THE JOURNAL OF
INVESTMENT CONSULTING

A reprinted article from Volume 24, Number 1, 2025

THE MASTERS SERIES

Andrew W. Lo, PhD: A Renowned Scholar on Investing and Innovation

INVESTMENTS & WEALTH INSTITUTE®

Andrew W. Lo, PhD

A RENOWNED SCHOLAR ON INVESTING AND INNOVATION

Andrew W. Lo is the Charles E. and Susan T. Harris Professor, a Professor of Finance, and the Director of the Laboratory for Financial Engineering at the MIT Sloan School of Management.

Lo's current research spans four areas: evolutionary models of investor behavior and adaptive markets; artificial intelligence and financial technology; healthcare finance; and deep-tech and impact investing. Recent projects include an evolutionary explanation for bias and discrimination and how to reduce their effects; a new analytical framework for measuring the influence of impact investing; the potential for large language models to provide trustworthy financial advice to retail investors; new statistical tools for predicting clinical trial outcomes; incorporating patient preferences into the drug approval process; and accelerating innovation in deep-tech via novel business, financing, and payment models.

Andrew W. Lo, PhD

Lo has been published extensively in academic journals. His most recent book is The Adaptive Markets Hypothesis: An Evolutionary Approach to Understanding Financial System Dynamics. *His awards include Batterymarch, Guggenheim, and Sloan fellowships; the Paul A. Samuelson Award; the Eugene Fama Prize; the IAFE-SunGard Financial Engineer of the Year; the Global Association of Risk Professionals Risk Manager of the Year; the Harry M. Markowitz Award; the Managed Futures Pinnacle Achievement Award; one of TIME's "100 most influential people in the world"; and awards for teaching excellence from both Wharton and MIT. His book* Adaptive Markets: Financial Evolution at the Speed of Thought *won the 2018 Professional and Scholarly Excellence Award for Excellence in Social Sciences from the Professional and Scholarly Publishing Division of the Association of American Publishers.*

He is a fellow of the American Finance Association, Academia Sinica, the American Academy of Arts and Sciences, the Econometric Society, and the Society of Financial Econometrics. Lo is also a principal investigator at the MIT Computer Science and Artificial Intelligence Laboratory, an external faculty member of the Santa Fe Institute, and a research associate of the National Bureau of Economic Research.

He has co-founded several asset management and biotech companies and sits on the boards of several for-profit and nonprofit public and private healthcare organizations. Lo earned a BA in economics from Yale University and AM and PhD degrees in economics from Harvard University.

In October 2025, Inna Okounkova, Journal of Investment Consulting *editor-in-chief; Mark Anson, The Commonfund; Edward Baker, EthicalFin; Ludwig Chincarini, University of San Francisco; Philip Fazio, Merrill Lynch; and Matthew Morey, Pace University spoke with Lo about the state of financial innovation and what this means for the global economy.*

INNA OKOUNKOVA: What were the major factors that helped shape your career and bring you to where you are today?

ANDREW LO: One of the major influences of my career is my upbringing. I'm the youngest of three children, and my older brother and sister are seven and five years older, both scientists. So you can imagine the kinds of discussions we would have at the dinner table throughout my childhood. Being much older than me, my siblings obviously knew more than I did, so it was a struggle for me to keep up, but I had a lot of motivation. My brother became a mathematician and now works at Caltech's Jet Propulsion Lab as a real rocket scientist. And my sister became a developmental biologist, now at the University of Pittsburgh Medical School. So I had to get a PhD just to keep up with them, and being intellectually challenged throughout my childhood (and even today) has been a key component to my current occupation. But the main inspiration was my mother, a single mom who raised three children by herself as an immigrant in New York City, and put them through college all on a secretary's salary with lots of overtime.

Another formative element of my educational journey were some extraordinary mentors that shaped who and what I am today. My undergraduate advisor, Sharon Oster,[1] who passed away a couple years ago from cancer, was a significant reason that I chose to major in economics, despite my STEM [science, technology, engineering, and mathematics] background in high school. Until meeting Sharon, I had planned on pursuing a degree in STEM.

Bob Merton, one of my graduate school professors, was another huge influence.[2] I was a Harvard PhD student who cross-registered for his introductory finance course at MIT, and his lectures were transformative. That class led to my decision to become a financial economist.

INNA OKOUNKOVA: What was your biggest achievement and biggest disappointment?

ANDREW LO: It feels a bit presumptuous to be speaking about any of my achievements as being "great" or "the greatest." So instead, let me just tell you what I'm proudest of and why.

If you asked me this question ten or twenty years ago, I would have said it was the adaptive markets hypothesis [AMH], the idea that markets behave more like complex biological systems subject to the principles of evolution rather than physical systems governed by a small number of immutable mathematical "laws of motion." One of the earliest renditions of this work is contained in an article I published in the *Journal of Investment Consulting* [2005].

The reason I'm so proud of this body of work is because it reconciles the traditional neoclassical economics notion of rational expectations, general equilibrium, and market efficiency with the seemingly contradictory discipline of behavioral finance.

Early in my career, I struggled with the notion of market efficiency—a beautiful theory with logic so compelling that its truth seemed inescapable—in light of the empirical evidence rejecting the random walk hypothesis that Craig MacKinlay and I developed [Lo and MacKinlay 1988,1990]. It was like trying to square the circle. I'd been taught throughout my graduate program that markets are efficient, yet all the data we generated suggested the opposite.

I only began to uncover the deep answers to this conundrum once I started looking into behavioral finance, and then neuroscience and artificial intelligence [AI], and then evolutionary biology and ecology. It was this process of searching for the answers within each of these different disciplines that helped me develop an explanation for how to reconcile the theory of efficient markets with the reality of various empirical violations of efficiency.

