
THE NEXT
AI UNICORN
BROWNRIDGE
RESEARCH

DAY ONE
INVESTOR

The Next AI Unicorn

BY JEFF BROWN
EDITOR, *DAY ONE INVESTOR*

Dear *Day One Investor* Reader,

Welcome to the newly relaunched *Day One Investor*. As the name suggests, here we invest in private, early-stage companies with 10X, 100X, and yes, even 1,000X potential.

We will do this by investing in the most promising Regulation Crowdfunding (Reg CF) and Regulation A/A+ (Reg A/A+) deals, which are open to all investors regardless of accreditation status.

My name is Jeff Brown, and I will be your editor.

For readers who don't know, I spent nearly three decades as a high-tech executive at blue-chip tech giants like Qualcomm, NXP Semiconductors, and Juniper Networks.

I've earned degrees from Purdue University and the London Business School. I've also received professional certificates from MIT, Stanford, and the University of California, Berkeley, School of Law. And I'm an alumnus of Yale University's School of Management.

I'm an active angel investor in early-stage technology companies. I've invested in more than 400 private deals. And that number continues to grow. This is where I've earned my largest returns as an investor, often in the thousands, sometimes tens of thousands of percent.

I understand the potential returns of this space firsthand. Investing in "day one" private companies has changed my life and ensured my family's financial security for generations to come.

In the months and years ahead, I will share with you my highest-conviction recommendations in the Reg CF and Reg A/A+ space. And I look forward to sharing my research on some of the highest-growth investments available to all investors.

To get started, I encourage subscribers to read my *Day One Manifesto*. This special report will provide some background on the real potential of private investing. The *Manifesto* will also outline our investment philosophy and show what types of projects we will be targeting.

I couldn't be happier to relaunch this service for my subscribers and to continue the mission I'm so passionate about. It's great to be back.

Jeff

SYNTENSOR

Crowdfunding Offering Type: Regulation CF
Security Type: SAFE (Simple Agreement for Future Equity)
Amount Available: $5 million
Minimum Investment Amount: $100
Offering Pre-Money Valuation Cap: $20 million
Availability: Open to U.S. Domestic and International
Industry: Generative AI/Biotechnology
Offering Platform: Republic (https://republic.com)
Link to Deal: https://republic.com/syntensor/

In the early days of the semiconductor industry, the tools available to design integrated circuits were limited. It was a labor-intensive, manual, and highly inefficient process that created a bottleneck in innovation and created uncertainties in how a semiconductor would perform once manufactured.

A great example is the Motorola 6800, an 8-bit processor designed 50 years ago. The chip, shown below, had 4,100 transistors, which was the leading edge at the time.

Motorola 6800 | Source: Motorola

Just to put things in context, NVIDIA's most advanced GPU, the Blackwell (B100), is designed with 208 billion transistors.

But back then, it would have been impossible to build a semiconductor like NVIDIA's. The tools simply didn't exist.

To design the Motorola 6800, the design team printed out the photomask – the blueprint for the semiconductor – onto paper to be laid out on the floor for analysis and testing. It took up an area of 15 meters by 30 meters, about the size of a basketball court.

And here is the crazier part, Motorola assembled a large team of employees to stand on the paper blueprint. Then the design lead would blow a whistle to simulate the clock cycle of the semiconductor. (The clock cycle is like a metronome for an integrated circuit. With every cycle comes a sequence of actions in the circuit).

Individuals would raise or lower a flag every time the whistle was blown. And those who were standing on specific functions in the circuit would step to the next gate in the circuit design.

This was simulation without software.

This may seem elaborate or ridiculous, but it was necessary to simulate how the integrated circuit would work and to debug any problems in the circuit design. It was an imperfect process that was prone to errors.

The problem was that it took millions, often tens of millions, of dollars to design a semiconductor and send it to manufacturing. And it was impossible to know with certainty whether or not the semiconductor would perform as intended until it was produced.

If the design didn't work, it was back to the drawing board.

THE ADVENT OF EDA

It wasn't until the 1980s when things radically changed.

