EONVERA CAPITAL

# EONVERA CAPITAL



## Audrey 1

Dear investor friends, hello everyone! ❤

I'm Audrey Lim, the person you all know as "The Golden Envoy."

Monday always fills us with hope. Usually, we'd be discussing market trends, company earnings reports, and investment opportunities.

Today, we're not going to start from an index or a stock. First, I'd like to think through a question with you: when technology is advancing faster and faster, capital investment keeps growing, and the future still can't be fully predicted, how do we make judgments?

Recently, AI infrastructure investment and corporate financing have once again become a market focus. Bond investors are starting to scrutinize tech companies' borrowing needs, cash flow, and investment returns.

This week, the U.S. midterm elections and the jobs data could also shape market expectations.

All of these events point to one thing: how do we

understand value and identify risk amid change, without being led around by short-term emotions?

Today, Dr. Campanelli, Founder & Chief Executive Officer of Eonvera Capital, will share his thinking on capital, intelligence, and decision-making, drawing on his education, his experience on Wall Street, and the building of EONVERA ADI.

I hope everyone will hear not only about technological iteration today, but also the thinking behind the technology.

Now, I'll hand the time over to our "Grandmaster of Capital," Dr. Campanelli.

## Pietro 2

Good morning, guys.

Before we get into today's content, I'd like to express some gratitude first.

Thank you to Audrey and the investor education team for the work you've done over the years.

Truly valuable conversations never happen by chance. They are organized over time, professionally and with restraint.

That I'm able to meet with you all in this way today is thanks to your efforts.

Thank you!

First, a simple question: if a technology has a bright future, does that mean every dollar invested in it today is worth it?

My answer is: no.

AI may change production, scientific research, and business. But building data centers, buying chips, and signing power contracts all take real capital. Capital has a price, debt has a maturity, and investment must come with returns.

Stock investors may focus on future growth; bond investors will ask: where will the interest come from? How will the principal be repaid? If customer demand falls short of expectations, will the company still have enough cash flow?

Bond market attention on AI financing doesn't mean the AI

industry has failed. It reminds us to separate the technology vision from financial reality.

Great technology is worth looking forward to, but every investment still has to pass the test of the cost of capital. Don't stop doing the math just because the AI story is grand.

## Pietro 3

Imagine a company deciding to build a large data center. It has to raise financing, buy chips, build facilities, sign power contracts, and finally wait for customers to rent the computing power. Every step comes with conditions.

When financing costs rise, project returns get squeezed; when construction is delayed, revenue is pushed back; when chips update too fast, equipment may depreciate ahead of schedule; when customer demand falls short of expectations, utilization drops.

A change in one link can spread to others. That is systemic risk.

So I don't just ask, "Is AI demand still there?" I also ask: Who pays for the demand? How long will the contracts last? What is the equipment utilization rate? Can cash flow cover the debt?

The bond market isn't announcing a verdict on behalf of the stock market. It's reminding us that we need to check the ability to repay behind a company's commitments.

Real risk awareness isn't about being afraid of change. It's about understanding how changes transmit to one another.

## Pietro 4

Guys,

To understand the AI industry, you have to tell two terms apart: training and inference.

Training is how a model learns parameters from large amounts of data and computation, and it usually takes intensive computing power and a long time. Inference is how a model, once it's put into use, provides answers,

predictions, or generated content to customers, and every call can generate ongoing costs.

Training gives a model its capability; inference brings that capability into real-world use. But capability isn't revenue, and call volume isn't profit.

Even a company with an advanced model has to answer: How much are customers willing to pay? How much does each inference cost? Will unit costs fall as the business scales? Can revenue cover computing power, energy, and operating expenses?

Whether a technology works is one thing; whether customers are willing to keep paying is another.

When researching an AI company, you can't just look at how big the model is. You also have to look at whether it can turn that capability into sustainable business value.

## Pietro 5

This week, there are two important things to watch in the U.S. market: the November 3 midterm elections, and the

October jobs report coming out on November 6.

The election could affect market expectations for fiscal, regulatory, trade, energy, and tech policy. But election results don't automatically turn into corporate earnings, and they won't immediately decide any industry's fate. Policy still has to go through institutions, legislation, and implementation before it gradually affects the economy.

Jobs data works the same way. Job gains, the unemployment rate, and wage growth shape the market's view of economic momentum and inflation, but a single month's data can't determine the direction of rates on its own.

I hope everyone will remember this: events are facts, expectations are interpretations, and prices are the result of market participants acting together. The three are not the same thing.

Don't let a headline think for you, and don't let the market's first-minute reaction make a long-term judgment for you.

## Pietro 6

Guys,

I trained in mathematics, probability, operations research, and economics at Princeton University. Those courses taught me how to handle complex problems, but what truly changed the way I think were the questions my professors asked.

Burton G. Malkiel was my professor. His research on market efficiency, random walks, and long-term investing had a deep influence on me.

He made me face a fact squarely: publicly available information is quickly reflected in prices. Investors can't assume they have an edge on confidence alone.