I concluded that the efficient markets hypothesis [EMH] isn't wrong—it's just incomplete. We have to consider exceptions to the rule, a rule deemed valuable, generally well-accepted, and supported by economists. But the exceptions are incredibly important and interesting. And over the past two or three decades, we just happen to be living in a time when exceptions occur with increasing frequency, thanks to technological breakthroughs such as algorithmic trading, high-speed networked telecommunications among financial institutions, computation, and artificial intelligence.

This is how I would have answered your question twenty years ago about what, among all my research accomplishments, I'm proudest of. But I've changed my mind since then, partly due to personal challenges. About fifteen years ago, friends and family members were dealing with various kinds of cancer, and over a six-year period, five people close to me all died of this terrible disease. I'd never dealt with death up close and personal before then, so it was a rude awakening for me.

During that time, I felt completely useless and helpless; a cancer patient needs a financial economist about as much as a fish needs a 401(k) plan. It wasn't until I started learning more about how cancer drugs were developed and delivered to patients that I realized something that was obvious to biomedical experts: Finance plays a huge role in the healthcare industry. I was naive enough to think that, if a patient was in need and a particular piece of science or medicine could help, money would magically appear to develop a treatment for that patient. Nothing could be further from the truth. It can be incredibly challenging for even the most brilliant physician-scientist to convince investors to take on the risk of biomedical R&D [research and development] given the "triple whammy" of drug development: low or unknown probability of success, outsized capital needs—often hundreds of millions to billions of dollars per drug—to develop and clinically test a single drug candidate, and an investment horizon that sometimes can stretch beyond a decade before any cash flows are generated. Biomedical investments, especially at the early stages of development, are simply too risky for most stakeholders.

So, over the past two decades, I've re-oriented my focus toward developing new financing models to accelerate biomedical innovation, and this research area has now become the most meaningful to me. I've applied the principles of finance that all of us financial economists know instinctively and use daily—diversification, portfolio optimization, securitization, risk management, and so on. Being able to take those ideas and work with collaborators in the medical industry to put them into practice in specific settings has been some of the most meaningful work I've done in my entire career.

This research agenda is especially satisfying because it shows how finance doesn't always have to be a zero-sum game. There are ways of doing well by doing good and helping patients along the way. Thanks to the progress my collaborators and I have made in this area, I now see so much more potential for how what we do in finance can have a huge positive impact on the world.

Along the same lines, I guess my greatest disappointment is that a set of ideas I proposed for large-scale funding has not yet been implemented by the healthcare industry. In 2012, I published my very first article in a leading biomedical journal (*Nature Biotechnology*) on novel ways to fund drug development, proposing the use of securitization techniques to unlock billions of dollars in capital. I called this idea the "megafund" model, and there were two pieces to it. One was portfolio theory and diversification, the idea that taking "multiple shots on goal," to use a hockey/soccer catchphrase, would improve the odds of at least one or two wins when compared to making bets one at a time. That part has been implemented successfully.

But the second and, I thought, more elegant and important piece, was to finance the development of these therapies via securitized debt rather than just equity. This second part has not yet been done. But I haven't given up.

I'm hoping that within the next few years, we will get to that second part, where we start using securitized debt to finance drug development, unleashing an order of magnitude more capital than what's available today.

LUDWIG CHINCARINI: Why are you keen on debt financing?

ANDREW LO: There are three reasons why I proposed debt financing. First, debt markets are one to two orders of magnitude larger than equity markets—there's a ton of money waiting to be deployed. If you just look at corporate bonds, that's certainly a risky asset class, but bondholders are aware of what the default risks are. They model it, and therefore they are able to invest large amounts of capital for various purposes, including mortgage-backed securities that financed trillions in commercial and residential real estate. The 2008 financial crisis made it obvious to me that a lot of money could be raised in a relatively short period of time for a very specific purpose, if you properly model and manage the risks. That's one reason.

But the second, and probably more relevant, reason is that the typical cost of capital in drug development among biotech companies is estimated to be anywhere from 20–35 percent and higher, depending on the specific stage of R&D. Today's corporate bond yields are in the mid to high single digits for investment-grade bonds. It's a huge efficiency gain if you can finance biotech via securitized debt, assuming the risk can be accurately measured and managed, which in turn is a function of analytics and scale.

Finally, it's because one of the key concerns of biotech entrepreneurs is dilution. Drug development requires multiple rounds of funding, and at each round these entrepreneurs face greater dilution. In fact, it's not at all unusual for a founder—the person who came up with the drug and got it all the way to FDA [U. S. Food and Drug Administration] approval—to own less than 5 percent of the equity of the company by the end of that very expensive process. Securitized debt is a source of non-dilutive financing, allowing equity holders to maintain their equity stake and get "more bang per buck" due to increased leverage.

From multiple perspectives, it seems like everyone benefits. The problem is that we don't yet have a widely accepted set of analytics that will allow bondholders to accurately measure the risk and reward of a given investment in biomedical assets. We're getting closer. I do believe that will happen sometime in the next few years. But I've been feeling that way for the past five years, so who knows whether it will really come to fruition.

EDWARD BAKER: I continue to be impressed by your work on the AMH and believe it offers an important perspective on financial market dynamics, most deeply reflected in your paper "The Origin of Behavior" [Brennan and Lo 2011]. You first published your work on the AMH in its most formal version in the *Journal of Portfolio Management* in 2004, their 30th anniversary edition.