That's when electronic design automation (EDA) software was created. It simplified the process of designing semiconductors and also analyzed and tested the circuit design using software to dramatically increase the likelihood that the semiconductor would function as intended once it went to manufacturing.

The company that led this transformation in semiconductor design was Cadence Design Systems (CDNS), itself the result of a merger between two leading EDA companies at the time. Cadence's EDA software changed the entire industry.

Rather than each semiconductor company having to develop their own software design tools and processes, they could now rely on Cadence to provide them with the software to automate the process of semiconductor design, and to test the semiconductors before they went to manufacturing.

It was a game-changing technology for the industry. Cadence went public in 1987 with $40 million in annual revenue and about a $100 million valuation. This year, Cadence will generate more than $4.5 billion in revenue and is worth about $86 billion. It's one of the greatest success stories in the history of high-tech.

And Cadence is a perfect example of what we eventually see in all industries. Technology companies always spring up to develop software to help solve critical problems that exist throughout an industry.

Whenever we see the same research and development (R&D) or operational tasks being done in-house by companies – that are the same for other companies in an industry – we know there is an incredible business opportunity to develop technology to automate those tasks for all companies in the industry.

Doing so accelerates R&D, reduces risk, and increases the probability of successful product development.

SOFTWARE DESIGN IN BIOTECHNOLOGY

I see the same dynamics playing out right now in the biotechnology industry. The same shift that happened in the semiconductor industry decades ago is about to happen in biotech.

The catalyst for this shift is artificial intelligence (AI).

AI, as it pertained to biotechnology and life sciences, started to show promise back in 2018. The inflection point was when DeepMind – a subsidiary of Google – applied neural network technology (a form of AI) for the application of predicting protein structures.

Understanding the structure of proteins has long been a grand challenge in life sciences… with limited success in solving using physics-based modeling.

Proteins are the building blocks of life, and understanding their structure enables us to understand their functions, as well as how to optimize therapies to improve health and save lives.

DeepMind's first attempt with its protein prediction AI, AlphaFold, delivered the best performance compared to all other attempts to date. It was ironic considering Google is an internet advertising company, and it was able to best all other research teams in the life sciences industry in 2018.

But the performance wasn't accurate enough to be useful.

The incredible breakthrough came in the summer of 2022 when DeepMind released an updated version of its AlphaFold AI – AlphaFold 2. It came with a database – free for all to see – of protein structure predictions of more than 200 million known proteins.

AlphaFold 2 not only "predicted" more than 200 million known protein structures... it did so with a near-perfect prediction.

