

Training AI to be as Irrational as Humans: Can AI Beat the Market?

As a college CFO I had much to learn when I met with my college's investment committee. We had a small endowment of around \$100 million, and the committee consisted of some of New York City's top banking and investment leaders. Our meetings usually began with presentations by one or two or our three investment managers, followed by a closed discussion when we would debate shifting our levels of investment among the managers or, perhaps, pulling out of one firm and begin seeking another.

One of our managers had a reputation for making "interesting" investments, putting money with firms for reasons that were not intuitive, but, generally, sound. At one meeting the head of this firm presented his latest thinking on the market and informed the committee of the rationale for his current investments, as well as for his next steps. His returns had been good, as good as the other firms with whom we had been working. In his presentation he projected confidence and pulled us into his thinking. I liked the idea of putting some of the endowment with a firm that went against standard Wall Street approaches.

After he had ended and left the room, we began our discussion. After a bit of positive evaluation, a committee member with a long and successful career in investing said, "Well, Mr. X has done a fine job for us, but I am very uncomfortable when he stops making sense." That set me back. I had enjoyed the presentation and had imbibed the feeling of confidence and even camaraderie that the head of the firm projected. I hadn't bothered to question whether he was making any sense, but, when I thought back on the presentation, I had to agree with the committee member. There were a lot of fine words and interesting ideas, but the man was no longer making any sense.

The committee fired the manager and never regretted it. The moral of the story is: "Don't let your feelings override your logic or your values."

I used to work in Harlem. On a couple of occasions, a man would stop me on my way to my office. Sometimes the man might be wearing a suit and a tie and, in the accents of someone with education, would beg my pardon for interrupting my walk, and would begin a story of a recent mishap where he lost his wallet and needed a few dollars to get home and set things right. Now, as you are reading this, you don't get the feeling I had. He was, of course, a con artist and with his dress, his speech, and his courtesy, he was making me feel that he was "one of us," and I should relax and help this unlucky man. I could easily be in his shoes.

He had a "tell," as the gamblers say, however. His shoes were a mess. Please note that I am not saying don't give money to a con artist. I'm saying, "Don't let your emotions make a decision that your logic and values should make." I want you to decide right now whether you would give a couple of bucks to a con artist. I want you to decide right now whether you would give your extra kidney to your brother or sister. I don't want you thinking about whether your relation is a lazy crumb or has ruined your life when they need a kidney. Decide now whether you would be able to do it. Find your values before the emotional pressure of a situation dictates a poor decision. Learning your values is not easy when you are under pressure.

(As an aside on shoes as "tells," I used to ride the NYC 7-train every workday. Often on the way home, I might see a man asleep sitting up. His clothes would be dirty with dust, sometimes in

Training AI to be as Irrational as Humans: Can AI Beat the Market?

need of mending. Almost always when I checked his shoes, however, they would be tan work boots with steel toes and solid soles. He was not homeless. He was probably a Jamaican skilled tradesman working in construction. Plaster dust and sawdust can make a mess of you, but they are marks of hard work.)

Investing

To begin, let's stick to simple cases: buying and selling stock, both regular and short. In a short sell, you sell the stock before you buy it, hoping that the price will fall between the time that you sell (what you don't have, of course) and when you buy it (to make good on that ephemeral previous sale).

Old wisdom has it that the price of stocks, and thus of the market itself, is six months ahead. If a stock is increasing in price, that means, according to this wisdom, that in six months, the company will be doing better. This is certainly true when a company releases information, like its sales for the last month, that indicates that it is (and, perhaps, will be) doing better. This wisdom assumes a rational market. In a rational market, according to the members of the investment committee mentioned above, the market (for a stock or the market itself) goes up with good news and goes down with bad.

Thus, an irrational market must be such that the market goes down when good news comes out and up when bad news appears. So-called "bull markets" thus behave irrationally. No matter what the news may be, the market goes up: good news, bad news, no news.

Thus, if this wisdom is true, we can beat the market by first determining whether the market is behaving rationally or irrationally, and second, selling or buying depending on the news and the rationality of the market. This investment methodology has worked reasonably well for long periods in the past.

But... it no longer works. Rational and irrational people no longer do most of the trading in the stock market: algorithms do. An algorithm is not artificial intelligence, no matter how smart we try to make it. An algorithm does not learn over time. An algorithm is merely a set of rules. These can be very numerous and complex, but they change (that is, "learn") only as their programmers learn and change the rules.