At the end of that paper, you said you were looking forward over the next thirty years, and hoped your paper received a lot of mentions in the journals to come. Has it achieved what you hoped, at least in terms of bridging the gap between the EMH and the various quirks and anomalies identified by behavioral finance?

ANDREW LO: Well, I believe it's still early days for the AMH. As the saying goes, the proof of the pudding is in the eating, and that will require other researchers to weigh in and generate additional theoretical and empirical analysis. I don't think there's been enough done yet.

Part of the reason is the way the AMH frames the problem. It's not the traditional approach that economists have borrowed from physics—develop a mathematical hypothesis for a given law of motion, collect data, and then test the hypothesis. Testing the AMH is not quite as simple as testing whether the intercept or alpha is zero in a standard one-factor CAPM [capital asset pricing model] regression.

It's more akin to an ecological research problem. If you ask an ecologist whether a particular ecosystem is sustainable or not, the ecologist doesn't usually begin by testing any single mathematical formula, even though mathematics does play a big role in ecology and evolutionary biology. They typically start by spending some time doing field studies, living in that particular ecosystem, measuring and cataloging the different species and biomasses of flora and fauna, and documenting how they relate to each other, e.g., via predator–prey relationships. They'll basically immerse themselves in that ecosystem. After six months to a year, or however long is needed to complete their field studies and create case studies of the various species, these scientists then will turn to modeling the mathematical relationships among those species. Only then will they be able to say: "Yes, if you eliminate this portion of the forest, you'll lose certain keystone species, which may then lead to a collapse of the food chain, and the entire ecosystem could be in jeopardy due to a sudden decline in biodiversity."

Although the AMH has now received additional exposure in the mainstream finance literature, it doesn't have nearly the same level of empirical validation that the EMH does. The AMH does have growing empirical validation in the context of providing evolutionary explanations for behavioral biases such as loss aversion, anchoring, and framing. But because the AMH is a much more complicated way of looking at things, testing it requires a fundamentally different approach. In fact, I only came to the AMH after taking a two-year leave of absence to start my own hedge fund, AlphaSimplex Group, in 1999.

It was through the process of putting money on the line—developing trading strategies, seeing those strategies generate significant profits initially, and then watching those profits decline over time as other trading groups reverse-engineered them and competed away those profits—that I realized markets don't work the way the EMH tells us they should. They work more like the way biologists and ecologists think about how the flora and fauna of an ecosystem evolve in the face of competition, innovation, mutation, and natural selection.

EDWARD BAKER: It seems that a lot of behavioral research is attacking the assumptions of the EMH, such as the assumption of rationality, rather than attempting to disprove it directly. Beyond that kind of thinking, has the AMH achieved the prominence you hoped for? Is it getting the publication mentions hoped for? Are other academics using it in their work?

ANDREW LO: Not as much on the academic side. It's funny—practitioners are more familiar with it than my academic colleagues. Once a practitioner reads about the AMH, they usually respond positively, saying, "Of course, this is exactly what's been going on." Not surprisingly, this is because the AMH emerged as my attempt to reconcile theory with practice, not so much as an academic exercise. However, the rejection of the random walk hypothesis that Craig [MacKinlay] and I wrote was the beginning of the academic journey, and I believe that economists are beginning to appreciate that body of work a bit more in the context of the AMH.

But because I didn't frame it as a simple alternative, like a two-factor alternative to the CAPM, or the Fama-French five-factor model, I didn't work out all the implications of the AMH for the existing academic literature. I think that's why it may not have gotten more acceptance or visibility among academics. In fact, the first time I published anything about the AMH was in a practitioner journal. And the first complete exposition of the AMH was published as a trade book in 2017 [Lo 2017], written for a general audience, not my academic peers. I felt that the multi-disciplinary nature of the AMH required a broader and more accessible approach, so that even non-specialists could understand the overall logic and vision of the AMH.

However, more academics are starting to take an interest in it, and in 2024 I published a much more technical exposition of the AMH [Lo and Zhang 2024] that contains a lot of the mathematical and statistical implications that academics like to focus on, which were deliberately omitted from the 2017 version.

In fact, "The Origin of Behavior," the paper you kindly mentioned [Brennan and Lo 2011], led to one of the biggest aha moments I had in my entire career. That one paper crystallized for me the solution to the decades-old controversy between the EMH and its behavioral critics, and when Tom and I completed it, I said, "Now I finally get it." In my 2024 monograph, published by Oxford University Press as part of their Clarendon Lectures in Finance, I hope to reach my academic colleagues and show them that market efficiency and behavioral finance are not competing theories but merely opposite sides of the same coin of adaptive human behavior in a complex and highly dynamic environment.

EDWARD BAKER: Have there been any surprises along the way in your further research?

ANDREW LO: Well, I certainly didn't expect this research, starting with the simple idea that individuals' behaviors are shaped by their environment, would generate so many different, rich implications. Nor did I expect that some of those implications would be relevant for the readers of the evolutionary biology literature. I didn't start out looking to contribute to that literature; I liked my chosen field of finance just fine and had no intention of moving into other fields. It just turned out that, in seeking answers to questions about financial anomalies, I discovered certain contradictions in some of the things that evolutionary biologists had concluded.

One example is the idea of "group selection": Can natural selection operate at the level of a group rather than at the level of an individual gene? The work I published on the AMH ended up having implications for the theory of group selection in evolutionary biology. To make sure I wasn't totally off base, I published that research in peer-reviewed scientific journals, and the referees, who were evolutionary biologists, confirmed that the AMH was a different but valid way of understanding group selection.

THAT ONE PAPER CRYSTALLIZED FOR ME THE SOLUTION OF THE DECADES-OLD CONTROVERSY BETWEEN THE EMH AND ITS BEHAVIORAL CRITICS, AND WHEN TOM AND I COMPLETED IT, I SAID, 'NOW I FINALLY GET IT.'

