

TRIAL and ERROR

ANDREW LO ON WHY IT TAKES A THEORY TO BEAT A THEORY

BY CHRISTINA GROTHEER

A natural born skeptic, Andrew Lo takes nothing at face value. So it should come as no surprise that regarding a polarizing topic such as the efficient market hypothesis (EMH) — with many academics in favor but many practitioners opposed — Lo has engineered his own alternative view. “My interpretation is one that I think integrates and resolves any conflicts between the critics of the EMH and its proponents,” says Lo, Harris & Harris Group Professor of Finance at MIT’s Sloan School of Management and director of MIT’s Laboratory for Financial Engineering. “It is what I call the biological view of efficient markets.”

Lo has done psychological surveys of online day traders, monitored the physiological profiles of professional traders during live trading sessions, and soon plans to do MRI scans of traders’ brains. Seem a bit far-fetched?

For Lo, crafting a clever theory is not enough; his ideas must stand up to extensive field tests. Says he, “Any theory is just that — it’s a theory. And all economic theories are incorrect; they’re all approximations to a very complex reality.”

The efficient market hypothesis elicits strong reactions from people, whether for or against. I’m wondering where you fall along the spectrum.

I think that I probably fall *outside* the spectrum in the sense that I think the efficient market hypothesis has been largely misinterpreted and misunderstood. And that’s one of the reasons why there’s such controversy between academia and the industry. I’d like to propose a third interpretation — what I call the “biological” view of efficient markets. Under this new interpretation, an efficient market is one that is highly competitive and highly adaptive — nothing more, nothing less.

Efficient markets doesn’t imply that it’s impossible to generate profitable trading strategies, but it does imply that it may be very difficult. The only way to make money on a consistent basis is to evolve and adapt to changing market conditions — in other words, to innovate on a regular basis.


D.E. Shaw has spent more than US\$100 million developing its technological toolbox. How can quant shops make any money with the expenditure required to find and exploit inefficiencies?

This is exactly the kind of competition I was referring to. There’s no doubt that hedge funds like D.E. Shaw, Renaissance Technologies, BNP Cooper Neff, Morgan Stanley’s Process-Driven Trading group, Prediction Company, and others have made the hurdle much higher for the rest of the industry. So, without a doubt there are many forces at work in making markets more efficient. But, at the same time, new sources of inefficiencies are also being created because of the natural ebb and flow of market participants and market conditions. With as dynamic a system as the economy, there are bound to be significant changes over time in the financial landscape, and new opportunities arise even as others disappear.

Do you think technological innovations have altered the support level for the EMH?

Absolutely. I think technology plays an incredibly important role in determining the *degree* of market efficiency. For example, D.E. Shaw spending US\$100 million on its technology infrastructure is a good example of how technology can change the playing field. Suppose that you spent US\$100 million 50 years ago on infrastructure. Well, you would have cleaned up back then.

Nowadays, US\$100 million doesn’t go as far, and you’ve got several other organizations that have at least as much money and who are willing to make the same level of commitment. So, I think technology — particularly for early adopters — plays a very important role in determining the speed with which markets become efficient. And not just investors’ technology, but *market* technology as well. For example, Electronic Communications Networks have changed the landscape for US





equities simply because now buyers and sellers are able to find each other much more quickly and at a lower cost than ever before.

Many practitioners are mystified by why the EMH is so important to academics. Can you shed some light on this?

Sure. There is a very good reason why the EMH has played such a prominent role in academic circles. Prior to the 1950s and '60s, there wasn't a very clear understanding of what drives financial markets.

There were all sorts of financial theories about how to value real and financial investments, what determines a corporation's "optimal" capital structure, when companies should issue new equity, and how to forecast stock prices, but these theories were often contradictory, internally inconsistent, and certainly not very useful from a practical perspective.

Then, the Efficient Market Hypothesis came along, first proposed by Paul Samuelson and then extended by others at MIT and the University of Chicago. Like a breath of fresh air, the EMH cleared out a lot of the previous theories by arguing that market prices are actually pretty good indicators of value, precisely because of the trading of many investors attempting to profit from small informational advantages. The EMH was an incredibly succinct way of summarizing lots of different forces that were at work. It was a stunning and surprising conclusion from a rather simple set of hypotheses.

Are you still running a hedge fund? I've seen very few mentions in the press.

Yes, my partners and I at AlphaSimplex Group have been managing money exclusively for Paloma Partners during the past four years. We'll be launching some new products over the next few months that will be open to other investors as well. We prefer to stay under the radar screen and to focus on what we're good at, which is research, development, and trading.

How does your work at AlphaSimplex differ from your research at MIT?

Well, I try to keep my business activities separate from my academic research. Obviously there are some synergies between the two. For example, I have a much better appreciation now of some of the more practical aspects of financial markets than I did before starting AlphaSimplex, and a few of the challenges we face at AlphaSimplex have been the motivation for some of my academic research. But by and large, the kind of research that academia tends to be interested in is quite different from the kind of research that's relevant for practical investment management.