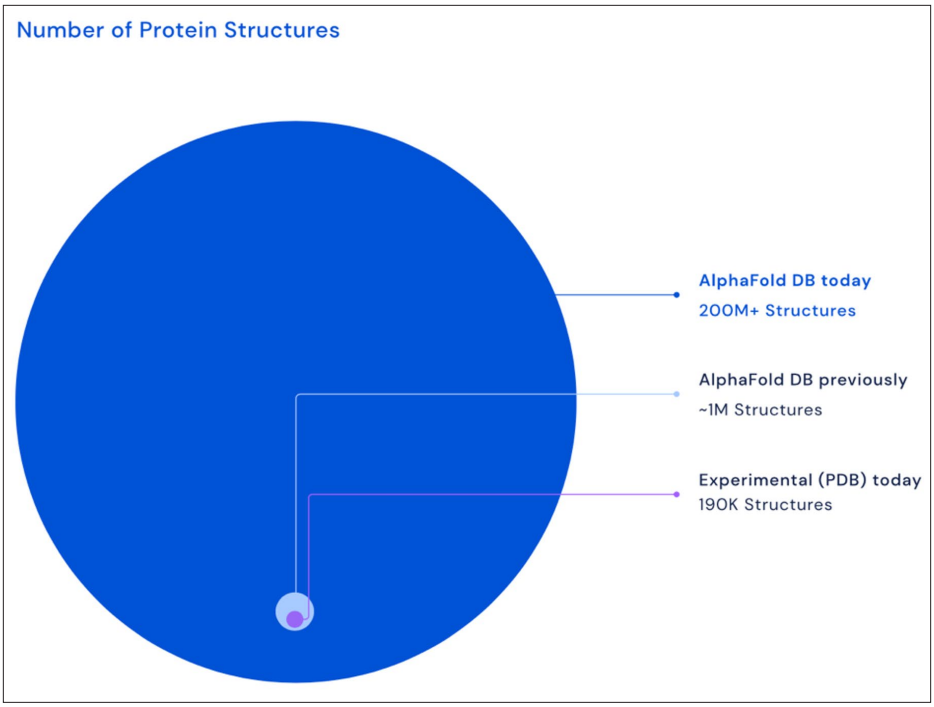

Source: DeepMind

To visualize what a remarkable achievement this was, the image above shows how greatly the knowledge base of proteins expanded... with just this one release of a single AI tool.

DeepMind went from about 1 million predicted protein structures in 2020, to more than 200 million protein structures in 2022.

And just a few months ago, DeepMind released AlphaFold 3.

AlphaFold 3 has predicted the structure – not just protein structures – of all molecules of life.

AlphaFold 3 is capable of accurately predicting the structures of proteins, DNA, RNA, and ligands, which are "binding" molecules that create bonds of various strengths with other molecules and ions. It even predicts how they interact.

This is likely the most valuable scientific tool and repository of life sciences information that the biotech industry have ever asked for. And it's free.

A COMPANY AT THE CENTER OF IT ALL

We're at a unique moment in time. The latest generation of semiconductors – specifically GPUs and other AI application-specific semiconductors – is the reason we've seen breakthroughs like DeepMind's AlphaFold.

This has led to a wide range of life sciences-related developments, resulting in large data sets related to human biology. And when we combine large data sets with powerful semiconductors and artificial intelligence, we're able to create powerful, generative AI.

Which leads me to our opportunity to invest in an incredible early-stage, generative AI company focused on the biopharmaceutical industry – Syntensor.

I first discovered Syntensor more than a year ago, and since then, I've had several meetings and calls with Syntensor's CEO, Clayton Rabideau.

Rabideau is a PhD candidate in chemical engineering and biotechnology at the University of Cambridge. He almost completed his doctorate work there, but the idea of building Syntensor was too compelling to wait.

While there, he co-authored a remarkable research paper that demonstrated a genomic foundation model trained on the human genome. An oversimplistic description of the work is that it is like a ChatGPT for the human genome.

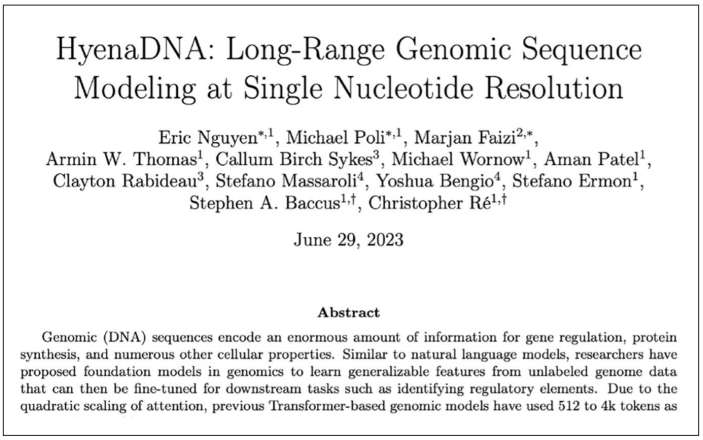

HyenaDNA: Long-Range Genomic Sequence Modeling at Single Nucleotide Resolution

Eric Nguyen[*,1], Michael Poli[*,1], Marjan Faizi[2,*], Armin W. Thomas[1], Callum Birch Sykes[3], Michael Wornow[1], Aman Patel[1], Clayton Rabideau[3], Stefano Massaroli[4], Yoshua Bengio[4], Stefano Ermon[1], Stephen A. Baccus[1,†], Christopher Ré[1,†]

June 29, 2023

Abstract

Genomic (DNA) sequences encode an enormous amount of information for gene regulation, protein synthesis, and numerous other cellular properties. Similar to natural language models, researchers have proposed foundation models in genomics to learn generalizable features from unlabeled genome data that can then be fine-tuned for downstream tasks such as identifying regulatory elements. Due to the quadratic scaling of attention, previous Transformer-based genomic models have used 512 to 4k tokens as

Some of us might notice one the of names on the paper shown above – Yoshua Bengio. He is one of the godfathers of AI and is known for his work on deep learning and neural networks.