EDWARD BAKER: In what ways would you like to see your research impact others' thinking?

ANDREW LO: I guess I've always felt there should be some consistency between what academics and practitioners do. To give you a little bit of a window into my own psychology, one of the reasons I decided to start my own asset management company in 1999, AlphaSimplex Group, was because I'd been teaching investments at MIT for years. I taught my students the EMH because that's exactly what I was told to teach by my older and wiser colleagues. I took their teaching notes as my starting point and tried to teach the subject to the best of my abilities. But I have to tell you, I felt like a total fraud for two reasons.

First, I didn't believe what I was teaching. But second, and more relevant, I actually didn't know very much about investing. Of course, I understood all the academic models of investment management, and had authored a number of those models in my research career by then. But I had never invested personally before I became an assistant professor of finance and started teaching investments. I did do some consulting for a few large financial institutions, but that was different. The institutions were the ones doing the investing, not me. So when I started AlphaSimplex Group, I felt I was taking an important step forward because I was putting my money where my mouth was, both literally and figuratively, from a career perspective.

That was a hugely formative experience and I learned a tremendous amount. The point of the AMH is that, now, I believe practitioners and financial economists in academia can be on the same page. Differences may still exist because we look at the world through different lenses and use different vocabulary to describe what we see. But there's a lot more agreement than academics and practitioners realize when you look through the lens of the AMH. It's like the old adage about five blind monks who encounter an elephant for the very first time. The monk who grasps the elephant's trunk says the elephant is just like a snake, while the monk who feels the elephant's leg says the elephant is just like a tree. They're not wrong, they just don't have the complete picture.

This aspect of the AMH was, and still is, particularly satisfying. I finally felt like I wasn't talking to two completely different groups of stakeholders who had nothing in common. We academics have a lot in common with our practitioner colleagues, but it may not seem so because we don't have a common language for exchanging ideas. The AMH provides that language.

PHILLIP FAZIO: I read *A Non-Random Walk Down Wall Street*. Certainly, for most of us, the math in that book is intense. There was an update if I recall. What are your observations today about non-randomness? Are there significant revelations that modify your work?

ANDREW LO: That was indeed an intense book, and I apologize for your having to wade through the thicket of equations and tables and figures. It took ten years of our academic careers for Craig and me to work all of that out, and I wouldn't expect too many people to read it for pleasure.

But that book was important for my professional development because I desperately wished for the EMH to be true. That was how I was taught. I learned about the EMH first from Bob Merton, and then from Fischer Black, and yet again from Stew Myers. All the teachers I had at Harvard and MIT believed that the EMH was true. That's why Craig and I thought we might have had a programming error in our code.

And the reason it took a decade for us to finally conclude that stock prices aren't random walks is because we tried to explain away the rejections time after time, using various explanations. We developed a model of non-synchronous trading. We developed a model of illiquidity. We developed a model of price discreteness (in those days, prices traded on discrete intervals of $0.125). We tried every possible angle to conclude that markets were in fact efficient and we were incorrectly rejecting the random walk due to one of these attributes.

In the end, this book—aptly titled *A Non-Random Walk Down Wall Street*—became a compendium of all our failed attempts to explain away the evidence as anything other than market inefficiency. Ultimately, what finally convinced me to stop doing my testing and just acknowledge markets can be inefficient at times was when I attended a talk at MIT by a hedge fund manager who was invited by the MIT Computer Science and AI Lab. He was a computer scientist and one of the first to create a statistical arbitrage, or "stat arb," hedge fund. I'd never met him before his visit and hadn't heard of him because I wasn't involved with the industry at the time. I was just writing papers and trying to get tenure.

He came up to me after the talk and seemed to know a lot about my research with Craig on testing the random walk, including the key equations in our 1988 paper. When I asked him how he knew so much about our work, he said that his first trading strategy made use of our research.

I was blown away. Some people have asked whether I felt bad that I didn't do it first. I said, "Of course not." I was in no position to do what this talented trader did. The fact that he was a computer scientist meant that he had no interest in academic theories—he just wanted to develop a profitable trading algorithm. That's not something Craig and I could have done when we published the paper in 1988.

But his experience told me that Craig and I were right. Markets don't follow random walks at times. If you have the right hardware, software, telecommunications, financing, people, and circumstances, one could profit handsomely from our ideas. And this talented hedge fund manager did just that, until it stopped working.

This notion of market efficiency as a black-and-white phenomenon basically went out the window when I met this impressive trader. He acknowledged this was their first trading strategy, and eventually everybody else figured it out and copied it with minor variations, at which point he turned to another set of strategies.

So my publications with Craig may have contributed to his success, but only a very small portion. However, his willingness to share his experience with me was invaluable, and gave me a clear sense of how dynamic financial markets must be and how practitioners look at the world differently because they're forced to. They're just trying to make a living, and lots of other talented hard-working professionals are competing with them to do the same.

Thus began my path down this AMH journey.

MATTHEW MOREY: I sometimes use *How Markets Fail* by John Cassidy when I teach my courses. The book was published in 2009. I really like the book. I think it's an intellectual tour de force. One of its basic central theses is that the Chicago style or concept of "utopian economies" equates speculation to stabilizing. If the market is overvalued, people will see that and push it downward and vice versa. He counters this is not the case at all. Speculation is in fact destabilizing. What are your thoughts on this, especially given where we are in the markets today?

