www.joim.com

Journal Of Investment Management, Vol. 23, No. 2, (2025), pp. 46–60

THE UNREASONABLE EFFECTIVENESS OF PORTFOLIO THEORY IN THEORY AND PRACTICE*

Andrew W. Lo[a]

As one of the founding fathers of modern finance, Harry M. Markowitz changed the way financial economists think about financial markets and institutions, and transformed the practice of finance from art to science. To honor his memory, I provide three specific examples of how portfolio theory played central roles in my own research, one involving the Adaptive Markets Hypothesis and two related to practical applications in biomedicine and fusion energy. These examples convinced me of the "unreasonable effectiveness"—to borrow a phrase from the great physicist Eugene Wigner—of portfolio theory in theory and practice.

1 Introduction

I want to start by thanking Ken Blay for the very generous introduction and for Gifford Fong for putting together this extraordinary gathering to honor Harry Markowitz. It's a great honor for me to be included in this program of financial giants, both academics and practitioners. I've been motivated, encouraged, and inspired by all of you, and it's been a wonderful privilege to be a part of this family at JOIM for the last 20 years.

It's difficult to follow all the wonderful presentations that we've seen over the course of the last day and a half. In fact, it was an intimidating task preparing for this conference because of all the luminaries that would be speaking before me. What can I say that will be additive to what has already been said?

So, let me start by observing that finance is a surprisingly young field. In much the same way that middle age first becomes real for each of us when we lose one of our parents, we're now reaching the middle age of our field now that we just lost the founding father of modern finance. Of course, there are other founding fathers of finance theory and practice who are still with us—we saw

*This article is an edited version of a presentation given at the 2024 Spring JOIM Conference in Celebration of Harry M. Markowitz, March 24–26, at the Rady School of Management, UCSD, San Diego, CA. I thank Neil Kumar and Dennis Whyte for participating in this session and adding their remarks, and Jayna Cummings for editorial support.

[a]MIT Sloan School of Management, Cambridge, MA, USA; MIT Laboratory for Financial Engineering, Cambridge, MA, USA; MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, USA; Santa Fe Institute, Santa Fe, NM, USA.

them yesterday, and a number of them are still in the audience today. It tells us how young this profession really is. So I'm going to use this opportunity to thank those people who have influenced me, not only Harry Markowitz—obviously he'll figure quite prominently—but also the others in the audience today as well as those who spoke yesterday. Thank you all for giving birth to a field that has given me and many others meaning and purpose to our lives.

My specific connection to Harry came quite late, and was due primarily to my co-author, Steve Foerster. He and I published a book a few years ago titled *In Pursuit of the Perfect Portfolio*. Steve was my first Ph.D. student at the Wharton School and is now a finance professor at Western University in Canada. Both of us teach investments, and at some point, we came up with the same question: "What would the founding fathers of modern finance choose as their perfect portfolio?" We decided to ask. Over the course of several years, we interviewed a number of these founding fathers—many of whom were here yesterday and today—including Harry, and these interviews eventually turned into a book.

This book was originally meant for our MBA students, but I have to admit, the reason we took on the project was hero worship. We wanted an opportunity to meet these giants, find out from them firsthand what their story was, how they got into finance, and how they ended up doing the extraordinary things they did.

Harry was the first person we interviewed. Afterwards, we went out to dinner with Harry (Figure 1). He chose his favorite Chinese restaurant in San Diego, Jasmine Seafood, and before Steve and I had finished looking at the menu, Harry had ordered the entire meal, something like 10 different dishes for the three of us and his wife Barbara. He turned and smiled at us, and said, "You know, I want to diversify." It was an incredible meal, and an even more incredible conversation. Our dinner lasted three or four hours, an extraordinary event that I'll never forget it.

Figure 1 From left to right, Stephen Foerster, Barbara and Harry Markowitz, Andrew Lo, taken 5 January 2013 at Jasmine Seafood Restaurant, San Diego, CA.

(*Source*: Andrew Lo).

I realized at that moment that Harry was, and is, such a larger-than-life figure—influencing so many people and institutions—that we often lose sight of the humanity that characterized him as an individual. I think the last day and a half of this conference really captured his gentle spirit and warmth, highlighting the true magnitude of our loss.

Now let me turn to the topic of my presentation, "The Unreasonable Effectiveness of Portfolio Theory in Theory and Practice," a title borrowed from the great physicist Eugene Wigner. (We all have physics envy in finance, and a number of physicists in the audience can confirm this claim.) Wigner, a Nobel Prize-winning physicist responsible for our understanding of the structure of atomic nuclei among other things, published "The Unreasonable Effectiveness of Mathematics in the Natural Sciences" in 1960. In that publication, he wrote that "the mathematical formulation of the physicist's often crude experience leads in an uncanny number of cases to an amazingly accurate description of a large class of phenomena" (Wigner, 1960, p. 8). He backed up this claim with three examples: Newton's laws of gravity, quantum mechanics, and quantum electrodynamics. This publication generated an enormous amount of discussion in math, physics, and the philosophy of science, and became one of his most cited papers.

So I thought I would try to do the same for portfolio theory. In my view, portfolio theory is everywhere. Once you learn and understand it, you can't help but see it in virtually every field and context, even if in somewhat different guises.

Looking back on my own career, I realize that there are so many instances where I use portfolio theory, it's hard to come up with just a few examples to illustrate the point. But, like Wigner, I focus on three examples—one in theory and two in practice—where portfolio theory lays at the heart of my contributions.

