Xiao-Ping (Steven) Zhang

To the Victor Go the Spoils: AI in Financial Markets

Artificial intelligence (AI) for finance is a hot topic these days. Top hedge funds are waging a talent war over AI experts the way sports teams compete for pro athletes. Recently, a friend of mine, an AI expert working at a tech firm, got an offer from a top hedge fund. From a financial perspective, it is "an offer he can't refuse," as Don Corleone says. In the sci-fi film *Transcendence*, Johnny Depp applied his newly digitized brain to the markets, pocketing hundreds of millions overnight. We've witnessed the invincibility of machines—AlphaGo, Watson, DeepBlue—and can't help but wonder what's next. In January of this year, a new AI named Libratus joined the club by defeating four of the world's best professional poker players. How close are we to AI's conquest of the game of finance? Will AI experts become the new titans of Wall Street?

Note that AI and signal processing (SP) have a natural connection. Indeed, AI for SP recently has achieved pioneering success in speech and image processing through deep neural network (DNN) technologies. There exists the unbridled expectation people and companies might have on this connection (SP + AI) to make money.

To properly align our expectation, we need to understand the nature of financial markets and, within the laws of that universe, what AI can do for us as a community and for you as an individual investor. From this perspective, I will show that AI cannot change the fundamental structure and dynamics of the financial market. Like all other technologies and innovations introduced to finance, AI may advance our markets and economy, automate away jobs, and even bring a select few vast wealth. But the vast majority of people—no matter their level of AI expertise—will not achieve large excess returns. So don't plan for an early retirement just yet.

A fundamental function of the financial market is resource allocation, i.e., for businesses or people to secure resources for promising investment projects. Everyone knows, or quickly learns, that investment returns are not guaranteed. Investors assume risk when allocating resources because they have imperfect knowledge of the future. A good investment has expected rewards that outweigh these potential losses. Investors of high-risk projects such as start-ups demand high expected returns. To minimize the risk or achieve the optimum risk-return investment, people seek information that helps them see the future more clearly. The financial market is driven by information.

Financial market innovations often help resources flow to the projects that we think are the most promising and needed. An example related to our daily lives is the invention of the credit card, which allows individuals easy access to capital (with high interest/risk) when people are in need. Many economists [1] believe that consumers' continuous spending using the debt accessible through credit cards in economic down times is an important cause for the Great Moderation in the United States in the 1990s. It is easy to take for granted how much technology and the capability to evaluate personal credit through data are a driving force for the popularity of credit cards.

The stock market processes an incalculable amount of data, converting it to information about the future, and boiling everything down into a single number: price. Due to the fluctuations of security prices, many people liken the stock and bond market to one giant casino and attempt to profit by speculating on what the asset prices should be. The best of quantitative hedge funds, such as Renaissance Technologies, D.E. Shaw, Two Sigma, and Citadel, have minted some of the world's wealthiest people, encouraging those of us who are working in SP and AI to fantasize about our shot at a big win. If we can use our skills to predict tomorrow's stock price, or even the direction of the market, we could make big money. It can feel as if we are just one DNN away from riches. Indeed, many open-source DNN tools are available for everyone to use in analysis.

DNN and future AI tools undoubtedly may help us to gain new information,

(continued on page 171)

Digital Object Identifier 10.1109/MSP.2017.2738038
Date of publication: 13 November 2017

 (continued from page 176)

but can they be a crystal ball to predict the stock market and earn a large excess return? Before we dream of the AlphaGo for market prediction, let us review a few basic principles on asset pricing to better understand the role AI could play. I refer to the stock market in this discussion, but the same applies to other assets and securities such as currencies, commodities, and bonds.

Stock price reflects a market equilibrium when supply meets demand, a state in which sellers and buyers agree

In a free market, the price "clears the market," which means the price at which the quantity people buy and the quantity they sell are the same. Why do prices move? The basic concept is represented in the supply-demand curves. When the quantity that sellers are willing to sell does not meet the quantity that buyers are willing to buy, prices change until a new equilibrium is reached. See Figure 1 for the supply-demand curves.