ANDREW LO: I spent a lot of time thinking about this idea as an undergraduate and graduate student. Given my high school background in STEM studies, I naturally gravitated to mathematical economics, due in large part to Sharon Oster's mentorship. I studied general equilibrium theory and the beautiful mathematics involved in that literature, and wrote my senior thesis on that topic. I was fortunate to have studied under some of the very best mathematical economists such as Herb Scarf, Pradeep Dubey, and Martin Shubik.

What I learned from that experience is that neoclassical economics is akin to the simplified world of frictionless physics. It assumes that markets are always in equilibrium and there is no cost to engaging in the price discovery process. The EMH conveys the same idea, that prices always reflect all available information, because it was implicitly assumed that markets equilibrate instantaneously. [Eugene] Fama didn't say it explicitly when he formulated the EMH, but it's embedded in the theory.

This is an important point because, in the event of a supply and demand imbalance, prices for that moment may not reflect all available information. But if you assume all supply and demand imbalances are eliminated immediately, then you're back in the happy world of efficient markets and rational expectations. The question remains, "Why do prices equilibrate at all?"

This question was asked and answered over a century ago by the French economist Léon Walras, who came up with a clever idea now known as the tâtonnement or "groping" process. Imagine a hypothetical auctioneer whose job it is to raise or lower prices in response to market demand. In the face of excess demand, the auctioneer raises prices until such demand is eliminated. If there's excess supply, the auctioneer will lower prices until it's eliminated. Through this process of successive approximations, the Walrasian auctioneer will reach market-clearing prices quickly and efficiently. This tâtonnement, or iterative process of trial and error, is how economists typically think about equilibrium prices, at least in theory. Of course, practice can be quite different.

When I looked at price movements on the New York Stock Exchange and Nasdaq, I didn't observe tâtonnement happening instantaneously. And there are days such as October 19, 1987, where it doesn't happen at all. Markets occasionally break down completely.

Walras was thinking like a mathematician, not a Wall Street trader or market maker. The process of equilibrating markets takes time, effort, and skill. If you acknowledge that aspect of the traditional economic paradigm you quickly come to the AMH, where prices don't always and everywhere fully reflect all available information at all times.

… I don't believe markets are anywhere close to being in equilibrium at all times. Disequilibrium is the rule, not the exception.

On the other hand, given the intense competition in financial markets, it's not trivial to exploit market inefficiencies when they do arise. Even if there's an imbalance, to arbitrage supply/demand gaps requires a lot of capital. In my view, you need to understand what the underlying fundamentals are to be able to exploit temporary disequilibria.

In fact, I don't believe markets are anywhere close to being in equilibrium at all times. Disequilibrium is the rule, not the exception. The most talented hedge fund managers—individuals like Warren Buffett, David Shaw, Jim Simons, and George Soros—are able to identify the nature of this disequilibrium at different horizons and take advantage of that knowledge.

MATTHEW MOREY: I remember reading about the Shleifer noise trader model. It reflects the idea that it's effectively better to ride the bubble than to fight it. It seems consistent with what you're talking about, absolutely.

ANDREW LO: If you can manage the risk, and know when to get out, or at least get out without too big a loss, then you can keep some of those profits. I would call that alpha. That's where the reality of financial markets meets the fictional theory that we all know and love. It's also why I believe you can integrate the two into a single narrative that practitioners and academics can agree on.

LUDWIG CHINCARINI: I interviewed Chi-Fu Huang for my book, *The Crisis of Crowding* [2012], when I was trying to understand all these peculiarities happening in the markets. I think he was a co-professor with you at MIT at one point and left for LTCM [Long-Term Capital Management].

He said the following: "The big innovation in finance was separating alpha from beta and that investors should focus on alpha and leverage." In fact, he believes that the 2008 financial crisis erased thirty to forty years of academic theory. What is your opinion on this statement?

ANDREW LO: Well, I think the statement needs to be qualified, specifically around the conclusion that it doesn't work. It doesn't work for whom?

If it doesn't work for Chi-Fu, or John Meriwether or George Soros, I get that. But I think it does work for most retail investors.

Here's what I mean. If you don't have the time, inclination, and resources to start your own hedge fund and be able to trade against institutions like Citadel, D. E. Shaw, Renaissance, and Two Sigma, then looking at the world in terms of alpha versus beta is a pretty good approach, along with putting your money in a passive index fund or ETF [exchange-traded fund]. Investing in passive low-fee financial products during a twenty- or thirty-year period prior to retirement has demonstrated tremendous value-added for generations of workers since Jack Bogle created the very first index fund in 1976.

So, I would venture to guess that when Chi-Fu talked about alpha versus beta, he was probably thinking more like a hedge fund manager than he was as a retail investor with no specific financial training.

The reason I love the field of finance so much is because what we do broadly applies to real-life situations. This notion of systematic versus idiosyncratic risk applies equally to retail investors and pension funds that don't want to engage in active management. Both parties want to index to a benchmark and grow their financial wealth in a responsible way.

At the same time, finance also allows for the existence of extraordinary people like Warren Buffett, David Shaw, Jim Simons, and George Soros. They have skills that not all of us possess. Those skills have value. So alpha does exist in those circumstances.

LUDWIG CHINCARINI: Can I just add to your statement? When I finished at MIT, I went to Wall Street and so did a student of yours, Toshi Baily, I think you remember.

ANDREW LO: I do indeed.

LUDWIG CHINCARINI: We upset a lot of people by leaving. But I thought, "Okay, I'm going to use my training and blow away these Wall Street guys."

Instead, every model I built could not beat the market. All of them failed and I was utterly depressed. Given your mention of these successful investors, what do you think "it" is? With AI [artificial intelligence], do you think people who claim you can use AI to pick the best stocks for the next week are right? I have a friend who is doing that. What do you think is the skill you need? Is it understanding quantitative investing? Is it combining a quantitative approach with other factors? What is a skill you need to be one of those extraordinary guys?