What I learned from him wasn't to give up on research, but not to treat prediction as the whole of investing ability.

John Y. Campbell's research in financial economics and asset pricing also helped me understand the links among risk premiums, macro variables, and asset prices.

Math isn't there to prove that you're always right. Its value

is in spelling out assumptions clearly, stating probabilities clearly, and showing us under what conditions a conclusion could fail.

Being rational isn't being sure you're right. It's carefully checking why you might be wrong.

## Pietro 7

After I got to Stanford, my research went deeper into stochastic optimization, dynamic decision-making, and computational mathematics.

William F. Sharpe was my doctoral advisor. His research on portfolios, asset pricing, and risk-adjusted returns profoundly shaped my understanding of capital allocation. Returns can't be evaluated in isolation from risk. That's a principle I've always held to.

Myron Scholes's research in mathematical finance showed me that complex financial relationships can be expressed rigorously; John Hennessy's work in computer architecture made me realize that algorithms ultimately have to run on

real processors, storage, and networks.

I've also seen how different fields change the world through the data and algorithm practice of fellow alumni Sergey Brin and Larry Page, Jennifer Rexford's research on network systems, Jeff Bezos's path of building platforms, and Jensen Huang's journey of advancing high-performance computing.

To me, my mentors and fellow alumni represent ideas worth studying, engineering ability, and entrepreneurial practice.

Ideas must be able to land, systems must be able to run, and capability must be tested against reality.

## Pietro 8

After earning my doctorate at Stanford, I experienced different investment research environments at D. E. Shaw, Bridgewater, and Two Sigma.

At D. E. Shaw, I saw how data, models, and systematic research enter a real investment process. An idea can't

just look beautiful on paper. It has to be tested by data and the market.

At Bridgewater, I came to understand macro variables, cross-asset relationships, and risk budgets more deeply. Even when a directional call is right, if the position is too large or liquidity is insufficient, it can still force an investor to exit at the worst possible time.

At Two Sigma, I got further exposure to data science and machine learning, and I also saw that data quality, out-of-sample performance, and overfitting have to be handled seriously.

Those experiences made me give up an illusion that's easy to have when you're young: finding a model that's always right.

Markets change, participants change, institutions change. Models must be tested too.

Investing ability isn't never making mistakes. It's spotting mistakes earlier, and not letting a single mistake destroy the entire system.

## Pietro 9

After many years of study and work, I began to wonder: if an institution's performance depends entirely on one person's judgment, what is left of the institution when that person makes a mistake, leaves their role, or the market environment changes completely?

In 2019, I founded Eonvera Capital, hoping to build not an organization that revolves around one person's opinions, but an institutional framework capable of continuous research, discussion, testing, and risk management.

An institution's long-term capability shouldn't be defined only by one year's results. Whether research can be reviewed and verified, whether risk has boundaries, whether responsibility is clear, and whether the team can learn from its mistakes matter just as much.

Different people can hold different views. Disagreement isn't the problem. A lack of evidence and accountability is.

I founded Eonvera not to have everyone follow one person's judgment, but to build a culture in which

judgments are put to the test.

Capital deserves respect, because behind it there is often the time, labor, and expectations a family has built up over many years.

Managing capital is, first and foremost, a responsibility.

## Pietro 10

In 2020, we began systematically building EONVERA ADI: Adaptive Decision Intelligence, an adaptive, evolving investment decision intelligence system.

Why do we need it? Because markets are too complex for a single model to answer every question.

One model may be good at identifying trends but unable to explain why things change; another may be good at valuation but overlook liquidity risk; and some models perform well on historical data but can fail in a new market environment.

That's why ADI isn't a tool that only hands out stock tickers. It has to connect data, model validation, multi-agent

deliberation, causal and scenario analysis, computing resources, and risk governance.

I place particular importance on counterevidence. A system that only looks for evidence supporting its own view will gradually turn into an echo chamber.

Truly valuable research has to actively seek out information that could overturn the original judgment.

ADI isn't meant to have machines always give answers. It's meant to make sure every answer can withstand questioning.

## Pietro 11

Now, I'd like to talk specifically about ADI's computing infrastructure, which is the foundation of ADI's most essential value.

We use NVIDIA GPUs and the CUDA ecosystem to support parallel computing, model training, inference, and simulation tasks that are compatible with it.

GPUs can process a large number of computations at the

same time, but a GPU won't decide for us what to research. It provides computing power; research questions, algorithm design, and task scheduling are what determine how that power gets used.

ADI also uses cloud computing resources from Amazon Web Services, or AWS, including EC2 compute instances, S3 storage, and batch processing services.

Why do we need both GPUs and cloud computing?

Because research tasks aren't all the same. Some need parallel acceleration, some need elastic scaling, and some need stable data storage and batch processing.

What I care about more isn't how impressive a device's name sounds, but whether tasks can be assigned correctly, whether resources are used effectively, whether costs are managed, and whether results can be reviewed and verified.

Computing power is a tool. What determines a tool's value is the problems we use it to solve.