Now add this to the mix: assume that an AI tool can predict tomorrow's stock price. Say the price is going to be US$10 higher tomorrow, represented by the red curves and P(t+1). What will happen?

If everyone in the market has a magic AI tool or at least knows the tool's prediction, then today's price will immediately rise to the predicted point. Why? Because no one would sell at a lower price. Indeed, as long as a few market participants have this knowledge, the price will change immediately. There are many examples of this in the real world such as an acquisition announcement. Company A is trading at US$25. Company B announces it has entered into an agreement to purchase Company A in one month's time for US$35. Company's A stock immediately jumps to US$35 in anticipation of the deal close price. As a result, no one can make profit from a known prediction.

An informal way to think about it is that today's price is the average price of the many possible prices we think we might see tomorrow, the "expected" price. And furthermore, when determining today's price, people discount tomorrow's expected price because, in general, people have an aversion to future uncertainty. The dark shadow curves in Figure 1 represent another possibility tomorrow.

When no one else but one person, say, Alice, has this magic AI tool, then there are two kinds of people trading. Alice, who knows the future price, and ignorant people trading at the wrong price today (see the aforementioned first argument) but will gain this knowledge tomorrow. So if Alice buys the stock today, she would make a free profit when the price rises to where she knows it will be tomorrow. Note that as Alice buys the stock, the price would be moving up along the supply curve, and the demand curve would be moving slightly right toward its correct place. The more Alice is buying, the more the demand curve will move to the right. Another compounding effect would be that, though other people (suppliers) do not have Alice's information, they may quickly infer from Alice's order for a large quantity of shares that there is something there and raise their expected price, moving the supply curve to the left. In such a case, the free profit Alice can make is apparently limited. Of course, if Alice is not trading in a large quantity and no one pays attention to her trading, she may make a one-time profit for herself by keeping her information secret. Real-world examples such as undetected insider trading, or, arguably, the opaque strategies employed by Jim Simons at Renaissance Technologies, can only profit if others do not know or infer what they know.

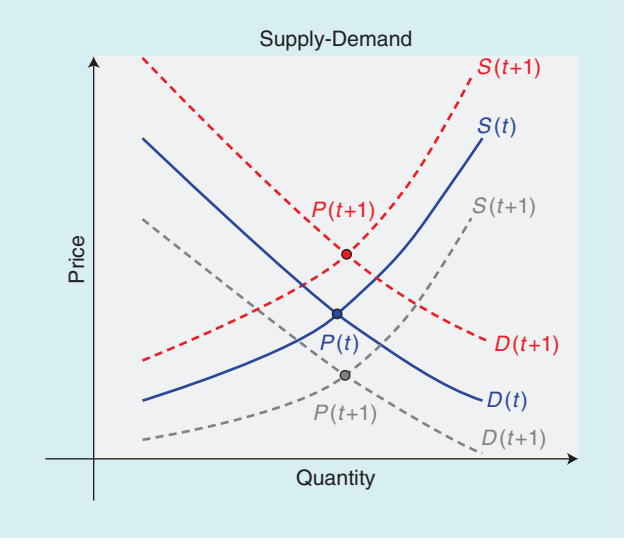

FIGURE 1. An example of the supply-demand and price equilibrium movement at today (time t) and tomorrow (time t+1).

Risks or uncertain future outcomes are an inherent nature of the financial market and the underlying economic activities

Stock price is the result of a market equilibrium when market participants optimize risk-returns. The current stock price represents the best knowledge of sellers and buyers for the expected (average) future of the stocks/firms. The excess expected return in the stock market comes as the reward of holding market risk, also called *systematic risk*, which cannot be eliminated.