ANDREW LO: That's a question I've been asking myself for a long time. I ended up developing an answer that satisfied me, but maybe not all of my academic colleagues, in an article I published called, "Where Do Alphas Come From?" in the *Journal of Investment Management* [2008].

In that paper, I tried to reconcile exactly the point you're making. How does someone like Warren Buffett coexist with someone like the day trader Steve Cohen? How do I think about what they do? Are they so different that they have nothing to do with each other? There is no way finance can explain or even be able to frame what they do, other than to say that they are anomalies. They don't fit our model.

But I came up with what I thought was a satisfactory mathematical framework to help academics and investment managers understand where alpha comes from. The answer is so simple and obvious, which is why I think others hadn't looked at it this way before. All active investment returns stem from the ability to predict the future. In statistical terms, this can be characterized by positive correlation between a manager's portfolio weights and asset returns In other words, what we mean by "investment alpha" is simply the ability to do "market timing" but at the asset level, i.e., buying an asset when it's about to go up and selling or shorting an asset when it's about to go down.

This idea may seem trivial, but it leads to a very concrete metric for capturing investment alpha that isn't based on the CAPM or any other asset-pricing model—it's based on the ability to predict future asset returns and take advantage of it by changing one's portfolio weights to exploit this predictability. [A more technical explanation of this measure is included in the appendix.]

Now getting back to the original question of a common framework to understand how Warren Buffett and Steve Cohen both make money. Both are predicting future asset returns and adjusting their portfolio weights to exploit these predictions. But Cohen's predictions are being made at the five-minute level whereas Buffett's are at the five-year level. They have different information sets that update at vastly different time horizons, but they are both doing the same basic thing.

All investors make predictions. The key to making reasonable predictions, as instantiated by their portfolio weights, is to use an appropriate time horizon. This gives us a way to analyze the skills of all investors as a function of their weights/returns covariances.

MARK ANSON: Your company QLS [Quantitative Life Sciences] has developed a proprietary model to estimate the probability of success of clinical drug trials. To what extent have you embedded AI into your probability of successful models?

ANDREW LO: Thank you, I'm really glad you asked this question because it's completely outside of finance, so I wouldn't have normally brought it up myself. However, as I said earlier, there is a link.

Recall my previous comment about my great disappointment that the biotech and pharmaceutical industries have not yet embraced the idea of using securitization for drug development. I said that the issue was missing analytics—we need to be able to calculate default probabilities.

One of the reasons the mortgage-backed securities markets grew so rapidly was an innovation in analytics: a formal mathematical model of prepayment risk, the so-called "Gaussian copula formula."[3] This model enabled all these mortgage-backed securities traders to be able to assess the risk of their positions and generate alpha or beta through that framework.

When I first started asking why securitization isn't being done, I learned there are no equivalent models for calculating the default probability of a portfolio securitized by biomedical assets. To assess default risk, you have to be able to forecast the likelihood that the FDA will approve or deny a particular drug candidate—an estimated probability of success for a given clinical trial. Basically, you need to create a Barra model for drug development.[4] And that didn't exist when I first started working on healthcare finance in the 2010s.

THE KEY TO MAKING REASONABLE PREDICTIONS, AS INSTANTIATED BY THEIR PORTFOLIO WEIGHTS, IS TO USE AN APPROPRIATE TIME HORIZON.

I was kind of shocked by that because the biopharma industry has been around a lot longer than I have, and there are a lot of smart, quantitative people in the field of drug development such as scientists and engineers, MDs, PhDs, and MD/PhDs. I was even more surprised to learn that, although there were some efforts to calculate historical rates of success in clinical trials, nobody had ever used AI before—particularly machine learning—to make forecasts of drug approvals.

I'd been using artificial neural networks for a long time to predict defaults for consumer credit, and published a number of articles using random forests, multi-layer perceptrons, radial basis functions, and other AI techniques. Apparently, no one in the biopharma industry had ever thought to apply these same models to predicting drug development outcomes.

To help fill this void, I published several papers on how to do so using various machine learning techniques [Wong et al. 2019; Lo et al. 2019;

Siah et al. 2021; Xu et al. 2022]. I figured once I published these tools, biomedical experts would start using them, and then we'd have a Barra model for the biopharma industry. But this didn't happen. The work was greeted initially with silence and indifference.

So together with my former PhD student/post-doc, Shomesh ["Sho"] Chaudhuri, we decided to take on this challenge and co-founded a company called QLS Technologies. Borrowing from my experience in the financial industry, in starting AlphaSimplex and building risk models for our internal use, and Sho's PhD research with me on healthcare finance, we created a comprehensive quantitative portfolio management framework for investors in biomedical assets. Our models are based on machine learning techniques and, most recently, large language models, and we now provide this analytics platform to biotech and pharma companies, as well as to investors. We hope that in the near future, these analytics will change the way investors look at risk and reward in the healthcare industry and create new sources of capital for this important sector.

EDWARD BAKER: You've written a lot about neuroscience. Your paper "The Origin of Behavior" touches on the issue of non-rationality and how it fits into the behavioral framework. What conclusions have you come to about the importance of non-rationality as an element of investment decision making? How does it blend with rationality? Do you think that part of what investment advisors should do is to serve as coaches and help clients manage these psychological tensions, so they make better investment decisions?