2 The Unreasonable Effectiveness of Portfolio Theory in Theory

In studying the decades-old debate between the Efficient Markets Hypothesis (EMH) (Fama, 1970; Samuelson, 1965) and the various behavioral anomalies that contradict this influential theory—including my own work on rejecting the Random Walk Hypothesis (Lo and MacKinlay, 1988, 1990, 1999)—I put forward an alternative called the Adaptive Markets Hypothesis (AMH) (Farmer and Lo, 1999; Lo, 2004, 2005, 2012, 2017; Lo and Zhang, 2024). The AMH provides a theoretical framework for reconciling the EMH and behavioral finance by showing that the EMH is not wrong but merely incomplete. To use language from ecology and evolutionary biology, under certain environmental conditions and given certain species in the ecosystem, the EMH holds. But if those conditions change, then behaviors may adapt to those changes and, in doing so, can violate the EMH. Ultimately, it is the interactions between environmental shifts and our subsequent adaptations that give us the financial dynamics that, from time to time, violate the EMH. The evolutionary biologist Theodosius Dobzhansky once said, "Nothing makes sense in biology except in the light of evolution" (Dobzhansky, 1973). I would argue that nothing makes sense in financial markets except in the light of the AMH.

So, what does all this have to do with portfolio theory? Let me explain by way of a behavioral anomaly that comes out of the behavioral literature from the 1950s (Grant *et al.*, 1951; Hake and Hyman, 1953; Herrnstein, 1961, 1970, 1997; Vulcan, 2000), an old example called "probability matching." Imagine playing the following coin-tossing game with an opponent. The opponent tosses a potentially biased coin repeatedly and your task is to guess the outcome—heads or tails—before each toss. A correct guess yields a $1 payment to you and an incorrect guess requires you to pay $1 to your opponent. After 20 or 30 tosses, you notice that the relative frequency of heads is 75%. With this observation in hand, what should you do to maximize your expected future winnings? The answer is simple: guess heads for every toss.

Experimental psychologists and behavioral economists ran this exact experiment with human subjects and discovered a surprising result (Grant *et al.*, 1951; Hake and Hyman, 1953; Herrnstein, 1961, 1970, 1997; Vulcan, 2000): people do not do this. Instead, experimental subjects seemed to randomize, but in a very specific way: they guessed heads 75% of the time, and they picked tails 25% of the time. In other words, they engaged in "probability matching" behavior.

This anomaly is so persistent that if, in the middle of this game, the opponent surreptitiously switches coins to one biased 60% in favor of heads, after multiple rounds with this new coin, most subjects change the relative frequency of their guesses to 60% heads as well.

What is even more puzzling is that this kind of behavior has been identified in other species, including ants (Deneubourg *et al.*, 1987; Pasteels *et al.*, 1987; Kirman, 1993; Hölldobler and Wilson, 1990), bees (Harder and Real, 1987; Thuijsman *et al.*, 1995; Keasar *et al.*, 2002), fish (Bitterman *et al.*, 1958; Behrend and Bitterman, 1961), pigeons (Graf *et al.*, 1964; Young, 1981), and primates (Woolverton and Rowlett, 1998). It seems that any animal capable of choosing between two actions, A and B, can be shown to engage in probability matching in one form or another. How irrational!

But the fact that probability matching seems to be found across so many species and over eons of evolutionary change suggests that it may not be irrational at all—there must be a good reason that such behavior exists and persists. This was the hypothesis that my co-author Tom Brennan and I conjectured when we first encountered this phenomenon and decided to try to explain it from first principles. And we found an explanation (Brennan and Lo, 2011).

Rather than describe the formal mathematical model, which can be found in Brennan and Lo (2011), let me describe a simple example from that publication, which captures the basic idea. Imagine a hypothetical ecosystem in which there are only two states of the world: rain or shine. It rains with 25% probability and shines with 75% probability.

Located in this environment are two specific regions: a valley with a beautiful river running through it, and a dry, rocky plateau on top of a mountainous region (see Figure 2). Now imagine a hypothetical individual, not necessarily human, populating this environment and it makes one important choice in its life. That choice is where to build its nest—the valley or the plateau—before having offspring.

The outcome of this important decision for the individual's offspring depends entirely on whether it rains or shines. If it builds its nest in the valley and sunshine is the state of nature, the offspring have access to plenty of water from the stream, and the valley is shaded by the canyons so they will be protected from the sun's deadly heat. Therefore, in this case three offspring will survive. But if the individual nests in the valley and it rains, the river will flood, drowning all the offspring—no survivors.

The situation is exactly the opposite for individuals nesting on the plateau. If sunshine is the state of nature, the individual's offspring will have no water or protection from the sun's heat, so no offspring will survive. But if it rains, the offspring will have access to water without the threat of drowning (excess water will simply drain into the valley), and the rain clouds will shield the offspring from the sun's heat. In this case, three offspring will survive. The payoffs of the individual's decision are summarized on the right side of Figure 2.

Now suppose we make one further assumption about this individual's behavior: whichever decision it makes regarding its nesting location, let us assume that all of its offspring make the exact same decision. In other words, if the individual chooses the valley (call this choice A), then all of its descendants will choose A; if the individual chooses the plateau (choice B), all of its descendants will choose B. And most importantly, if the individual chooses to randomize between A and B with some probability f, then all of its descendants will also randomize with the same probability f.