Rabideau assembled a cross-functional team at Syntensor with backgrounds in biology, artificial intelligence, machine learning, bioinformatics, quantum chemistry, and other AI-related expertise.

I have been looking for a company with a mission like Syntensor for years. As I've been researching the cross-section between semiconductors, artificial intelligence, and biotechnology, I knew the time had come for a company like Syntensor to become the Cadence of the biotechnology industry.

Just like the semiconductor industry years ago, the biotech industry has tended to apply machine learning, a form of artificial intelligence, in-house.

What this means is that biotech companies hire software engineers to work with machine learning in hopes of accelerating drug development alongside a drug development team.

These two skill sets are quite different, and they are often at odds with each other. Early-stage biotechnology companies, even those that are well-funded, can rarely afford to build out a top-notch AI team empowered to build a powerful model for optimizing drug discovery and development.

Financial constraints force biotech companies with small, in-house AI teams to develop a narrow focus with their AI, which hands over drug development candidates to the drug development team for analysis and decisions to move forward.

Once the biotech company has a short list of drug candidates for development, it tends to de-emphasize the AI team and shift focus on working the drug through the clinical trials process.

The problem with this narrow focus of applying machine learning to a smaller subset of data is that it limits the benefits of using AI for development and predictive capabilities.

THE HARD REALITIES OF THE BIOPHARMA INDUSTRY

The industry reality is that there is about a 90% failure rate for drugs that enter clinical trials. And the cost to develop a new drug and receive FDA approval is around $2 billion.

Worse yet, the fourth-leading cause of death in the U.S. is adverse reactions to drugs. That's worse than the number of deaths from diabetes.