That's actually a good thing because otherwise there

would be some potential conflicts of interest between what I do at AlphaSimplex and what I do at the MIT Laboratory for Financial Engineering. And I'm very careful to provide full disclosure of all of my affiliations in any forum that I speak or publish in.

But, cutting-edge research in academic finance is often far removed from the more practical aspects of investment management (otherwise it wouldn't be considered "cutting edge" from the academic perspective), while much of the work we do at AlphaSimplex is a very specialized kind of financial engineering that academic journals are not likely to be interested in. Nevertheless, there's no doubt that our academic background has given us great advantages in the scope and depth of our investment activities.

I've read that you are "politely disparaging" of behavioral finance. But a lot of your research focuses on irrationality and emotion. How do you distinguish between the two?

I want to be very clear that I have a great deal of respect for the behavioral economists and psychologists who pioneered the work that is now called "behavioral economics and finance." Maybe the reason that some of my views have been characterized as "politely disparaging" is that I've been somewhat disappointed that the behavioral literature doesn't go far enough in providing useful alternatives to the current EMH paradigm. In my view, it takes a theory to beat a theory. If we're going to criticize the existing paradigms of modern finance — which have been enormously successful despite their limitations — then I think there's an implicit obligation for us to provide better alternatives. Behavioral finance has not yet offered any such alternatives.

I do think that the behavioral finance literature has documented some valuable anomalies and departures from the classical EMH. That kind of documentation I view as important "grist for the mill" that will ultimately generate a new paradigm, and frankly, one that I hope to develop over the next few years along biological and evolutionary lines.

What have you discovered from your biological approach to markets?

I now feel that I have a consistent paradigm for thinking about a lot of these anomalies as well as the original EMH — an intellectually satisfying resolution of the controversy between the two perspectives. The resolution is simply that people engage in heuristic approximations to the theoretical optimal behavior that economists predict using their neoclassical models, and that these heuristics are shaped over time through the forces of evolution: mutation, adaptation, and natural selection.

Many years ago, Herb Simon, the Nobel Prize-winning economist, came up with the notion that people don't optimize, but use "rules of thumb" instead. He dubbed these rules of thumb "satisficing," arguing that people don't actually optimize to the n^{th} degree; they have to compromise because they have limited computational abilities. But his work was roundly

criticized by most economists because he wasn't able to explain where to draw the line between optimizing and satisfying behavior: How do individuals figure out when to stop optimizing if they don't know how much better the optimal solution will be relative to a merely satisfactory one?

In my own research, I think I've found a compelling response to Simon's critics: Where you draw the line is determined randomly, based on the particular environmental conditions that exist at the time. That is, people arrive at rules of thumb through trial and error. What works, they'll continue to use. What doesn't work, they'll discard. And very quickly, the forces of natural selection, especially in the fast-paced world of financial markets, can yield nearly optimal behavior in many different contexts.

Can you give me a specific example?

Sure. The equity risk premium and asset allocation or market timing — the controversy that Peter Bernstein stirred up recently by arguing that investors should reconsider market timing.

looks upon life and finances — her risk preferences were permanently altered by her experiences during the Great Depression. By that same token, the great bull market of the '80s and '90s created a generation of investors who never lost money in the stock market, and the only lesson they ever learned was to buy on dips. That experience also must have changed their perspective and risk preferences.

The point is that the history of market conditions and economic cycles actually determines the risk preferences of individuals in the marketplace today, so of course the equity risk premium is not a universal constant like gravity or the speed of light. Gravity may not have changed a lot in the last 100 years, but financial markets have changed a great deal. And so, I think we need to be asking what the appropriate risk premium is *today* given the experiences in the marketplace over the recent past, and how it's likely to change over the next 10 years as the population changes. By looking at it from an evolutionary perspective, you end up asking different questions and ultimately coming up with different answers.

You received a lot of press when you wired traders to monitor their biological responses to the market. What were you hoping to accomplish?

One of the things that my co-author Dmitry Repin and I are attempting to do in our broader quest to develop a formal theory of behavioral finance is to determine whether or not emotion plays an important role in financial decision making. In some of this research, we've attached sensors to professional traders that measure physiological characteristics — blood pressure, skin conductance, breathing rate, and body temperature — while they trade.

What we're finding is that emotion is quite a significant factor in the decision-making process, even for the most seasoned professional traders, who, by all accounts, are among the most rational of financial-market participants. This really calls into question the dichotomy between fear and greed on the one hand and rationality on the other.

Why didn't you track what the traders' brains told them to do to see exactly how emotions affected their decisions?

That's our next phase. When you're dealing with humans as experimental subjects, the necessary experimental design and controls become extremely challenging. The question you're asking is a great one: What part of the brain gets triggered during the financial decision-making process? But to be able to answer that question, we first need to have a sense of where to look.