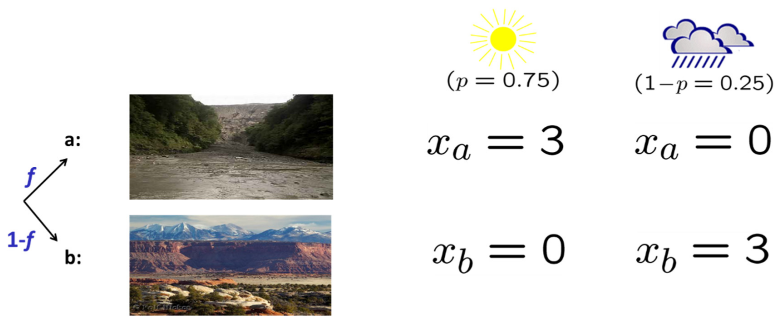

Figure 2 Illustration of Brennan and Lo (2011) binary choice model of probability matching. Individuals in this example choose to nest either in the valley (choice a) or the plateau (choice b) with probability f and $1 - f$, respectively, and this choice leads to x_a or x_b offspring, respectively, where x_a and x_b are Bernoulli random variables given by the 2×2 matrix to the right.

What we have just concocted is a very simple ecosystem with multiple "species," each characterized by their nesting behavior, and we have captured the role of genetics by assuming that each individual's specific nesting behavior is passed down from parent to child (for simplicity, the transmission of this trait from one generation to the next is assumed to be perfect—there is no mutation from one generation to the next).

So, in this ecosystem, what is the individual's optimal nesting behavior? In other words, what is the economically rational choice that optimizes the number of offspring for the individual, which we denote by f^r? This is a life-or-death question for the individual's offspring and, at the population level, for the entire species.

The obvious answer is the economically rational choice for the individual to nest in the valley 100% of the time, i.e., $f^r = 1$. But this choice ignores the fact that the first time it rains, all of the offspring in the valley will be wiped out. In other words, this particular behavior of choosing the valley all the time is not sustainable—such behavior will become extinct rather quickly.

So which behavior(s) are sustainable? To find out, consider a simple simulation where we start with an equal number of individuals of each type f and then simulate multiple generations of these individuals facing random outcomes of rain/shine. The results of this simulation are summarized in Figure 3 for behaviors $f = 0.20, 0.50, 0.75, 0.90, 1$. The outcomes are clear: the behavior that maximizes the size of the population is to randomize with $f^* = 0.75$—probability matching! More formally, survival in this simple setting is equivalent to maximizing the geometric growth rate of the population, which is given by (see Brennan and Lo, 2011, Section 4.1):

$$\mu(f) \equiv \log 3 + p \log f + (1 - p) \log (1 - f) \tag{1}$$

and this function attains its maximum at $f^* = p$. Intuitively, probability matching equalizes the expected number of surviving offspring come rain or shine which, in turn, maximizes the growth rate of the population.

Although randomization is clearly the correct answer from the perspective of the population, it can still seem highly irrational from the individual's perspective—$f^r \neq f^*$. In particular, imagine a purely fictional conversation that might take place among all individuals of type $f^* = 0.75$: "We all understand that with a 75% chance of sunshine, the best bet is for us to nest in the valley, but if we all do that, our entire species will

Generation	$f = 0.20$	$f = 0.50$	$f^* = 0.75$	$f = 0.90$	$f = 1$
1	21	6	12	24	30
2	12	6	6	57	90
3	6	12	12	144	270
4	18	9	24	387	810
5	45	18	48	1,020	2,430
6	96	21	108	2,766	7,290
7	60	42	240	834	21,870
8	45	54	528	2,292	65,610
9	18	87	1,233	690	196,830
10	9	138	2,712	204	590,490
11	12	204	6,123	555	1,771,470
12	36	294	13,824	159	5,314,410
13	87	462	31,149	435	15,943,230
14	42	768	69,954	1,155	0
15	27	1,161	157,122	3,114	0
16	15	1,668	353,712	8,448	0
17	3	2,451	795,171	22,860	0
18	3	3,648	1,787,613	61,734	0
19	9	5,469	4,020,045	166,878	0
20	21	8,022	9,047,583	450,672	0
21	6	12,213	6,786,657	1,215,723	0
22	0	18,306	15,272,328	366,051	0
23	0	27,429	34,366,023	987,813	0
24	0	41,019	77,323,623	2,667,984	0
25	0	61,131	173,996,290	7,203,495	0

Figure 3 Monte Carlo simulation of binary choice model for valley/plateau ecosystem in which rain occurs with 25% probability and sunshine occurs with 75% probability. The deterministic behavior to nest in the valley ($f = 1$) is not sustainable and is eliminated from the population the first time it rains. The optimal behavior (from the perspective of the survival of the species) is to randomize the nesting choice, with $f^* = 0.75$, which matches the probability of sunshine. After only 25 generations, this behavior completely dominates the entire population.

become extinct the first time it rains. So I need 25% of you to leave the valley and make your nest on the plateau. Chances are that none of your offspring will survive, but we need you to do this anyway so as to hedge our extinction risk. Any volunteers?" Miraculously, 25% of the population does nest on the plateau and as a result, the $f^* = 0.75$ types proliferate faster than any other individual types. Before you know it, they'll be writing the histories of how they came to dominate the ecosystem. It seems that altruism pays off.

Of course, no such conversation takes place. The "behaviors" in this example are completely mindless actions of these hypothetical individuals. But the process of evolution and natural selection causes one very specific behavior, $f^* = 0.75$, to emerge and eventually take over the entire population. Even so, this simple model offers a surprising general explanation for how probability matching might have arisen in so many animal species, and why altruism exists despite the fact that it certainly does not benefit the individuals engaging in the altruistic behavior. In fact, the simplicity of the model is a strength rather than a weakness, supporting the universality of this phenomenon precisely because so few assumptions are needed to generate it.

In case you haven't noticed, this is portfolio theory in action. Diversification is what allows the behavior $f^* = 0.75$ to persist and thrive. By not putting all our eggs in one basket—having all members of our population nesting in the valley—we are able to survive the devastations of periodic rains and flourish instead of facing extinction. It might be said that "Nature abhors an undiversified bet."[1]

Now, I realize that Markowitz had no intention of applying portfolio theory to ecology and evolutionary biology. But it does apply, nonetheless, and I credit his work with giving Tom Brennan and me the inspiration to understand and fully develop this theory of how seemingly irrational behaviors cannot only emerge organically but also come to dominate the entire population under certain circumstances.

