**Quantitative investing and the limits of (deep) learning from financial data**

J. B. Heaton, Managing Member, Conjecture LLC

The idea of quantitative investing – using powerful computing power and algorithms to trade securities – inspires both awe and fear. Reality is less exciting. With a tiny handful of exceptions, most quant funds have been unimpressive. I explore some limits of quantitative investment, with a focus on the promise – or lack thereof – of techniques from deep learning and artificial intelligence. These limitations help explain the disappointing performance of many quant strategies and cast doubt on the promise of artificial intelligence techniques for improving returns. The main problem is that financial market data is unlike the data that machine learning works well on in computer vision, speech recognition, and natural language processing. While deep learning and artificial intelligence are changing the world in many ways, they are unlikely to generate fortunes for investors, who will continue to remain best-served by inexpensive and passive index products that themselves will be augmented by machine learning techniques to drive costs even lower.

jb@conjecturellc.com
forthcoming, *Journal of Financial Transformation*
This Draft: March 2, 2018

## 1. Introduction

Quantitative investing – using powerful computing power and algorithms to trade securities – inspires awe and fear. Awe arises from the idea that the use of mathematics, statistics, and computer learning might be the closest thing possible to a real-world Philosopher's Stone. Fear comes from worry that traders using computer algorithms to trade securities – often at the "high frequency" of microseconds and faster – can disrupt capital markets and delink prices from fundamentals. Investors have responded more to the awe than the fear. Investment industry sources estimate that quant funds managed more than U.S.$1 trillion by the beginning of 2018. The strategies include everything from relatively simple regression-based "factor models" from firms like AQR (founded by fellow finance Ph.Ds. from the University of Chicago) to highly-complex computer learning models employing the latest ideas from artificial intelligence, including the deep learning models I refer to in my title.

The reality is less impressive than awe or fear suggest. One of the biggest fund implosions of all time was Long-Term Capital, a quant fund run by former Salomon Brothers proprietary traders and Nobel Prize winners Robert Merton and Myron Scholes. More recently, BlueTrend, a Geneva-based fund and Winton Capital, a large London-based fund, have had unimpressive years, as has AHL, one of the main funds of the Man Group, and Aspect Capital, another large quant fund. With only a handful of exceptions – most notably Renaissance Technologies, the hedge fund that mathematician James Simons founded in 1982 - most quant funds do not have consistently impressive performance. Consider Citadel LLC, a huge hedge fund headquartered in Chicago. The hedge fund returned only about 13 percent in 2017, short of the S&P 500 Index's 20 percent gain, a return available quite inexpensively to anyone with the money to open an account at Vanguard Group. The explanation for the difference in returns is not lower risk. Citadel fell nearly 60 percent in 2008, far more than the S&P 500 index. More recently, when market turmoil hit in early February 2018, some of the biggest names in quant fell hard again, including Winton and AHL, and a large quant fund managed by Lynx Asset Management.

In this article, I explore some limits of quantitative investing, with a focus on the promise – or lack thereof – of techniques from deep learning and artificial intelligence more generally. As deep learning – a subset of machine/computer learning – has achieved more and more success in image and speech recognition, product recommendations, and self-driving vehicles, hope has escalated that the techniques allowing advances in other domains will pay off for quant investors as well. Deep learning is a form of machine learning, the use of data to train a model to make predictions from new data. Recent advances in deep learning have dramatically improved the ability of computers to recognize and label images, recognize and translate speech, and play games of skill, in each case sometimes at better than human-level performance. In these applications of deep learning, the goal is to train a computer to perform, even better, certain tasks - such as recognizing the content of an image - that a human usually is able to do quite well.

Financial markets present far different problems than those presented in computer vision, speech recognition, and natural language processing. Unlike recognizing an image or responding appropriately to verbal requests, humans have no innate ability to, for example, select a stock that is likely to perform well in some future period. These limitations are related to the disappointing performance of many quant strategies and cast doubt on the promise of artificial intelligence techniques for improving matters.

## 2. Statistical modeling, machine learning and deep learning

The idea behind many quant strategies is that some variable of interest – say, a move in a security price – can be modeled as some function of available data. The function of the data could be something as simple as a factor regression model or something as complex as a many-layered deep learning computation. We can conceive of the problem as approximating for the variable some existent, but unknown, true function of the data. The problem is to use available data to estimate the relationship or "train the model" and then test the model on new data or data that we set aside from the initial estimation. The basic tools of quant strategies are mathematics (including optimization, probability, information theory, and statistics) and numerical methods using computers.

The problem is that financial markets generate data that is not like the data on which machine learning works well. Machine learning is incredibly powerful with data patterns that are stable. A vision system may take a while to learn all the ways to recognize a dog in an image, since there are so many different angles from which to capture the dog and so many objects that might interfere with the dog's image (e.g., the dog might be partially hidden by a fire hydrant!). But a dog is a dog and remains so, and the algorithm can learn to distinguish dogs of different types from cats and cactuses. The same is true of chess patterns, the game of Go, speech recognition, you name it. The computer is allowed to train on data and learn relationships that remain relatively stable and eventually gets it right.