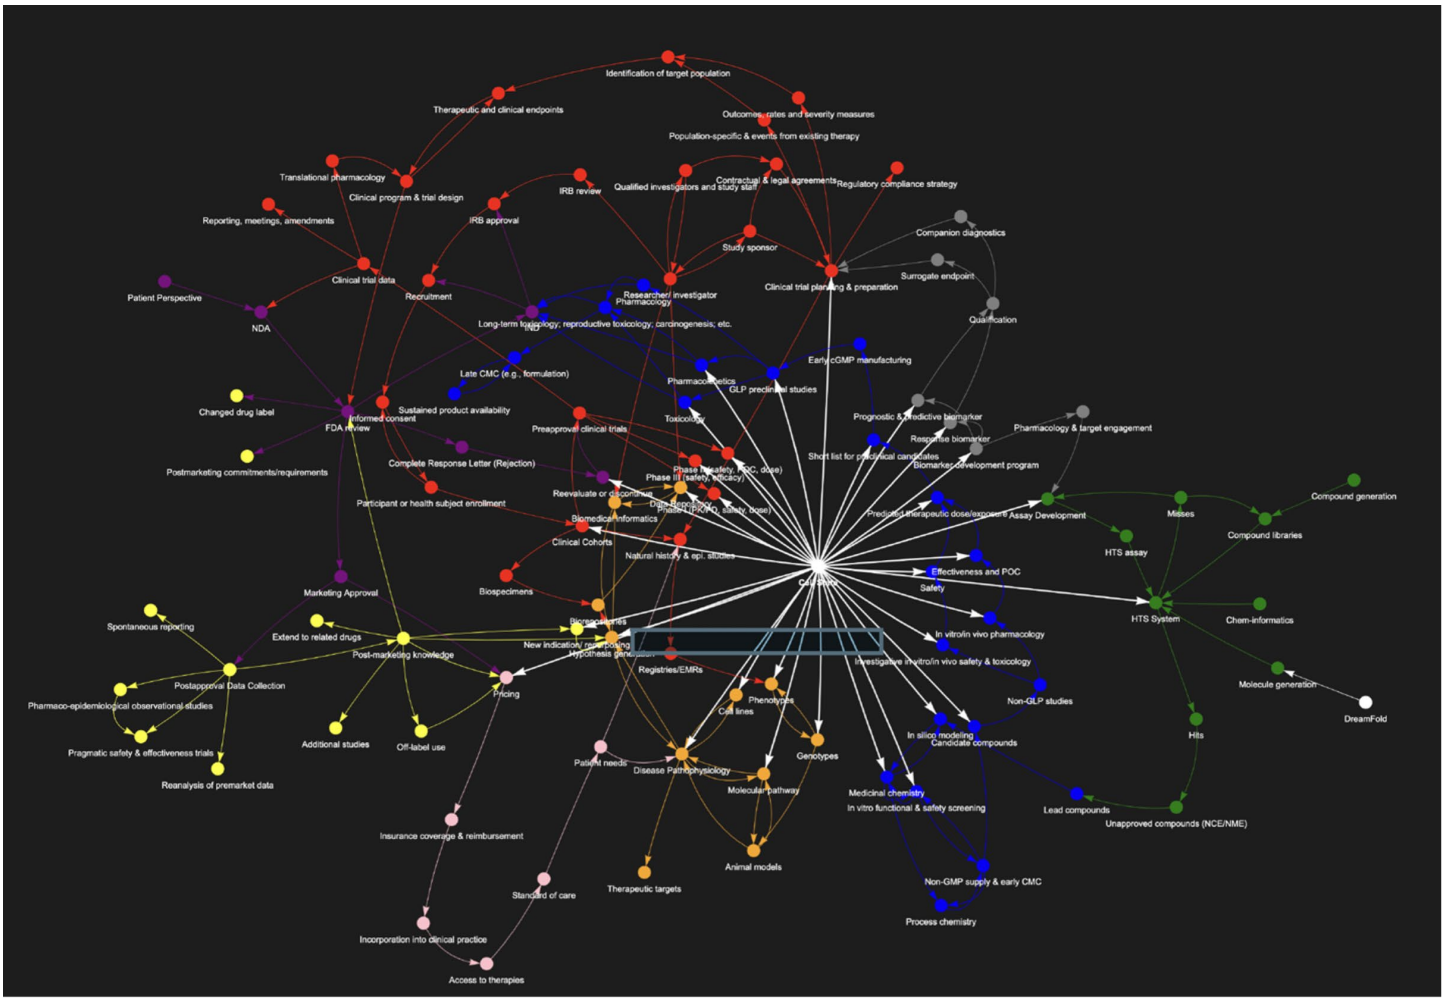

The reality is that the drug discovery and development process is incredibly complex, with so many interconnections and dynamics that are interlinked. It's precisely the kind of complex, multivariable problem that AI is good at solving… given enough data. (See image above.)

The problem with the narrow approach is that it does not take advantage of the full power of artificial intelligence. Without a large, general AI model for human biology and physiology, there will be an absence of predictability in drug development.

A FOUNDATION MODEL FOR BIOTECH

This is exactly the problem that Syntensor is solving. Just like Google's DeepMind was able to accurately predict how more than 200 million proteins fold by building a massive general model with artificial intelligence, Syntensor is doing the same for human biology, physiology, and drug development.

It's the same concept behind the incredible success that we've seen with large language models like Chat-GPT, Google's Gemini 1.5, Anthropic's Claude, and xAI's Grok 1.5 in the last 18 months. They all built what's called *foundational models* for language using more than 100 billion parameters for training their AIs.

Syntensor is building a foundational model for drug discovery and predictability for the biopharmaceutical industry. This is a multibillion-dollar opportunity and we're able to participate on the ground floor.