More generally, Brennan and Lo (2011) show that if we change the assumptions of the model, the resulting behaviors that emerge will also change—the agents in this model adapt to shifts in environmental conditions. In some cases, probability matching gives way to other types of non-probability-matched randomization, especially if the outcomes listed in Figure 2 change.

Most of these emergent behaviors contradict the rationality of standard economic analysis applied to an individual agent—choose the option with the highest expected number of offspring—because natural selection does not favor individual optimality; its focus is the population. This is the lesson of $f^* = 0.75$ and $f^r \neq f^*$.

There is, however, a very important special case where selection and economic rationality coincide. Brennan and Lo (2011) consider a slightly different ecosystem in which each individual faces its own "microclimate" of rain or sunshine. This means that the environment experienced by one individual is statistically independent of the environment experienced by another individual in the same generation. Even though all agents still have a 75% chance of sunshine and a 25% chance of rain, the weather experienced by one individual may be completely different from that experienced by the other on any given day because they are assumed to live in different microclimates.

This one seemingly minor change to our example changes everything, as Figure 4 illustrates. In this Monte Carlo simulation, the exact same starting point with the exact same set of behaviors yields a completely different outcome. In this case, the economically rational individual behavior of always choosing to nest in the valley, $f^* = 1$, dominates the population. The reason is simple: because each agent lives in its own microclimate, the chances that it will rain for all individuals that choose the valley become vanishingly small as the population grows.[2]

The assumption of individual microclimates is tantamount to assuming that reproductive risk in this ecosystem is—using an idea and terminology I first learned from Sharpe (1964)—idiosyncratic, not systematic as in the original example simulated in Figure 3. As a result, in this ecosystem, the individually optimal behavior is dominant, $f^r = f^*$. Under these environmental conditions, selfishness wins out over altruism.

To bring this back to the EMH debate, our example shows that under certain circumstances, rational behavior is indeed observed in the ecosystem, but when those circumstances change (i.e., systematic instead of idiosyncratic reproductive risk), seemingly irrational behavior—probability matching—emerges and can persist. The key

Generation	$f = 0.20$	$f = 0.50$	$f = 0.75$	$f = 0.90$	$f^* = 1$
1	12	9	18	27	27
2	6	15	42	72	54
3	3	27	87	177	120
4	6	45	168	357	270
5	3	60	300	717	588
6	3	84	591	1,488	1,329
7	0	141	1,074	3,174	2,955
8	0	207	2,007	6,669	6,555
9	0	315	3,759	14,241	14,748
10	0	492	7,152	29,733	33,060
11	0	705	13,398	62,214	74,559
12	0	1,053	25,071	130,317	167,703
13	0	1,635	46,623	273,834	377,037
14	0	2,427	87,333	575,001	849,051
15	0	3,663	163,092	1,206,849	1,910,031
16	0	5,433	305,091	2,536,023	4,296,213
17	0	8,148	570,852	5,325,852	9,666,762
18	0	12,264	1,069,884	11,188,509	21,755,844
19	0	18,453	2,007,642	23,494,611	48,959,286
20	0	27,711	3,763,281	49,346,967	110,148,060

Figure 4 Monte Carlo simulation of binary choice model for valley/plateau ecosystem in which rain occurs with 25% probability and sunshine occurs with 75% probability, but where the ecosystem consists of microclimates so that reproductive risk is idiosyncratic. In this case, deterministic behavior to nest in the valley ($f = 1$) is sustainable and coincides with the individually optimal behavior. After only 20 generations, this behavior completely dominates all other behaviors, including probability matching.

question is not whether individuals are rational or irrational but what are the environmental conditions that give rise to their behavior?

In a series of publications summarized in Lo and Zhang (2024), my co-authors Tom Brennan and Ruixun Zhang and I provide a comprehensive and rigorous treatment of the mathematical and statistical foundations of the AMH. We provide evolutionary explanations for a number of human behaviors in addition to probability matching including risk aversion, loss aversion, bias and discrimination, group selection, and other tendencies that, in isolation, may seem economically irrational. In fact, these puzzling behaviors are almost always manifestations of Nature optimizing her portfolio according to principles first described by Harry Markowitz, Bill Sharpe, and the other founding fathers of modern finance. Despite the fact that none of them had ecology, evolutionary biology, and the AMH in mind when they made their contributions, the sheer power of their ideas explains the unreasonable effectiveness of portfolio theory in a wide range of theories across multiple disciplines.

3 The Unreasonable Effectiveness of Portfolio Theory in Practice

Let me now turn to the second and third examples of the unreasonable effectiveness of portfolio theory, both involving practice.

The first example has to do with the financing of drug development programs to treat rare diseases. According to the US Orphan Drug Act of 1983, a rare disease is defined as any disease that affects fewer than 200,000 patients in the country. Examples include cystic fibrosis, hemophilia, sickle cell disease, Duchenne muscular dystrophy, and over 10,000 other diseases, many of which are caused by rare genetic mutations. Although each disease may impact a relatively small patient population, over 30 million Americans are affected by one of these conditions, so this is not a small problem.

For concreteness, consider Canavan disease, a fatal neurological condition caused by a mutation in a single gene responsible for maintaining white matter in the brain, which is critical for the proper transmission of nerve signals. Patients affected by this mutation—mostly infants and young children—gradually lose their white matter, causing a variety of symptoms including the inability to crawl, walk, sit, talk, and control their muscles. Children with Canavan may suffer seizures, have trouble swallowing, and may become blind, deaf, and paralyzed. They rarely live past the age of 10. There are currently no effective approved therapies for Canavan, so the only treatments available are those addressing the symptoms.