Financial market data is different. First, the algorithm may identify a relationship that does not actually exist. "Signals" in financial markets come with enormous amounts of "noise." A quant system may falsely identify a signal that does not actually exist or may overestimate the effect size of an actual signal. Second, even if the relationship did exist at one time in the data, it may disappear quickly. The behavior causing the relationship might change as investors learn something that alters their expectations. Alternatively, arbitrage – the actions of the quants themselves - might cause the signals to disappear as investors who see the relationship compete them away. Anyone who has taken an introductory economics course has encountered the model of perfect competition where buyers and sellers are equally well informed and there is no market power. Quants are all trying to hire the same kinds of people, educated at the same institutions, and trained in the same methodologies. As a result, quant trading is largely (Renaissance Technologies excepted?) a commodity business and competition to exploit signals looks a lot like competition to buy and sell identically-graded wheat. Third, if a relationship does *not* disappear – that is, a signal continues to "work" – the relationship may indicate the presence of a risk-return tradeoff, not an arbitrage opportunity. If a relationship that everyone can see continues to exist, then – like the fact that stocks on average return more than U.S. treasuries – it is more likely that the return is compensation for real risk. Indeed, the marketing genius of firms like Dimensional Fund Advisors and AQR is that they implemented simple academic factor model-based investing as if it offered superior risk-adjusted performance when that superior performance is probably just compensation for economic risk that occasionally bites the investors who bear it.

**3. Fooling the machines**

A human – though perhaps quite fallible – is often able to notice or "get the feeling" through some mechanism of intuition we do not yet understand that something is not quite right. That can allow a human trader to put the brakes on a strategy. Computers are not always good at this. Sometimes this is what allows algorithms to make better decisions, as they are not prone to wrong intuitions. But sometimes intuitions are right. Livermore (1940), the famed speculator of the early 20th Century, writes: "A speculator of great genius once told me: 'When I see a danger signal handed to me, I don't argue with it. I get out! A few days later, if everything looks all right, I can always get back in again. Thereby I have saved myself a lot of worry and money. I figure it out this way. If I were walking along a railroad track and saw an express train coming at me sixty miles an hour, I would not be damned fool enough not to get off the track and let the train go by. After it had passed, I could always get back on the track again, if I desired.'"

Because computers interpret information differently than humans, they may miss some trains coming. Recent analysis of image-recognition deep-learning algorithms, for example, reveals that tiny errors – the change of a single pixel in an image, for example – can lead algorithms to fail miserably. This raises the possibility that some market participants may create just such tiny errors in financial data on purpose as a way to change the trading behavior of other active algorithms, using a tiny perturbation of the price or other data to shift a competing trading algorithm from buy to sell or no action at all, or to increase or decrease the size of buy and sell orders.

This is an interesting turn of events. Research suggests that manipulating the prices of securities through mere trading (as opposed to fraud) is quite difficult, at least when humans oversee trading decisions [Fischel & Ross (1991), Kyle and Viswanathan (2008)]. But we know far less about how price manipulation might work in a computer-driven market, and there are reasons to believe stock manipulation is more widespread than recognized [Comerton-Forde and Putnins (2014)]. Quant trading is likely to raise important regulatory issues in the future [Korsmo (2014), Mahoney and Rauterberg (2017)]. The temptation to manipulate markets may be particularly large for some quants as they find themselves unable to generate promised returns legitimately.

**4. Cognitive biases, or, how to make money as a quant without really trying**

One way to understand quant funds – like most hedge funds – is to realize they are not about superior investment performance. Most hedge funds charge enormous fees to deliver performance that consistently underperforms passive (and very inexpensive) index funds. Investors in hedge funds may be the dumb money in the market. They are optimistic and overconfident gamblers who think they can pick a winner notwithstanding the failures of similarly-situated investors to do so. To them, hedge funds – including quant funds – are like casinos. And like the real casino business, hedge fund investors like to frequent the shiniest and brightest. Hedge funds build their casinos accordingly: hiring lots of people with no proven ability to beat the market, but who look awfully smart. This includes data scientists.

Take a firm like AQR. Their results speak for themselves; it is not a performance powerhouse. But its founder, Cliff Asness, is a master marketer of academic research. He has an unparalleled advantage at making investors feel he is implementing for them the lessons of tried and true academic research. By hiring (co-opting?) incredibly smart academics and appealing to academic journal results, he builds a casino that attracts the gamblers who want a University of Chicago patina of academic rigor. But is he adding much value over Vanguard's cheaper index funds? It doesn't appear so. BlackRock is the world's largest asset manager, but it makes far more money on its most expensive products than the passive products that are best for customers (I think of passive index customers as the non-gamblers who come to the casino for the great food at rock-bottom prices). BlackRock has now set up a group to research artificial intelligence in investment. That seems more about marketing than anything else.