This is a more involved argument that can be best illustrated by the Capital Asset Pricing Model (CAPM). The argument also holds for other asset-pricing models. A CAPM (see Figure 2) is an equilibrium relationship between stock prices and the market portfolio. It is obtained when people optimize the expected return (mean of the stock return) for a given risk represented by volatility, i.e., the standard deviation of the stock return. This optimization leads to a market equilibrium, i.e., a mean-variance-efficient (MVE) tangency portfolio T, such that the prices of all stocks satisfy the CAPM relationship according to their correlation to the market portfolio, the optimal portfolio. A market index, such as the S&P 500 index,

is often used as a proxy of the market portfolio. Note that the MVE tangency portfolio is the optimal portfolio in that it gives maximum expected return for a given volatility and risk-free rate. A CAPM is an equilibrium relationship that stock prices must satisfy with regard to their risks and correlations. In a CAPM world, the market portfolio is the MVE tangency portfolio. No one can hold a better portfolio than the market portfolio—the MVE tangency portfolio T. Readers can refer to [2] for an introduction to the CAPM with a signal processing perspective as well as finance jargon I use here.

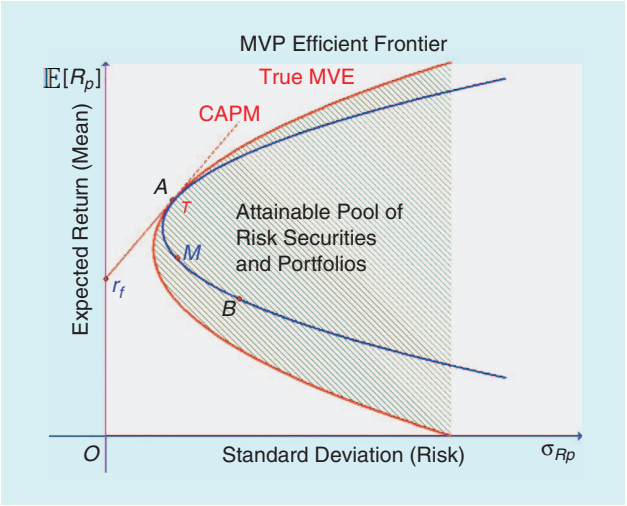

FIGURE 2. Portfolio risk-return optimization and market equilibrium: r_f represents the risk-free rate. The risk-free rate is the theoretical rate of return of an investment with zero risk. In practice, it is often represented by the interest rate on U.S. Treasury bills.

Of course, such market equilibrium and optimal market portfolio will only hold when major market participants agree on the knowledge, i.e., mean and covariance, of all stock returns. If an AI tool provides better estimate of the knowledge, as long as everyone in the market knows, no one can beat the market portfolio in a risk-return sense.

Now, what if only one group of investors, represented by Alice, has the better knowledge, and the other group of investors, represented by Bob, disagrees and stubbornly holds and acts on the wrong knowledge of the firms? In such a case [3], the informed Alice will hold the true MVE tangency portfolio T, while the misinformed Bob in aggregate holds a non-MVE portfolio B. The market portfolio M is now the value-weight portfolio of T and B. Let w denote the proportion of the total value of risk asset owned by Alice, and R_A and R_B denote the expected returns of portfolios A and B, respectively. Then the market portfolio return R_M is

$$R_M = wR_A + (1-w)R_B.$$

The relationship among portfolios is illustrated in Figure 2.