The good news is that physician–scientists have recently developed a number of remarkable ways to treat and, in some cases, cure these terrible afflictions. One such technique is known as "gene therapy," which involves the radical approach of injecting patients with a virus containing the correct version of their mutated gene. The virus replaces the defective gene with the correct one, which then gets re-copied correctly from that point forward by the treated cells, eventually restoring the function that was lost due to the mutant gene. Like magic, the patient is cured.

The not-so-good news is that, because of the smaller patient population affected by a typical rare disease, big pharmaceutical companies are reluctant to undertake the costly process of developing and testing drugs for these conditions. To maximize shareholder value, they tend to focus on potential "blockbuster" products—defined as drugs generating annual sales of $1 billion or more—like therapies for diabetes, heart disease, and obesity.

However, this lack of financial incentive can be addressed by portfolio theory: combining multiple rare-disease drug development programs into a single financial entity. In addition to increasing the potential sales to blockbuster status in aggregate, this approach can also reduce the risk of failure. In the biopharma industry, the phrase "multiple shots on goal" is often used to describe the advantages of large pharmaceutical companies, but this soccer or hockey analogy can also be applied to developing a portfolio of rare-disease therapies.

What makes the application of portfolio theory so compelling in this case is the observation that because rare diseases can be so different from one another, the success or failure of one rare-disease program is generally uncorrelated with that of another program. In other words, it is possible to construct a portfolio of rare-disease drug programs that are pairwise uncorrelated. A diagonal covariance matrix is something rarely seen in practical portfolio optimization problems, and to underscore this point, I sometimes challenge my MBA students to name five financial investments—any assets are fair game, including stocks, bonds, currencies, commodities, venture capital, private equity, and all types of hedge fund strategies—that are pairwise uncorrelated. The most thoughtful answers usually involve two or, at most, three assets (e.g., gold, equities, and cryptocurrencies), but over recent periods, even these examples no longer work. It is an empirical fact that over the past decade, it is impossible to name any nontrivial[3] collection of five pairwise-uncorrelated financial investments. However, I can easily name 20 rare-disease drug development programs that have nothing to do with each other, hence their pairwise correlations are negligible, if not zero.

This implies that the return volatility of such a portfolio will decline with the number of programs, n, at a rate of $\sqrt{n}$, which, in turn, implies that the Sharpe ratio increases at that same rate. For example, if the Sharpe ratio for a single drug development program has an expected excess return of 15% above the risk-free rate, but an unattractively high volatility of 75%, the Sharpe ratio is $\frac{15\%}{75\%} = 0.20$, but a portfolio of 5, 10, and 25 identical, uncorrelated programs have Sharpe ratios of 0.45, 0.63, and 1.00, respectively. For comparison, the Sharpe ratio of the S&P 500 from January 1926 to July 2024 is 0.30.[4]

A simple application of Markowitz mean-variance analysis shows that a portfolio approach to developing treatments for rare diseases can be highly attractive to investors at sufficient scale. In a collaboration with researchers at the National Institutes of Health's rare disease division,[5] I published a simulation of a rare disease portfolio that generated a hypothetical expected return of 22% with moderate risk (Fagnan *et al.,* 2015), a surprising result that exceeded the performance of most hedge funds during the same period.

Prior to publication, I presented these results at a healthcare finance conference I co-organized at MIT called CanceRx,[6] and it caught the attention of a former MIT student of mine, Neil Kumar, a chemical engineering Ph.D. who attended my introductory finance course years ago. He approached me after the conference and asked whether I could re-run the analysis with his own assumptions, and over the next few months, we formulated a more realistic simulation that incorporated sharper estimates of the programs' probabilities of success, development costs, revenue forecasts, and other parameters.

Those results formed the basis of a rare-disease portfolio company that Neil, several scientists, and I co-founded called BridgeBio Pharma (Kumar *et al.*, 2024). Within the space of 9 years, this company has raised and deployed several billion in capital, and it currently has a market value in excess of $6 billion.[7] However, the company is most proud of the fact that they have had three drugs approved and several more in the pipeline, one of which is a gene therapy that can cure Canavan disease. Patient number 3 in BridgeBio's clinical trial for this treatment is Noa Greenwood.[8] In June 2022, Noa received the experimental treatment and weeks later, started improving noticeably. Two and a half years later, she is walking, talking, riding a bicycle, and doing everything a healthy toddler should be doing. One of the greatest honors I have ever had—both professionally and personally—was meeting Noa in March 2023. I will never forget my first sight of her, walking up the staircase in a local coffee shop holding her dad's hand. The phrase "the blind shall walk, the lame shall see" is usually associated with religious experiences. These experiences are happening today, thanks to the miracle of modern medicine. And portfolio theory.[9]

I want to emphasize that I had nothing to do with any of the science or medicine behind this and other BridgeBio drugs—this is entirely the domain of Neil and the extraordinary BridgeBio scientific and clinical operations team. Finance can never turn a bad drug into a good drug. However, the lack of finance, or the wrong kind of finance, can definitely ensure that a good drug never reaches a patient. BridgeBio's financing strategy and its impact on patients is yet another aspect of Harry Markowitz's remarkable legacy.

The final example of the unreasonable effectiveness of portfolio theory in practice involves energy transition. Most of us agree that climate change is real and to address it, we need to reduce our reliance on fossil fuels. But the term "transition" begs the question of what we are transitioning *to*? In fact, the most recent data regarding global energy consumption shows that, despite all the progress we have made in renewable energy such as wind and solar, we have only decreased fossil fuel consumption as a percentage of primary energy usage by 0.4% to 81.5% (Energy Institute, 2024, p. 4). To reduce our use of oil and gas, we need other forms of energy—and vast amounts of it—that do not emit greenhouse gases.

