avoid those conflicts if you're mediating transactions while also serving SaaS customers?

Well, to be frank, I am worried about it to some degree. We built the whole platform with a firewall between those two businesses. So the brokerage business is using—I mean, that part we run like a venture shop, basically. We use our tool to evaluate companies and underwrite them and produce memos and share those with our LP base. But we're just one user of the product.

We don't have any access to any of Velvet users' data. It's all stored on a secure enterprise workspace. So every fund that works with the AI analyst has a sandboxed environment. It's enterprise-grade encrypted, there's zero training on all of the customer data. That's in our policy. Nothing's being trained on a central model either. All the indexing and AI inference occur within each firm's private environment.

And that's key. Because if you're going to build an agent that can work with people, it has to be *their* agent. That's what they want. They want this personalized thing that understands everything about them, that can add value to them, and that can include liquidity.

We see the brokerage and the transactional business as something that can interact with the user's needs but in a user-led way. I think why Carta ran into that problem—and I'd like to think I understand it pretty well—is that nobody ever lived on Carta. I think they did the marketplace idea too late. It was tacked on top of a system of record for the middle-office and the back-office. They had all the ownership data, they definitely had the best chance at it. But they didn't have workflow tools that were really embedded leading into the transaction.

So it didn't really make a ton of sense for users to lead the demand and supply. And so you ended up with employees at Carta basically incentivized to go drive volume through CartaX. And of course, they're going to maybe use data if they had access to it to try and drive volume. That was their big mistake, and I think that was very publicly acknowledged.

I'm never going to do that. It doesn't work. You have to align the incentives from the ground up. So at best, I could have an agent that suggests when people should be interested in a specific category, either as a seller or buyer. And then from there it all has to be double opt-in by the users. That's already what over emails right now today. But it's all done privately.

Regarding horizontal AI assistants, how do you think about maintaining an edge? Is being a vertical AI tool focused on private markets enough to shield you from generalist AI?

I'm all-in on the vertical AI agent for venture capital. The horizontal agent things I think are only competitive in the buy-and-build scenario, which most venture capital firms aren't doing. Most, 80% to 90% of them would rather buy. And then the 10%, like the SignalFires, or the Insight Partners, the really big funds, are more likely to build their own stack because they have the resources to do that. So they might get a Glean or a Hebbia to work with enterprise knowledge search over their own dataset integrated into their own CRM that they probably built themselves. But the majority of the market's really not in that position.

There's not really any VC-specific vertical AI tools that are hard in the market right now, I think, other than us. You'll see there's platforms like Wokelo, which will write a really good memo. There's a little bit of human in the loop with it. They get the

images, they get really good market research for you. And it's mainly used by PE and investment banking. They have a few venture fund customers as well, the way I understand it. And then you've got other tools that are VC-specific, like Harmonic, that own sourcing and people signals really well. But I don't even see Harmonic as competitive to myself either. I'd actually like to integrate with them because they're basically all about pipelines. They may be a better signal tool than the historic Crunchbase and PitchBook data, which at this point have just become data warehouses. Harmonic has become a robust alternative sourcing tool for many GPs who want to leverage where they can find signal across thousands of companies.

Do I think Notion and Airtable are my competitors? The generalist AI stack? Not really. I think venture's begging for a tool that can basically work across their whole investment and data stack, grow with them, this sort of new idea of an agent that can do that, focused around decision-making, which is the wedge we've been working on for a while. That's been working really well for us.

What about AngelList? I saw in some material you shared that you mapped them as maybe your closest comp in terms of positioning. Is their AI chatbot closer overlap than the other adjacent companies we've talked about?

Yeah, 100%. I think AngelList is one of the only really product-led companies in the space. And they've always been that way. They're not a service wrapper. They're not just an SPV formation shop or just an admin. They are a tech stack.

AngelList historically, though, is very stuck in the retail vertical. They did one of the greatest things ever, which was they digitized the syndicate, basically. But what that attracted was mainly these rolling funds, the smaller GPs, these angel investors. It started as a list of angels. I know the AngelList story really well. And ever since, they've been trying to push upmarket, and they've never been able to be a very serious institutional-focused product.