Here's a simple rule: "Buy whenever a stock increases more than 1% in a day and sell when a stock decreases more than 1% in a day." Now imagine 1,000 stock trading firms who invest according to this rule. Suddenly, large swings are possible in a market because a small change triggers a massive sell-off by 1,000 firms all invested in the same stock, using the same algorithm.

Then the programmers add another rule: "The day after a massive sell-off, buy back twice the number of shares in the same company that were previously owned by the investment fund." Aha! Now we really have volatility.

Strange as it may seem, programmers who write algorithms tend to be rational. If things look good, buy! If things look bad, sell! If things look bad and you are the first to know, sell short,

Training AI to be as Irrational as Humans: Can AI Beat the Market?

then tell a lot of people the bad news (and disclose that you have hugely “shorted” the company) and watch the price crash so that you can buy it to complete your previous sale at a huge profit (even if the so-called news is a little shaky). People will hate you, but you’ll be rich.

So far, so good. Let’s now move into a little science fiction. Let’s go from algorithms to artificial intelligence. Now, we will pour a ton of historical information into our computer and have it learn what works to make a profit. We have a game that’s a little like chess, except the rules aren’t very clear and our history is based on a changing set of actors, using shifting sets of rules, usually rational, but sometimes not.

We can start it with the buy/sell algorithm above, but then we would make it go through all the information available that would help it learn when there were exceptions to this rule and exceptions to the exceptions. In the end we would have a new set of algorithms that are so complex that the whole thing is called a “black box.” The computer can’t explain what it knows, and what it knows keeps changing.

Let us also suppose that our 1,000 investment firms are all creating AI investment models. Can we predict some possible outcomes?

I would predict that the AI models would all figure out that the best way to make money is to create an irrational bull market. They might all figure that they should randomly buy stock, hold it for six months, then sell it, regardless of performance. If a few began doing this, the others would “learn” that it worked and would join in. Then everyone would start randomly buying stocks, holding them, selling later, and immediately buying new shares. If the market keeps going up, why would anyone do anything other than buy stock?

Could this go on forever? If you were a set of AI computers maximizing wealth, why would you doubt that it could go on forever? Every sale brings a profit. It’s an irrational market. Why pop the balloon?

A rational person might remember the parable of the tulip market. Everyone was getting rich buying and selling tulip bulbs as the price rose into the thousands. Then, some fool said, “This is crazy. These are just tulip bulbs.” Suddenly, no one was buying. The market crashed.

This crashing scenario depends on the existence of a few rational people. Do they really exist?

There is a more frightening scenario. Everyone, and this means everyone, realizes that the old ways of making money are more trouble than they are worth. Why not quit and just keep investing, handing money over to the AI on the corner?

So, the workers at the auto plants quit and start investing. The teachers quit and get into investing. The farmers put their tractors in their barns and put all their money into the market. Everyone starts to make a lot of money in the stock market, but there are no cars to buy, no gas at the pumps, the schools are closed, and there’s no food in the markets.

Training AI to be as Irrational as Humans: Can AI Beat the Market?

But that's the flaw with AI. It can be built to win games with a fixed set of rules, or it can be built to emulate human behavior. It can win at chess, or it can drive a car as well as an alert human. It can't, however, be gracious in its win, and it cannot make moral decisions, since being gracious takes sincerity, and, well, humans can't agree on what it is to be moral.

AI cannot tell the difference between rational and irrational behavior. All it knows is that X behavior occurs in Y situations when A, B, and C occur with a probability of D. We can put a rule override on the AI guidance: "Don't buy tulip bulbs when they cost more than \$10,000." Or, should that be \$1,000 or \$10? What's the right rule?

But let's push AI further. It can not only build a dynamic set of investing algorithms based on historical data, but it can also learn from feedback on its own decisions. Let's suppose that it buys a stock after bad news about the company comes out, and it loses money when it sells in six months. It learns, and it updates itself such that it is a little less likely to buy after bad news. In effect, it becomes more rational.

A company comes out with a new product that no one seems interested in. Our AI avoids it. Would it buy Apple in its first stock issue? No. AI would be rational, but too safe. Perhaps it would learn to randomly take chances. Then, would it be much different than humans?

Would it stop buying tobacco stock, oil stock, munitions stock, or big pharma? Could it learn to be something like moral? It might be able to predict when public opinion was reaching a point that made pulling out of a stock a reasonable value idea. So, it might learn to be a moral investor if the morality had a financial reward.