I have been told by energy experts at MIT and elsewhere that we have three options in addition to wind, solar, and hydroelectric power: geothermal energy, nuclear fission, and nuclear fusion. For reasons that I do not have time to cover here, the most politically feasible choice at this time seems to be fusion energy.

You may recall that in contrast to fission—which involves splitting atoms—fusion energy comes from *combining* atoms at extremely high temperatures to form heavier elements,[10] releasing enormous amounts of energy in the process. This is the same process that powers the Sun and all other stars, hence fusion energy is the most common and most abundant source of energy in the universe. Moreover, it can, in theory, unlock the most amount of energy per gram of fuel compared to all other forms of energy (according to Einstein's $E = mc^2$). And the primary fuel source—seawater—is largely free and, for all practical purposes, limitless! Moreover, unlike fission reactors, fusion power plants could theoretically do all this with minimal radioactive waste.

So what's the catch? The catch is that a number of technological hurdles must be overcome before a commercially viable fusion power plant (Figure 5) becomes reality: heating the plasma to a sufficiently high temperature for fusion to occur (around 100 million degrees Celsius), containing the plasma with a sufficiently strong magnetic field (around 20 T or more), housing this magnetic bottle of plasma in a larger structure safe for human operators, converting the fusion energy released from this bottle of plasma to electricity, etc. A number of these hurdles involve creating first-of-a-kind (FOAK) technologies that have yet to be invented and built, hence the risks are unusually high. But the rewards are even greater, not just to investors but to all of humanity.

The pace of progress seems to be picking up. In September 2021, MIT and its spinoff,

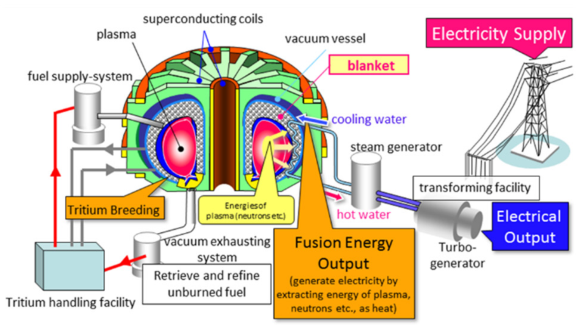

Key Components:

- Sustained high-temp plasma
- Confinement
- Materials
- Tritium breeding
- Energy extraction and transfer
- Stability and control
- Scalability
- Fuel cycle and waste management
- Economic viability
- Regulation and licensing

Figure 5 Diagram of key components of a nuclear fusion power plant.

(*Source*: https://www.qst.go.jp/site/qst-english/, accessed 23 January 2025, and author's analysis).

Commonwealth Fusion Systems, reached a key fusion milestone by switching on a powerful superconducting magnet that generated a sustained magnetic field of 20 T,[11] a key prerequisite for containing matter undergoing fusion reactions. In December 2022, researchers at the US Department of Energy's National Ignition Facility (NIF) imploded a miniscule pellet containing deuterium and tritium fuel using 2.05 megajoules of energy delivered via simultaneous pulses from 192 lasers, generating an energy output of 3.15 megajoules. This is the first time net energy gain has been achieved in a controlled fusion reaction on Earth, and the experiment has since been repeated several times by NIF, in each case with greater energy gain. And in January 2025, China's Experimental Advanced Superconducting Tokamak achieved a remarkable milestone by maintaining plasma at temperatures exceeding 100 million degrees Celsius for 1,066 s. This achievement not only set a new world record but also demonstrated the feasibility of sustaining the extreme temperatures necessary for fusion reactions over periods lasting more than split seconds.

So how are we to clear the remaining hurdles? Not surprisingly, financing is a key issue. In fact, one of the world's leading fusion experts—my MIT colleague Dr. Dennis Whyte—has observed that the biggest risk to achieving ignition is not the science or engineering, but the financing.[12] Without sufficient financing, science and engineering cannot progress. But given the outsized risks associated with FOAK technologies, how can investors be convinced to provide capital to this endeavor, especially with competing opportunities like new solar and wind energy technologies, hyper-Moore's-Law GPU chip technologies, generative AI, and so on?

One answer is provided by Markowitz: think of an investment in fusion not as a single bet, but rather as a portfolio of bets on multiple technologies (see the list of key components in Figure 5), most of which will have separate and independent value in their own right. For example, the superconducting magnets developed by Commonwealth Fusion Systems can greatly improve the efficiency of electric motors, magnetic resonance imaging medical devices, quantum computers, and magnetic levitation trains. When viewed as a portfolio of well-diversified intellectual property—and if fusion entrepreneurs are willing and able to deploy these properties in spinout subsidiaries—investors will respond to the reduced risk and greater return potential by deploying more capital.

Inspired by Markowitz's approach, Dennis and I and our students describe how to apply this portfolio approach to financing fusion energy in Alhamdan *et al*. (2023) through the same "megafund" structure that was used in creating BridgeBio Pharma. The analogy between fusion and the biotech industry is no accident—we argue in Lo and Whyte (2024a,b) that the fusion industry today is roughly where the biotech industry was in the mid-1970's.

But rather than taking 50 years to reach the same level of success as today's biotech industry, the fusion industry can accelerate its progress by adopting the lessons that biotech learned through trial and error, one of which is the importance of risk reduction via portfolio diversification. As if to underscore this point, 3 months after the demonstration of its 20-T magnet, which greatly reduced the risk to the company's investors, Commonwealth Fusion Systems raised an additional $1.8 billion of private-sector investor capital and is now well on their way to creating the first proof-of-concept tokamak called SPARC in Devens, Massachusetts (Figure 6).