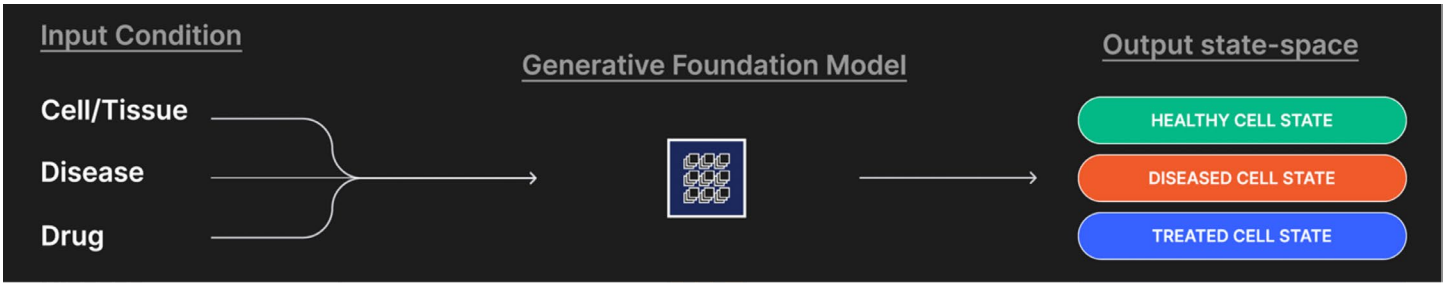

The simplest way to think about Syntensor's AI is that it can "ingest" data provided by a biopharma company that includes information about cells, tissues, the target disease, and the molecular structure of their drug.

This information is processed by Syntensor's generative AI – its foundational model – and the system outputs critical and predictive information about the state of the human cells given the inputs. (See image above.)

Syntensor's foundational AI model is so valuable because even small improvements in predictability can be worth tens of millions, if not hundreds of millions, of dollars.

As mentioned before, developing a drug from the drug discovery phase all the way through to FDA approval costs around $2 billion. Which means that it's not unusual for biopharma companies to get to Phase III clinical trials and find out that their drug has limited efficacy and/or unwanted adverse side effects after having already spent more than $1 billion on R&D and clinical trials.

Being able to predict early on that a drug candidate will likely not be effective can save unbelievable amounts of capital.

Better still, being able to accurately predict drugs that will be both effective and safe in advance can transform a company's therapeutic pipeline, accelerate drug development, and ensure that the biopharma company is spending its capital on the highest probability drug candidates.

THE CADENCE DESIGN SYSTEMS OF BIOTECH

Syntensor's business model is very simple and similar to Cadence. It licenses out the use of its AI and technology platform to its customers.

It is a classic picks-and-shovels play on the entire biotechnology and pharmaceutical industry. This is a key point, as we don't have to worry about the typical risks associated with drug development. Syntensor isn't a drug developer.

Its business will thrive as long as its technology improves drug development outcomes and helps biopharma companies avoid pursuing drug candidates that have low probabilities of success. And as long as there is a growing biotech industry, there will be incredible opportunities for growth for Syntensor.

And the best part is that we are on the cusp of what I call the Golden Age of Biotechnology. The vast availability of data related to genetics, proteins, disease, existing drugs, toxicity, and other factors – combined with the power of artificial intelligence – will drive radical advancements in drug development.

In fact, it has already started.

The costs for drug development will decline, the speed at which drugs can move through clinical trials will increase, and there will be a flood of investment into the biotech industry as a result.

Syntensor is going to be right in the middle of it all.

SYNTENSOR'S DEAL STRUCTURE

We're investing at a very attractive pre-money valuation cap of just $20 million, in the form of a simple agreement for future equity (SAFE). This is an investment that eventually converts to equity at a valuation of $20 million or less, upon change of control (acquisition) or an initial public offering (IPO).

[**NOTE:** A pre-money valuation "cap" simply means that in the very unlikely event that the next funding round happened at a valuation less than $20 million, it means that our investment would convert at that lower valuation. Put another way, it guarantees that our conversion, and thus our investment, will happen at a $20 million valuation or less.]

As a reminder, Cadence went public in 1987 at a valuation of around $100 million. Syntensor's offering is happening at a $20 million pre-money valuation, which is extremely rare these days for a company working on a foundation model in artificial intelligence.

As a comparable, Clarivate (CLVT) is an analytics and services company that operates in the life sciences and healthcare industry. It is publicly traded and now worth about $9.7 billion.