So I start quite a bit upmarket from AngelList in that most of our funds are serious funds with at least $20 to $50 million AUM all the way up to $2 billion AUM. And to their FinAI tool, it's interesting, but like a lot of products in the space right now,

it's pretty easy to replicate. It's basically GPT with Crunchbase and AngelList data. We have a very similar thing that we built a long time ago. Even their Relay product—we built an email forwarding tool in about a week.

In the end, I think an enterprise AI agent for venture, like ours, is going to be the thing that wins here. It's going to push me up-market quite a bit. In the end I would want to remain market-neutral to AngelList. If people want to syndicate on AngelList's network, I would build an integration the same way I'm building with Sydecar. In Q3 or Q4, Sydecar SPVs will be able to be created on Velvet platform.

I would integrate with all syndication, SPV platforms, all admins, same with Carta. We have a partnership moving forward there just to refer people to their SPVs. They see me as deal distribution. I would imagine AngelList sees me the same way. Whether people choose to live on the AngelList stack or the Velvet stack, I think will come down to the tech and the networks that live on it. Would you rather be a part of the larger family office venture capital fund network on Velvet or more of the rolling fund retail syndicate on AngelList? I think both have a place.

Let's talk about the future. Where do you see VC liquidity going? We hear a lot about secondaries booming right now, and maybe even the IPO window opening. What's the steady state for VC markets that you think will emerge, and how does Velvet fit in?

Yeah, I think in two to five years, you're going to see a more liquid Series B to IPO market—much more liquid, much more transparent, much more accessible. And that is actually being driven by institutions.

Probably the most famous acquisition in our space was BlackRock's acquisition of Preqin. I'm sure that BlackRock would love to build a Series C through IPO index of private companies because the alpha is there, the return profile is there. And they have the distribution and the assets. They have the balance sheet, basically, to drive capital through the private market. What's held all these firms up is the lack of standards across information and the lack of access. Because really, the only way to get access to this is through the customer base—the venture capital firms, the family offices, the private market investors. Facilitating first-party access

requires first-party data because you need to know where everybody's at and what they're doing.

By nature, my platform's doing that in the most real-time way possible. It's a diligence tool. So it's a tool that people are using every day to structure the data around their investments and analyze it, and then maybe offer it. So I see Velvet actually evolving as an access point in two to five years for this part of the private market—a data and an access point for institutions. This kind of living private market network, you could call it a more open capital network that could exist and that everybody wants. What's missing is access and standards, like I said before.

What would the institutional side come and do on Velvet? Can you make that concrete?

Well, I think with enough funds on the platform—and we've already processed 12,000 deals—enough deals and data, you could build an entire private markets portfolio through the platform at some point. So you could partner with managers. We could even reintroduce this sort of GP marketplace. You also see most LPs demanding GPs to do more co-investments, and things like that. So that's why we're leaning into this whole co-investment framework.

Another famous acquisition in our space is JP Morgan's acquisition of Aumni. They got bought for $400 million with not a lot of revenue. A good amount, but not a lot. Because JP Morgan wanted that first-party data. Aumni had all the legal documents across 300 venture firms—Tiger, Sequoia, a16z—because they took all those legal documents and processed them with AI, which was actually like 300 people in the Philippines.

So two years ago, we said, fuck it, we'll build Aumni, but with real AI. And we'll do it with more than just legal documents. We'll do it with all deal data and build a tool to analyze those deals. And then you have this sort of living knowledge base with all the deals and the investor behavior that goes into it. So if you're capturing deal data and then the workflow, the investor behavior and decision points, you're now starting to see this behavioral map of how the private market functions.

On top of that, you can build tools to help connect it better in a double opt-in way that gives institutions an access point for the

category, gives retail an access point for the category, gives that middle market—the family offices, the RIAs—all of those actors access points. Already in the secondaries business, we work with a lot of RIAs and family offices already on these high-signal names that nobody can get ahold of—the SpaceXs, the Groqs, the Lambda Labs.

That's how I see this transactional business and the SaaS product merging. The transactional business can lead a lot of that go-to-market and that signal that begets signal. For example, with AngelList, one of their big breakout moments was the Uber Series A was on AngelList. So you had this high-quality, high-signal network begetting high-quality, high-signal deals. So instantly, if you can get that kind of quality on a platform, it sort of flips the table and everybody's trying to get in. What's going on on Velvet? And I think that'll echo from retail all the way to institutions if you do it right.