## Pietro 12

ADI also uses Google Cloud's TPUs and related computing resources. A TPU is a specialized accelerator designed for specific machine learning workloads, and it isn't simply a substitute for a GPU.

In a mature computing system, CPUs, GPUs, TPUs, and HPC each have the tasks they suit. CPUs handle general-purpose control, data processing, and classical optimization; GPUs suit large-scale parallel computing; TPUs suit neural network workloads that have been adapted for them; and HPC can support large-scale simulation, backtesting, and batch computing.

That's heterogeneous computing.

I don't believe that putting machines from different brands side by side naturally creates a technological edge.

The real difficulty is identifying a task's characteristics, deciding what computing resources it needs, and then making choices among latency, cost, reliability, and data permissions.

ADI needs to master task orchestration and resource scheduling on its own, rather than being locked into one type of hardware.

It's like a team: collaboration doesn't mean everyone does the same thing. Different capabilities taking on the right tasks at the right time is what makes the whole more effective.

ADI's computing efficiency comes from the right division of labor, not from blind accumulation.

## Pietro 13

Guys,

I've always paid close attention to quantum computing, but I'm unwilling to package it as magic that solves every complex problem.

Three concepts need to be told apart here.

Quantum-inspired optimization is algorithms that borrow quantum ideas but can run on traditional CPUs or GPUs; real quantum computing uses quantum processors, or

QPUs, to run quantum algorithms; and hybrid quantum-classical computing has classical computing resources and QPUs work together on experimental tasks. ADI's quantum research and experimentation ecosystem uses IBM Quantum with Qiskit and Amazon Braket, and also uses NVIDIA CUDA-Q and cuQuantum for quantum-classical development and GPU simulation. These tools let us explore combinatorial optimization, complex constraints, and some financial numerical problems.

Any new method has to be compared with mature classical methods: Are the results better? Is the cost lower? Is the stability sufficient? Can it be repeated?

The real scientific spirit is being willing to let new technology be tested.

## Pietro 14

Now, let's put computing resources back into ADI's overall architecture.

ADI is made up of six core intelligence modules: the Financial World Model, the Adaptive Model Foundry, the Machine Investment Committee, Scenario & Causal Intelligence, Heterogeneous & Quantum Computing, and Intelligent Risk Governance & Evolution.

The Financial World Model helps us understand macro, corporate, and cross-asset relationships.

The Adaptive Model Foundry is responsible for model training, validation, and monitoring.

The Machine Investment Committee has different agents raise supporting and opposing views.

Scenario & Causal Intelligence studies different paths and their possible consequences.

Heterogeneous computing provides tasks with CPU, GPU, TPU, HPC, and experimental QPU resources.

Risk governance constrains authority, risk budgets, and execution discipline.

They aren't six unrelated boxes. When there's a gap in the data, research needs to fill it in; when a model's performance drifts, it needs to be re-validated; when

scenario analysis finds that risk has exceeded its limits, the governance mechanism has to step in.

Computing resources let tasks run, the system architecture lets tasks work together, and the governance mechanism lets results be used within boundaries.

What I hope ADI becomes isn't the model with the loudest voice, but a decision-making system that can listen to different evidence, accept rebuttals, reassess, and take responsibility.

## Pietro 15

Guys,

We started from the AI capital markets, passed through mathematics, finance, computational science, and investment practice, and finally came back to the simplest question: in a world that can't be fully predicted, what is worth holding on to?

I believe there are three things.

First, respect facts. Don't ignore opposing evidence just

because a view fits your expectations.

Second, respect risk. Don't assume any price and any financing structure is reasonable just because a technology has enormous prospects.

Third, keep the ability to correct course. Methods that worked in the past aren't guaranteed to keep working in the future; a truly mature system must allow itself to be tested and challenged, and must also allow itself to change.

That's also why I founded Eonvera and keep driving the ongoing buildout of ADI. We can't make the world fully certain, but we can make research more rigorous, the decision-making process more transparent, and risk-taking more clearly bounded.

If you remember only one sentence today, I hope it's this: true rationality isn't always knowing the answer. It's still knowing how to act responsibly when the answer is uncertain.

I'll hand the final time back to Audrey.

Thank you all, and see you next time!

## Audrey 16

Dear investor friends,

Thank you to Dr. Campanelli for today's sharing.

From the capital pressure in the AI industry, to his years of study at Princeton and Stanford; from his practice on Wall Street, to the founding of Eonvera and ADI, what we heard today wasn't only the development of technology, but an attitude toward facing uncertainty.

I especially liked one point Dr. Campanelli just made: rationality isn't always knowing the answer. It's still knowing how to act responsibly when the answer is uncertain.

That may be the meaning of our continued learning. We don't have to pretend we can predict every market change, but we can look at one more piece of evidence, ask one more why, and leave room for our own judgments to be corrected.

If you still have questions about ADI's computing infrastructure, the six intelligence modules, or the AI

capital issues we covered today, you can reach me through the following:

May you stay curious and stay prudent; embrace technology, and respect risk.

See you tomorrow! ❤