There are no direct comparables for Syntensor, but if the company simply grows into a $1 billion valuation, that would represent a 50x return for investors in this funding round. And if Syntensor can remain independent, the opportunity scales up from there. Syntensor has multi-billion potential.

This investment is open to all investors, accredited and non-accredited, and is open to both U.S. and international investors. In the U.S., ACH transfers are the easiest, but credit cards are also OK. And for international investors, both wire transfers and credit cards are fine for making the investment.

If any subscribers have issues reserving their investment, please reach out directly to the following email address: investors@republic.co. Alternatively, subscribers can speak with a Republic representative at +1-347-934-6876.

Also, please note that investors domiciled in the Canadian provinces of Alberta, Ontario, and Quebec will not be able to participate in this offering. That's because of restrictive regulations around private investments specific to those provinces. This is something that's completely outside the purview of Syntensor and Republic.

IMPORTANT: PRACTICE RATIONAL POSITION SIZING

Please remember that we should always keep our position sizes rational for these early-stage deals.

The absolute worst thing a private investor can do is go "too heavy" on one early-stage private investment. The smartest angel investors I know always balance their portfolios with rational investment sizes over a larger number of companies. This is the very best way to diversify risk in a highly risky asset class.

So let's make sure we are allocating our capital wisely to each position.

My guidance for Syntensor is for subscribers to invest no more than 20% of what they normally put into a small-cap tech stock.

For instance, if we normally invest $5,000 into a small-cap stock like the ones I recommend in *Exponential Tech Investor*, then we should invest no more than $1,000 in Syntensor.

If we'd normally invest $1,000 in a small-cap technology stock, our position size would be $200.

This speaks to something I've said before – every recommendation I'll make in *Day One Investor* is different. We are going to be strategic with our capital allocation plans.

That said, we should never invest any more than we are willing to lose in a single, early-stage private investment. No matter how much diligence we perform or how well I know the industry, company, and technology, there are simply too many unknown variables outside of our control when investing at such an early stage.

That's why our goal is to build a large portfolio of exciting, high-growth companies with fantastic founders and incredible growth potential over a wide range of industries.

These investments will take place over a number of years. And we want these companies to stay private, thrive, receive additional funding, and eventually grow into multibillion-dollar companies. All it takes is one unicorn ($1 billion valuation or more) to drive outsized returns in a portfolio.

Action to Take: Invest in Syntensor on Republic.

Please go to the following URL to reserve your place in Syntensor's Reg CF offering as soon as possible.

https://republic.com/syntensor/

This investment opportunity will fill up fast, likely within 24-48 hours of publishing my research. So please don't wait. Investing in private deals isn't the same as investing in a publicly traded company.

High-quality deals simply don't last long, and once the $5 million allocation is taken, there is nothing left.

These recommendations are the kind of companies that can easily raise capital from venture capital firms. In that way, we are "competing" with the VCs for deals. And because these companies have so much growth potential, the investments will fill up quickly.

Regards,

Jeff Brown
Editor, *Day One Investor*

RISK MANAGEMENT

Due to the nature of this investment, there will be no stop loss. I always encourage readers to use rational position sizing. That means we should only invest an amount that is appropriate given our portfolio size and tolerance for risk.

Each investor is responsible for determining what a proper position size is given their individual circumstances. And remember, I never recommend anybody go "all in" on any one investment.

Investing in private offerings is different from investing in public equities. Upon making an investment, we will not be able to exit our position.

Private investments can remain illiquid for extended periods of time. That means investors should only invest capital that they can afford to be without as our investment thesis plays out. On a similar note, investors should never invest any more than they are willing to lose on any single private investment.

If you have questions regarding this investment, please contact the team at Republic. The customer service representatives of Brownridge Research will be unable to answer your questions regarding this investment, as we do not have any business connection with Syntensor.

Finally, the role of Brownridge Research is that of a publisher of investment research. We are not associated with Republic in any way. And we never receive financial compensation for any recommendations we make. As an analyst, I provide general recommendations and guidance only for the benefit of my subscribers.