I wanted to touch really quickly too back on that IPO window thing. I think I'm interestingly exposed to both the IPO window staying shut down *or* opening back up. So the SaaS product is volume-based deal stuff. So suddenly, if capital gets frenzied like in 2021 again, people are going to need the AI tool even more because you've got to evaluate companies faster. There's more stuff going on, the market speeds up, deal velocity speeds up. You've got to solve that information asymmetry problem as quickly as possible as a VC firm.

If it doesn't—we have the liquidity platform, which is like, okay, if the IPO window is closed, companies are going to stay private longer. Okay, great. I would be a market-neutral secondary exchange. I would plug into the Forges, I would plug into the Zanbatos and the Hiives, the Augment Markets. And I would do that because all I'm really trying to own is data and decision-making. So see the deals and control the access point through data and decisions. Everything else can plug into and create a market-neutral open capital network.

What does the future look like when agents are handling matching, diligence, maybe even closing? Is it basically a few decision-makers and then agents doing all the work?

What I've seen over the last two years is there's two things that people don't really want to renounce, and that's end decision-making. We don't want to replace human decision-making. I don't think we should either, at least with the state of AI today.

And then we also don't want to replace relationships or the human element.

So what all investors want to spend all their time on, is finding the best companies, working with those founders, adding value to them, bringing their network to bear, which I think is a moat that's not replaceable post-AI. And then making decisions—the selection part. So it's about multiple agents that follow you and support you through that journey. So, for example, it can highlight the highest ARR company in a given firm's pipeline. Just automatically it can say, hey, "here's this thing that fits your thesis the best with the highest metrics. You should take a deeper look."

And then you've got this memo agent, research agent, that carries out all of the parts of the deeper look for you. And you can work back-and-forth with it through that conviction journey. Kind of like Cursor AI will rewrite code for you. You can do the same stuff rewriting each section of your memo expressing your belief in a company—the market, the competition, the risks, et cetera. So you can be this incredibly informed investor, even in just having great questions before the first call. Then you've got the first call agent as well.

And then on to the liquidity agent, that's where things get really interesting. So Sequoia has got a piece where they say that they think there's going to be an agent economy sometime in the next five to 10 years. And it is agents working with other agents. If I did that successfully for every venture capital firm, my idea of a liquidity-focused agent is one that could actually suggest both demand and supply across a proprietary dataset. So it becomes this sourcing tool where it's like, "13 investors on the Velvet platform are looking at fintech deals. Would you like to be friends with them?" There's a Bloomberg-style chat room. There's digital deal rooms that can be shared. It's all double opt in. And all the data is already there because everybody's working with their own agent everyday.

And then maybe these agents are suggesting connections to each other. I'm seeing a lot that can happen, but importantly you can maintain privacy in the double opt-in that while actually enabling liquidity. Whether that's like, I want to bring 20 investors to my syndicate or I want to fill a secondary very quickly and discreetly. The agents all know their investors' behavior, what they're looking at, their thesis, what they care about, because everyone wants a more personalized agent

that can carry out tasks the way their firm does, respecting their format and their standards.

So that's how I see it—that's the positive-sum game that you could drive here. But to do that you have to have all the data and behavior in one place, because that's how agents will understand you and grow with you. There's no tool that's ever been built that LPs and GPs use every day.

Do you need LPs to close the loop in this market so to speak?

I think LPs are a big part of the piece. That's the way I think about it. Because they're at the top of the capital stack.

Their current need is not a lot different than the GPs, but they basically just want to know what's going on across their managers and on to the portfolio companies. They don't know what's going on, especially the venture capital, private equity, and there's not really standards of reporting there. So with this tool, you're giving them better data, better information across all of their investments, and in turn, better access. So it goes back to what brings institutions in—information standards and participants being all in one place. That really includes LPs because GPs are a service provider for them at the end of the day.

GPs also have a willingness to pay constraint, when you look at them as a long-term, high-value customer. Can I really charge all of these funds 100K ACV with the best product ever? Maybe. But LPs are a better customer from that perspective, after all they're the capital provider.

Whether its GPs or LPs, the key is anchoring around the atomic unit of the decision, accelerating decision-making, which should accelerate capital flow, if all of the data is there. So focusing on decision-making, you get peak information on the decisions made around a company, and the information flow from founder to LP should be able to accelerate around that premise as well. That's how I look at the whole market. Can I centralize data and capital flow and increase the velocity of the latter? That's the big end goal.

Disclaimers