Imagine the impact that a multibillion-dollar fusion megafund could have on society. The investment managers in this audience and the

Figure 6 Commonwealth fusion systems SPARC facility, Devens, Massachusetts.

assets they are currently managing could bring us to the next level of human evolution by harnessing the power of the Sun here on Earth.

4 Conclusion

I hope these three examples illustrate just how wide Harry Markowitz's impact has been and will continue to be for the foreseeable future. He has been an important inspiration to me, to most everyone in this audience, and to the rest of my colleagues in academic finance and in the financial industry. Although we will greatly miss his larger-than-life presence, he lives on in our memories and in the unreasonable effectiveness of portfolio theory in many disciplines, including financial theory and practice.

I'd like to conclude by relating a scene from one of my all-time favorite movies, *Annie Hall*, a love story by Woody Allen. Allen plays a character named Alvy Singer and in one scene, Alvy is waiting in line at a movie theater with the eponymous star played by Diane Keaton. Behind them is an obnoxious pseudo-intellectual who is trying to impress his date by pontificating about Marshall McLuhan's theory on media and society. Alvy can't stand it because he understands the theory and the stranger is just bastardizing it, so he fantasizes that McLuhan—who plays himself in the movie—magically emerges next to him and reprimands the pseudo-intellectual, saying "You know nothing of my work...."

As I pontificate about biotech and fusion, I have concerns that someone in this audience will fantasize about their own Marshall McLuhan moment. So I have decided to pre-empt that fantasy with my own Marshall McLuhan moment by inviting Neil Kumar and Dennis Whyte to join me on the stage and tell you firsthand about the exciting work they are doing. Please join me in welcoming Neil and Dennis.

5 Closing Remarks

Let me wrap up by pointing out, first of all, how we got here. We can thank Markowitz, Sharpe, Merton, Scholes, Engle, and all of the other giants of finance on whose shoulders the rest of us stand. The financial models and methods we learned at the start of our careers actually work. Those of you who are on the front lines of practice know this from experience, but there are an even greater number of people in other fields who are not familiar with ours and desperately need financing, so greater collaboration between financial experts and the rest of the world is really critical.

But can we afford all this financing, across all the myriad challenges facing society today—cancer, Alzheimer's, diabetes, new infectious diseases, climate change, pollution, poverty, and so on? I believe this is the wrong question. A better one is "can we afford not to try?"

These challenges are real, and finance can help.

In fact, this very audience holds managers with enough assets as well as the expertise—just between Neil and Dennis—to be able to solve a number of these challenges.

With the proper financial structure, I believe anything is possible. There are tremendous opportunities that we can take advantage of, and by doing so, we can do well on behalf of investors by doing good. This may be the most important legacy that financial economists and practitioners leave to society. And the best way to honor the memory and spirit of Harry M. Markowitz.

Notes

[1] See Real (1980) and Real and Caraco (1986) for examples of how ecologists incorporate risk and diversification explicitly in models of animal foraging behavior.

[2] The assumption of statistical independence of the state of nature across individuals implies that the probability of rain for all m residents of the valley is $(1-p)^m$, which converges to 0 geometrically fast as m increases without bound.

[3] This qualifier rules out degenerate examples like long 100 shares of X, short 100 shares of X, etc. because those cases yield zero expected returns.

[4] The Ibbotson SBBI US Large-Cap Stocks geometric compound return and standard deviation are 10.4% and 18.6%, respectively, and the geometric compound return of the Ibbotson SBBI Intermediate-Term US Treasury Return is 4.8%.

[5] The National Center for Advancing Translational Sciences, https://ncats.nih.gov/research/our-impact/our-impact-rare-diseases (accessed 23 January 2025).

[6] https://lfe.mit.edu/events/cancerx-2013/ (accessed 25 January 2025).

[7] As of 23 January 2025.

[8] https://ntsad.org/support-for-families/meet-our-families/noa-greenwood-ntsad-impact-story/ (accessed 25 January 2025).

[9] See Neil Kumar's remarks also included in this volume for more on the intersection of drug development and finance.

[10] A common type of fusion is between deuterium and tritium, which are hydrogen atoms containing one and two neutrons, respectively, in their nuclei. Deuterium is plentiful and can found in seawater. Tritium is considerably rarer but can be created as a by-product of fusion itself, yielding a sustainable fuel generation cycle.

[11] For comparison, the strength of the Earth's magnetic field at the surface ranges between 25 and 65 μT, where a microtesla is one millionth of a Tesla.

[12] See his comments in this volume.

References

Alhamdan, A., Halem, Z. M., Hernandez, I., Lo, A. W., Singh, M., and Whyte D. (2023). "Financing Fusion Energy," *Journal of Investment Management* **21**(1), 5–51.

Behrend, E. R. and Bitterman, M. E. (1961). "Probability-Matching in the Fish," *The American Journal of Psychology* **74**(4), 542–551.

Bitterman, M. E., Wodinsky, J., and Candland, D. K. (1958). "Some Comparative Psychology," *The American Journal of Psychology* **71**, 94–110.

Brennan, T. J. and Lo, A. W. (2011). "The Origin of Behavior," *Quarterly Journal of Finance* **1**(1), 55–108.

Deneubourg, J. L., Aron, S., Goss, S., and Pasteels, J. M. (1987). "Error, Communication and Learning in Ant Societies," *European Journal of Operational Research* **30**(2), 168–172.

Dobzhansky, T. (1973). "Nothing in Biology Makes Sense Except in the Light of Evolution," *The American Biology Teacher* **35**(3), 125–129.

Energy Institute (2024). *Statistical Review of World Energy*. Energy Institute, London.