Remarks

- Apparently, Alice's portfolio T has the highest Sharpe ratio and outperforms the market portfolio M.
- The stubborn misinformed Bob's portfolio underperforms the market.
- Note that, if the misinformed investors in aggregate hold the tangency portfolio, the CAPM still holds and no one will outperform the market.
- When the overall weight of the misinformed investors is large enough, it can drag down the market portfolio. This may happen in a market with many nonprofessional stubborn active investors. But in the U.S. stock market, the institutional investment composes the majority. Therefore, M is expected to be similar to the true tangency portfolio T.
- If the nonprofessional retail investors are equipped with AI, they would act more like informed investors that will bring the market portfolio closer to M, making the market more efficient and practically losing less money compared to the informed professional investors.
- Here's what's interesting: if Bob knows that he is not as informed as Alice but doesn't know who Alice is or what her secret holdings are, his best strategy is to be a passive investor holding the market portfolio. If all misinformed investors know they are misinformed and therefore hold the market portfolio, the market portfolio again becomes the optimal MVE tangency portfolio T. However, if Bob uses a wrong AI tool or data and becomes overconfident that he's the informed investor in the market, guess what—he becomes the stubborn guy holding the wrong portfolio who's worse off, dragging the market portfolio down and hurting all the passive investors. Warren Buffett put it this way: "If you've been playing poker for half an hour and you still don't know who the patsy is, you're the patsy." The caveat is that it's easier said than done: you have to have enough unbiased data samples to draw statistically significant conclusion to know who the patsy is.

My conclusion? AI may provide new tools for information. But are you a seasoned financial professional who works closely with the market, or are you a casual investor? If you are one of us, the majority, don't expect AI to bring a quick buck. If you are the former, you have a tough job but there's a chance to make a fortune.

Summary and Q&A

Q: Can AI (or any other technologies) help us better evaluate the risk?

A: Yes, it is possible. But from society's perspective, it is not necessarily a good thing. We want smart people to create totally new knowledge, products, art, etc. that always involve high risk. Accurate risk evaluation may cause these high-risk start-ups to find themselves short of funders, just like an insurance company may not want to insure certain high-risk patients.

Q: Can AI help us better allocate capital in a financial market?

A: Yes.

Q: Can AI replace some financial analysts?

A: Possibly. That may happen in any industry and has happened before in the stock market. The famed trading pits in New York City and Chicago stuffed with floor traders shoving and shouting are no more. They've been replaced by electronic exchanges. Now automation

with AI is threatening financial managers who do work at their desktops.

Q: Can AI generate a few new successful hedge funds?

A: Maybe, there is always a chance to become the next Jim Simons who was a mathematician, or D.E. Shaw, who was a computer scientist. But the overall quantity and amount of successful hedge funds are unlikely to change. The top hedge funds today have no choice but to recruiting top AI talent. It's a Darwinian world where you retain your edge or the hedge fund dies.

Q: But can AI make every smart AI expert excessive risk-free money from stock market?

A: The answer is a solid "no!"

Author

Xiao-Ping (Steven) Zhang (xzhang@ryerson.ca) received his B.S. and Ph.D. degrees from Tsinghua University, Beijing, China, in 1992 and 1996, respectively, both in electronic engineering. He received an M.B.A. degree in finance, economics, and entrepreneurship with honors from the University of Chicago Booth School of Business in 2008. He is a professor of electrical and computer engineering and cross-appointed to the Finance Department at the Ted Rogers School of Management at Ryerson University, Ontario, Canada. His research interests include signal processing, machine learning, big data, finance, and marketing. He is the cofounder and chief executive officer of EidoSearch, an Ontario-based company offering a content-based search and analysis engine for financial big data.

References

[1] (May 24, 2014). Sea of tranquillity: The return of moderation. *Economist.*, vol. 411, no. 8888, pp. 70. [Online]. Available: https://www.economist.com/news/finance-and-economics/21602734-volatility-has-disappeared-economy-and-markets-could-be

[2] X.-P. Zhang and F. Wang, "Signal processing for finance, economics, and marketing," *IEEE Signal Processing Mag.*, vol. 34, no. 3, pp. 14–35, May 2017.

[3] E. F. Fama and K. R. French, "Disagreement, tastes, and asset prices," *J. Financial Econ.*, vol. 83, no. 3, pp. 667–689, Mar. 2007.

SP

Digital Object Identifier 10.1109/MSP.2017.2768458