The purpose of our physiological experiments was first to see whether emotion plays a role at all. And those results have shed considerable light on what parts of the brain we should focus on. In our next phase, we plan to run MRI scans on traders while they trade to see whether the regions that we've identified are indeed responsible for financial-risk processing. Ultimately, we're hoping to see if there's some kind of genetic basis for financial-risk-taking skills, or at least develop some



Where I agree with Peter, and where I see the evolutionary approach as having enormous value in helping us understand markets, is that the equity risk premium — that single number that people focus on for all sorts of reasons — is not a number at all. It's actually a quantity that varies through time with market conditions, with demographics, and with all sorts of phenomena that are determined by evolutionary forces.

So how do evolutionary ideas come into play?

Through the recognition that the equity risk premium is really determined by the nature of the specific population at hand, whose risk preferences are determined by their collective past experiences in a very path-dependent way. For example, I know someone whose grandmother lived through the Great Depression. She's the matriarch of a family that's currently worth about US\$50 million dollars, and she lives in a very modest, rent-controlled apartment. Even when it's raining, she'd rather walk than pay the \$5 cab fare to her son's apartment. When her children and grandchildren criticize her, she responds, "Did you ever have to wait on line every day for your dinner, not knowing whether there was actually going to be food left when you reached the front of the line? You can always lose what you have."

This woman was indelibly changed in terms of how she

sense for whether or not certain people are predisposed to making better risk-based decisions.

Were you surprised by the response that you got from technical analysts because of the research you published three years ago?

Yes, I was very surprised, and not completely prepared for it either. I guess I didn't realize there was such a groundswell of interest in the quantitative foundations of technical analysis, and am grateful for the following that our work seems to have attracted.

Have you made any progress toward detecting new technical patterns that might be even better market indicators?

Actually, I have, but I haven't published anything yet because I'm still working out some of the analytics. It does look like there are more sophisticated patterns that contain even more information about future prices than the standard ones. These seem to be largely driven by technological advances in the marketplace. Nowadays, you have people who are trading on much shorter time scales, and, as a result, the tell-tale patterns that they leave in market prices and volume tend to be somewhat different than the more traditional technical patterns like head-and-shoulders.

Has your research on the serial correlation of hedge fund returns been published?

Yes, it's forthcoming in the *Journal of Financial Economics*.

Given your findings of extensive smoothing in hedge fund returns, how do you set yourself apart from less scrupulous hedge fund managers?

At AlphaSimplex, we disclose all of our risk exposures to our investors using an extensive set of proprietary risk models so that, at all times, they have a clear understanding of the risks that we are taking and the risks that we are not taking. The genesis of my interest in serial correlation in hedge fund returns was our internal research process for developing a set of analytics to provide our investors with the most accurate measure of our investment performance.

One of the things that most hedge fund investors want to see is the Sharpe ratio. And, it occurred to me that in virtually all of the studies and computations that I've seen of the Sharpe ratio, there is an implicit assumption that there is no serial correlation in the underlying returns of the fund. But, having looked at a variety of hedge fund databases, I knew that in

“... the equity risk premium is not a universal constant like gravity or the speed of light.”



many cases, hedge fund returns can exhibit substantial serial correlation.

At first I thought this was a minor statistical anomaly that could easily be addressed by a somewhat more sophisticated estimator, which is what I constructed and reported to our investors in addition to the traditional estimator. And they responded, “We're glad that you computed it both ways, but can you explain the more sophisticated estimator to us?” And so, little by little, I got drawn into the topic and realized that the problem of serial correlation is really quite widespread, and is not at all well understood.

How rewarding is it to see that what you're doing is making a difference?

It's extremely exciting and gratifying. In fact, that's really the motivation for AlphaSimplex. I felt a need to be part of the market in order to teach our students exactly how they can make use of all of the various theories that we develop in academia. I've been studying financial markets and teaching investments for so long that, at some point, I simply had to take some of the ideas and implement them — not only because it's fun but also because the proof of the pudding is in the tasting. ■

Christina Grotheer is a contributing editor and an editorial consultant to CFA Magazine.

Andrew Lo, Ph.D., is the Harris & Harris Group Professor of Finance at the Sloan School of Management at the Massachusetts Institute of Technology and is the director of MIT's Laboratory for Financial Engineering. He is also a founder and chief scientific officer of AlphaSimplex Group, LLC, a quantitative investment management company based in Cambridge, Mass., USA.

Lo has published numerous articles in finance and economics journals and is the author of *The Econometrics of Financial Markets* and *A Non-Random Walk Down Wall Street*. He is an associate editor of the *Financial Analysts Journal*, the *Journal of Investment Management*, the *Journal of Portfolio Management*, and *Quantitative Finance*. Lo also serves as a governor of the Boston Stock Exchange, a member of the NASD Economic Advisory Board, and a research associate of the National Bureau of Economic Research. He received a Ph.D. in economics from Harvard University.