ANDREW LO: When I first encountered behavioral finance literature I was initially very excited because I thought it was going to offer an alternate theory to market efficiency. But I was ultimately disappointed because the behavioralists didn't come up with any alternative. Instead, they came up with a host of counterexamples of the EMH. Now, I'm not criticizing the pioneers of behavioral finance; they provided an important contribution to the profession by cataloging all the apparent failures of the EMH. But someone wise once said that "it takes a model to beat a model," and the behavioral finance literature had no model that could challenge the EMH. So what do you do with all their counterexamples? What should we replace the EMH with?

This dissatisfaction prompted me to spend time reading the neurosciences literature, in an attempt to understand why people behave the way they do. That was an eye-opener. I began to understand that what we call "rationality" is a delicate balance among the multiple components of our cognitive faculties. It involves being able to calculate, it involves memory, but it also involves emotion. You know, patients who have suffered damage to areas of the brain that are responsible for emotion—and therefore can't feel any emotion—end up acting quite irrationally sometimes. For example, they focus on one task completely, mindlessly, and not realizing they're late for an appointment or missing some other more-important task. Prioritization, it turns out, requires emotional states such as anxiety, fear, guilt, shame, excitement, enthusiasm, etc. If you eliminate emotion, you eliminate rationality. On the other hand, too much emotion also can make us irrational. For example, it's a well-established fact that love makes you stupid. I've re-learned that lesson firsthand several times through a number of bad relationships.

Rationality, from a neuroscientific point of view, differs from that of economics: It's a state of mind that requires the proper balance between IQ and EQ [emotional quotient].

Financial advisors could all benefit from a bit of neuroscience so they can explain to their clients why retail investors sometimes act in ways that are counterproductive to their financial health. Investors should focus on how they might be able to be more productive by changing their perspective, acknowledging that sometimes they enter emotional spirals that make it difficult for them to act rationally. In those circumstances, having professional advisors—or, nowadays, AI chatbots—may help restore balance to our temporarily misaligned cognitive functions.

INNA OKOUNKOVA: It's been a great pleasure to discuss your cutting-edge work, but we've unfortunately run out of time here. Is there anything you would like to add?

ANDREW LO: We've covered a lot. I just want to say that it's incredibly gratifying that colleagues whom I've known and respected for so many years would take so much time reading and thinking about my work. Thanks to you all, Inna Okounkova, Mark Anson, Edward Baker, Ludwig Chincarini, Philip Fazio, and Matthew Morey for this great pleasure and honor.

ENDNOTES

1. American economist Sharon Monica Oster (1948–2022) was the Frederic D. Wolfe Professor Emerita of Management and Entrepreneurship and the dean of the Yale School of Management, where she was the first woman to receive tenure as a professor at the Yale School of Management and the first woman to be named its dean, serving through 2008–2011 and the Great Financial Crisis. During her forty-three-year career at Yale, Oster became an expert in competitive strategy, microeconomic theory, industrial organization, the economics of regulation and antitrust, and nonprofit strategy. She strongly influenced the curriculum of the Yale School of Management, where she three times received the school's Award for Excellence in Teaching. See https://som.yale.edu/story/2022/sharon-oster-1948-2022 and https://www.nytimes.com/2022/06/14/business/sharon-oster-dead.html.
2. Robert C. Merton is the School of Management Distinguished Professor of Finance at Massachusetts Institute of Technology, and the John and Natty McArthur University Professor Emeritus at Harvard University. He is currently resident scientist at Dimensional Fund Advisors, where he is the creator of Target Retirement Solution, a global integrated retirement-funding solution system. See https://robertcmerton.com/about/.
3. *Wired* magazine's Felix Salmon (2009) published the article "Recipe for Disaster: The Formula That Killed Wall Street" in the aftermath of the 2008 financial crisis and specifically identified David Li's (2000) Gaussian copula formula as the culprit. The role of analytics in the financial crisis was subsequently studied in greater depth by MacKenzie and Spears (2014).
4. Barra refers to the multi-factor model, pioneered by Barr Rosenberg in 1975, that measures risk associated with a security relative to the market. These models help investors and portfolio managers forecast risk for

equity, fixed income, cash, and derivative instruments, at both the asset and portfolio level. See https://www.investopedia.com/terms/b/barra-risk-factor-analysis.asp.