Fagnan, D. E., Yang, N. N., McKew, J. C., and Lo, A. W. (2015). "Financing Translation: Analysis of the NCATS Rare-diseases Portfolio," *Science Translational Medicine* **7**(276), 276ps3.

Fama, E. (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* **25**(2), 383–417.

Farmer, J. D. and Lo, A. W. (1999). "Frontiers of Finance: Evolution and Efficient Markets," *Proceedings of the National Academy of Sciences of United States of America* **96**, 9991–9992.

Graf, V., Bullock, D. H., and Bitterman, M. E. (1964). "Further Experiments on Probability-Matching in the Pigeon," *Journal of the Experimental Analysis of Behavior* **7**, 151–157.

Grant, D. A., Hake, H. W., and Hornseth, J. P. (1951). "Acquisition and Extinction of Verbal Conditioned Responses with Differing Percentages of Reinforcement," *Journal of Experimental Psychology* **42**(1), 1–5.

Hake, H. W. and Hyman, R. (1953). "Perceptions of the Statistical Structure of a Random Series of Binary Symbols," *Journal of Experimental Psychology* **45**, 64–74.

Harder, L. D. and Real, L. A. (1987). "Why Are Bumble Bees Risk Averse?," *Ecology* **68**, 1104–1108.

Herrnstein, R. J. (1961). "Relative and Absolute Strength of Responses as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Behavior* **4**, 267–272.

Herrnstein, R. J. (1970). "On the Law of Effect," *Journal of the Experimental Analysis of Behavior* **13**, 243–266.

Herrnstein, R. J. (1997). *The Matching Law* (Cambridge, MA: Harvard University Press).

Hölldobler, B. and Wilson, E. O. (1990). *The Ants*. (Cambridge, MA: Belknap Press).

Keasar, T., Rashkovich, E., Cohen, D., and Shmida, A. (2002). "Bees in Two-armed Bandit Situations: Foraging Choices and Possible Decision Mechanisms," *Behavioral Ecology* **13**, 757–765.

Kirman, A. (1993). "Ants, Rationality, and Recruitment," *Quarterly Journal of Economics* **108**(1), 137–156.

Kumar, N., Lo, A. W., Shukla, C., and Stephenson, B. (2024). "Applications of Portfolio Theory to Accelerating Biomedical Innovation," *Journal of Portfolio Management* **51**, 213–236.

Lo, A. W. (2004). "The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective," *Journal of Portfolio Management* **30**(5), 15–29.

Lo, A. W. (2005). "Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis," *Journal of Investment Consulting* **7**(2), 21–44.

Lo, A. W. (2012). "Adaptive Markets and the New World Order," *Financial Analysts Journal* **68**(2), 18–29.

Lo, A. W. (2017). *Adaptive Markets: Financial Evolution at the Speed of Thought* (Princeton, NJ: Princeton University Press).

Lo, A. W. and MacKinlay, A. C. (1988). "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies* **1**, 41–66.

Lo, A. W. and MacKinlay, A. C. (1990). "When are Contrarian Profits Due to Stock Market Overreaction?," *Review of Financial Studies* **3**, 175–206.

Lo, A. W. and MacKinlay, A. C. (1999). *A Non-Random Walk Down Wall Street* (Princeton, NJ: Princeton University Press).

Lo, A. W. and Whyte, D. G. (2024a). "What Can Fusion Energy Learn from Biotechnology?," *Issues in Science and Technology* **40**, 68–70.

Lo, A. W. and Whyte, D. (2024b). "What Fusion Energy Can Learn From Biotechnology?," Available at SSRN: https://ssrn.com/abstract=4779516 or http://dx.doi.org/10.2139/ssrn.4779516.

Lo, A. W. and Zhang, R. (2024). *The Adaptive Markets Hypothesis: An Evolutionary Approach to Understanding Financial System Dynamics*, Clarendon Lectures in Finance (Oxford University Press).

Pasteels, J. M., Deneubourg, J. L., and Goss, S. (1987). "Self-organization Mechanisms in Ant Societies. I: Trail Recruitment to Newly Discovered Food Sources," *Experientia*. Supplementum.

Real, L. (1980). "Fitness, Uncertainty, and the Role of Diversification in Evolution and Behavior," *American Naturalist* **115**, 623–638.

Real, L. and Caraco, T. (1986). "Risk and Foraging in Stochastic Environments," *Annual Review of Ecology and Systematics* **17**, 371–390.

Samuelson, P. (1965). "Proof That Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review* **6**(Spring), 41–49.

Sharpe, W. F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance* **19**(3), 425–442.

Thuijsman, F., Peleg, B., Amitai, M., and Shmida, A. (1995). "Automata, Matching and Foraging Behavior of Bees," *Journal of Theoretical Biology* **175**(3), 305–316.

Vulcan, N. (2000). "An Economist's Perspective on Probability Matching," *Journal of Economic Surveys* **14**, 101–118.

Wigner, E. P. (1960). "The Unreasonable Effectiveness of Mathematics in the Natural Sciences. Richard Courant Lecture in Mathematical Sciences Delivered at New York University, May 11, 1959," *Communications on Pure and Applied Mathematics* **13**(1), 1–14.

Woolverton, W. L. and Rowlett, J. K. (1998). "Choice Maintained by Cocaine or Food in Monkeys: Effects of Varying Probability of Reinforcement," *Psychopharmacology* **138**, 102–106.

Young, J. S. (1981). "Discrete-Trial Choice in Pigeons: Effects of Reinforcer Magnitude," *Journal of the Experimental Analysis of Behavior* **35**, 23–29.

Keywords: Harry M. Markowitz; modern portfolio theory; biomedical innovation; fusion energy