So, now that we have a bit of understanding of AI and the stock market, what do I predict. I must admit that AI will be able to pull in and process far more information than humans are doing now. I predict that it *will* be able to "look at the shoes." That is, it will be able to evaluate better than would humans the quality of the information without the emotional misinformation that often comes with news.

I predict that, because of its ability to learn, it will be more rational in its methods than humans, while, however, appearing irrational because of the complexity of its rules.

I predict that the earliest deployments of AI in investing will be failures because too little feedback will be available, but that later developments will be successful and will tame some of the volatility of the current, simple algorithm-driven market.

I also predict, however, that widespread adoption of AI in investing will neutralize any advantages, and that irrational, human investing will occasionally pay off handsomely, making headlines, but being, in fact, extremely rare.

Thus, although stock investing is not ideally suited to AI with the market's lack of rules and drifting behavioral patterns of investors and companies to mine and mimic, it has a future. It can learn from more information than a human is capable of, and it puts algorithmic-driven investing to shame.

Training AI to be as Irrational as Humans: Can AI Beat the Market?

My next question, however, is: Does AI have a role in regulating the stock market? Hmm... maybe...

But... picking stocks is too simple. Even a human could do a reasonable job. What about managing a portfolio?

Now we would make AI not only pick stocks, but sell stocks, and build a portfolio with many kinds of investments: bonds, cash, treasury notes, derivatives, etc., each with its own expected rate of return, taxation exposure, and risk. We now need a set of rules that does more than maximizes profit. It must maximize profit while keeping risk at a level with which the firm feels comfortable.

We need more than an evolving rule on when to buy a stock, but an evolving set of rules on when to buy, when to sell, and how to balance the portfolio against risk.

Let's suppose that we build it. We take all the historical data and plug in a full firehose of current data. We tell it that we want to be 99% certain that we won't go bankrupt. It takes our current portfolio, grabs the historical data, and comes up with a set of rules. The black box then says, "Sell 300 shares of Y company and buy 400 of the preferred bonds of Z Inc." It is moving toward an ideal portfolio within the constraint we gave it.

The head of the firm looks at the recommendation and says, "That doesn't make sense. I know that we are too conservative now." The programmer replies, "You want to argue with a black box?"

Compare this with the current method, known as stochastic modeling. With this style of modeling, you take your current portfolio in broad strokes: \$X of blue-chip stocks, \$Y dollars in the stock of small companies, \$Z dollars of municipal bonds, etc. Each has a historical return that can be viewed as a probability distribution: sometimes the return is good, sometimes not so good, and once in a while it's terrible. For each period in the model's projection, the computer rolls the dice and picks a return from the probability distribution for each chunk of the portfolio.

You run this projection a thousand times and get a distribution of possible overall portfolio performances. Then the analyst gives the results to the head of the firm: "We have an average (over the 1,000 trials) performance of an 8% return per year with a 95% probability that we will not go broke."

The head of the firm says "Yikes!" (Actually, heads of firms never say yikes.) "That means there's a one in twenty chance that we will go out of business. I will lose all my money and never work on Wall Street again. Let's adjust the portfolio. Add more bonds and drop the small cap investments. Let's see if we can get the risk of failure down to one percent."

While you may not fully understand Monte Carlo simulations from this brief exposition, you can see the difference between interactive modeling and AI. Interactive modeling is used to help us learn about our values, our risk profiles, and the impact of options. AI gives us a solution without

Training AI to be as Irrational as Humans: Can AI Beat the Market?

a chance to learn. It may perform very well, but if we can't learn as we use it, we will find it difficult to trust.

I think the next evolution with artificial intelligence will be toward helping us learn, not just handing us a solution.

What do you think?

There's a third level of complexity: feedback. Every time AI evolves and shifts its rules, competitors and industries react. If AI begins buying more stock based on short-term results (like quarterly profits), competitors may move in the opposite direction, while companies will find ways to improve short-term results, perhaps harming the longer-term. The more investment firms react to short-term results, the more companies will slow investment in new plants and equipment and show greater short-term profitability, for example, at the expense of longer-term results.

Now AI must anticipate the possible moves of competitors and investment opportunities. Suddenly chess seems simple. Would investment managers trust AI's recommendations now?

AI: "Sure. Trust me. I'm smarter than you are."

Nathan Dickmeyer © 2023

<http://www.DickmeyerConsulting.com>

<http://www.amazon.com/author/ndickmeyer>

<https://www.thehelmsfordpress.com/>