This window dressing is unlikely to generate returns for the reasons I reviewed above, but it is likely to exploit investor optimism and overconfidence and make considerable sums for the quant managers, especially those who were fortunate to generate high returns long ago on much smaller asset bases and who therefore can still claim they have high "average" returns. Optimism is an important cognitive bias that draws investors to overpriced active management, including quant strategies. Much psychological research shows that individuals do not base predictions upon objective evidence, e.g., the evidence that the median active manager does not beat passive indices and the evidence that the only reliable persistence in returns is that really bad active managers tend to remain really bad. In a widely cited paper in Nature Neuroscience, Sharot et al. (2011) suggest that optimism may arise because desirable information is integrated into prior beliefs more readily than undesirable information. When newly encountered information – the underperformance of your hedge fund investment in Citadel or Winton, say - is worse than expected, people largely ignore it, perhaps consoling themselves with that "average" return that places significant weight on big returns from the 1990s.

## 5. The task is harder than you might think

Quant strategies face a bigger problem than the limits of data science with financial data. In prior work with co-authors [Heaton et al. (2016)], confirmed by related work [Bessembinder (2018)], I find that active managers are probably doomed to underperform large passive (and inexpensive) indexes, like the S&P 500, in most years because active strategies miss the handful of stocks that drive market results. An underemphasized empirical fact is that the best performing stocks in a broad index often perform much better than the other stocks in the index, so that average index returns depend heavily on a relatively small set of winners. Quant strategies that select subsets of securities from an index are likely to underperform it.

To illustrate the idea, consider an index of five securities, four of which (though it is unknown which) will return 10% over the relevant period and one of which will return 50%. Suppose that active managers choose portfolios of one or two securities and that they equally weight each investment. There are 15 possible one or two security "portfolios." Of these 15, 10 will earn returns of 10%, because they will include only the 10% securities. Just 5 of the 15 portfolios will include the 50% winner, earning 30% if part of a two-security portfolio and 50% if it is the single security in a one security portfolio. The mean average return for all possible actively-managed portfolios will be 18%, while the median portfolio of all possible one- and two-stock portfolios will earn 10%. The equally-weighted index of all 5 securities will earn 18%. Thus, in this example, the average active-management return will be the same as the index [see Sharpe (1991)], but two-thirds of the actively-managed portfolios will underperform the index because they will omit the 50% winner. Quant strategies face a daunting task in beating the odds of missing the best performing trades. And by constant trading, they create even more positions that are likely to underperform market indices.

## 6. Conclusion

In this article, I explore some limits of quantitative investment with a focus on the promise – or lack thereof – of techniques from deep learning and artificial intelligence more generally. In prominent applications of deep learning, the goal is typically to train a computer to do as well or better at a task - such as recognizing the content of an image - that a human usually does quite well. But financial markets present far different problems than those presented in computer vision, speech recognition, and natural language processing. Given the mostly unimpressive performance of quant funds – with a tiny handful of exceptions – it is more reasonable to view quantitative investment management as more marketing than effective trading technique. Moreover, there are empirical reasons that it is very difficult to beat large passive portfolios consistently, and those empirical facts are just as hard for quants to overcome as for other active managers. While deep learning and artificial intelligence are changing the world in many ways, they are unlikely to generate fortunes for investors, who will continue to remain best-served by inexpensive and passive index products that themselves will be augmented by machine learning techniques to drive costs even lower.

**References**

Bessembinder, H. 2018. "Do stocks outperform treasury bills?" Journal *of Financial Economics*, forthcoming.

Comerton-Forde, C. and T Putnins, 2014, "Stock price manipulation: prevalence and determinants," *Review of Finance* 18, 23-66

Fischel, D., and D. Ross, 1991, "Should the law prohibit 'manipulation' in financial markets?" *Harvard Law Review* 105, 503-553

Heaton, J. B., N. Polson, and J. H. Witte, 2017, "Why indexing works," *Applied Stochastic Models in Business and Industry, Applied Stochastic Models in Business and Industry* 33, 690-693

Livermore, J., 1940, *How to trade stocks*, McGraw-Hill Education

Korsmo, C., 2014, "High-frequency trading: a regulatory strategy," *University of Richmond Law Review*, 48, 523-609

Kyle, A., and S. Viswanathan, 2008, "How to define illegal price manipulation," *American Economic Review* 98, 274-279

Lin, T., 2017, "The new market manipulation," *Emory Law Journal*, 66, 1253-1314

Mahoney, P., and G. Rauterberg, 2017, "The regulation of trading markets: a survey and evaluation," working paper

Merton, R., 1980. "On estimating the expected return on the market," *Journal of Financial Economics* 8, 323-361

Sharot, T., C. W. Korn, and R. J. Dolan, 2011, "How unrealistic optimism is maintained in the face of reality," *Nature Neuroscience*, 14, 1475-1479

Sharpe, W. F., 1991, "The arithmetic of active management," Financial Analysts Journal 47:1, 7-9

Taylor S., and J. Brown, 1988, "Illusion and well-being: a social psychological perspective on mental health," *Psychological Bulletin* 103, 93–210

Weinstein, N., 1980, "Unrealistic optimism about future life events," *Journal of Personality and Social Psychology*, 39, 806-820