REFERENCES

Brennan, T. J., and A. W. Lo. 2011. The Origin of Behavior. *Quarterly Journal of Finance* 1, no. 1 (2011): 55–108.
Cassidy, J. 2009. *How Markets Fail: The Logic of Economic Calamities*. Farrar, Straus and Giroux.
Chincarini, L. 2012. *The Crisis of Crowding: Quant Copycates, Ugly Models, and the New Crash Normal*. Wiley/Bloomberg.
Fernandez, J.-M., R. M. Stein, and A. W. Lo. 2012. Commercializing Biomedical Research Through Securitization Techniques. *Nature Biotechnology* 30 (October): 964–975. https://doi.org/10.1038/nbt.2374.
Li, D. X. 2000. On Default Correlation. *Journal of Fixed Income* 9, no. 4: 43–54.
Lo, A. W. 2004. The Adaptive Markets Hypothesis. *Journal of Portfolio Management* 30, no. 5 (30th Anniversary Issue): 15–29.
——. 2005. Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting* 7, no. 2: 21–44. https://ssrn.com/abstract=1702447.
——. 2008. Where Do Alphas Come From? A Measure of the Value of Active Investment Management. *Journal of Investment Management* 6, no. 2: 1–29.
——. 2017. *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton University Press.
Lo, A. W., and A. C. MacKinlay. 1988. Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test. *Review of Financial Studies* 1: 41–66.
——. 1990. When Are Contrarian Profits Due to Stock Market Overreaction? *Review of Financial Studies* 3, no. 2: 175–205.
Lo, A. W., K. W. Siah, and C. H. Wong. 2019. Machine Learning with Statistical Imputation for Predicting Drug Approval. *Harvard Data Science Review* 1, no. 1. https://doi.org/10.1162/99608f92.5c5f0525.
Lo, A. W., and R. Zhang. 2024. *The Adaptive Markets Hypothesis: An Evolutionary Approach to Understanding Financial System Dynamics*. Oxford University Press.
MacKenzie, D., and T. Spears. 2014. The Formula that Killed Wall Street: The Gaussian Copula and Modelling Practices in Investment Banking. *Social Studies of Science* 44, no. 3 (June): 393–417.
Salmon, F. 2009. Recipe for Disaster: The Formula That Killed Wall Street. *Wired* (February 23). https://www.wired.com/2009/02/wp-quant/.
Siah, K. W., N. W. Kelley, S. Ballerstedt, B. Holzhauer, T. Lyu, D. Mettler, S. Sun, S. Wandel, Y. Zhong, B. Zhou, S. Pan, Y. Zhou, and A. W. Lo. 2021. Predicting Drug Approvals: The Novartis Data Science and Artificial Intelligence Challenge. *Patterns* 2, no. 8 (August) : 100312.
Walras, L. 1874. *Eléments d'Économie Politique Pure ou Théorie de la Richesse Sociale*. Lausanne/Paris: Corbaz & Cie.
Wong, C. H., K. W. Siah, and A. W. Lo. 2019. Estimation of Clinical Trial Success Rates and Related Parameters. *Biostatistics* 20, no. 2: 273–286.
Xu, Q., E. Ahmadi, A. Amini, D. Rus, A. Lo. 2022. Identifying and Mitigating Potential Biases in Predicting Drug Approvals. *Drug Safety* 45: 521–533.

Appendix

If you think about how to characterize the expected return of a given trading strategy, formulaically, how would you do that? Well, we can start by calculating the expected return of the strategy, which in turn is just the expectation of the sum of the products of individual portfolio weights times the corresponding asset returns. In finance textbooks, the formula usually appears as something like this:

$$R_p = \omega_1 R_1 + \omega_2 R_2 + \ldots + \omega_n R_n$$
$$\mathrm{E}[R_p] = \mathrm{E}[\omega_1 R_1 + \omega_2 R_2 + \ldots + \omega_n R_n] = \mathrm{E}[\omega_1 R_1] + \mathrm{E}[\omega_2 R_2] + \mathrm{E}[\omega_n R_n] \quad \text{(A1)}$$

So far, this derivation isn't particularly controversial or interesting. But the next step is both:

$$\mathrm{E}[\omega_1 R_1] + \mathrm{E}[\omega_2 R_2] + \ldots + \mathrm{E}[\omega_n R_n] = \omega_1 \mathrm{E}[R_1] + \omega_2 \mathrm{E}[R_2] + \ldots + \omega_n \mathrm{E}[R_n] \quad \text{(A2)}$$

The right side of this expression simply moves the portfolio weights outside of the expectations operator, a step that every finance academic has done countless times in calculating the expected return of a portfolio. But this operation is only correct if portfolio weights are fixed constants. In practice, portfolio weights are in fact functions of various inputs such as corporate announcements, macroeconomic news, fund inflows and outflows, and a portfolio manager's view of that asset's future prospects. These considerations imply that the weights are themselves random variables, hence they cannot be brought outside the expectations except under very restrictive statistical assumptions that generally aren't satisfied in practice.*

Because of the following statistical property of any two random variables and

$$\mathrm{E}[XY] = \mathrm{Cov}[X,Y] + \mathrm{E}[X]\mathrm{E}[Y] \quad \text{(A3)}$$

we can rewrite (A2) in the following interesting way:

$$\begin{aligned}\mathrm{E}[R_p] &= \mathrm{E}[\omega_1 R_1] + \mathrm{E}[\omega_2 R_2] + \ldots + \mathrm{E}[\omega_n R_n] \\ &= (\mathrm{Cov}[\omega_1, R_1] + \mathrm{E}[\omega_1]\mathrm{E}[R_1]) + \ldots + (\mathrm{Cov}[\omega_n, R_n] + \mathrm{E}[\omega_n]\mathrm{E}[R_n]) \\ &= (\mathrm{Cov}[\omega_1, R_1] + \mathrm{Cov}[\omega_2, R_2] + \ldots + \mathrm{Cov}[\omega_n, R_n]) + (\mathrm{E}[\omega_1]\mathrm{E}[R_1] + \\ &\quad \mathrm{E}[\omega_2]\mathrm{E}[R_2] + \ldots + \mathrm{E}[\omega_n]\mathrm{E}[R_n]) \quad \text{(A4)}\end{aligned}$$

The last expression in (A4) says something important. As an active trader, there are two ways you can make money for your investors. One way is to generate positive covariance between portfolio weights and asset returns, i.e., have higher weights in assets when their returns are higher, and lower weights in assets when their returns are lower. The second way is to have larger positive average holdings in assets with larger positive average returns, i.e., have higher average weight on assets with higher risk premia. The first source of expected return is from being able to predict future returns—a dynamic activity that implies changing weights over time. The second source is from holding higher-yielding assets. This decomposition provides a natural interpretation of active versus passive generation of investment returns.

* *In particular, portfolio weights must be statistically independent of contemporaneous returns. Even in the case of passive market-cap-weighted indexes, the weights will typically be correlated with contemporaneous asset returns.*

5619 DTC Parkway, Suite 600
Greenwood Village, CO 80111
Phone: +1 303-770-3377
www.investmentsandwealth.org