

What Do Trading Algorithms Know?

Ayan Bhattacharya

Baruch College, The City University of New York
New York, USA.

© 2017 Ayan Bhattacharya, All Rights Reserved.

Let us start with an intriguing puzzle at the heart of Epistemic Game Theory: Is it possible to have a configuration of beliefs such that,

“Ram believes that Kali assumes that *Ram believes that Kali’s assumption is wrong*”

(Brandenburger & Keisler, 2006)?¹ If Ram believes that Kali’s assumption is correct, then he believes that the italicized second half of the sentence – Ram believes that Kali’s assumption is wrong – holds. But that immediately leads to a paradox: we started with Ram believing that Kali’s assumption was correct! That we are able to conjure up such puzzles at will suggests that paradoxes like these are part and parcel of our thinking – tucked away deep in our minds and manifested in the many contradictory decisions that we seem to take in ordinary life. How troublesome are such impossible beliefs when we trade? Do trading algorithms, too, grapple with such paradoxical beliefs? Algorithmic trading, especially in advanced markets such as the US, is the first instance of large-scale, real-time, interactive, automated decision-making in an ecosystem outside of computer science. And the many fascinating questions the area has been throwing up, especially as algorithms mature, has left all parties – researchers, practitioners and market regulators alike – scratching their heads.

1. Explosion in Algorithmic Trades

Algorithmic trades account for a majority of the trades in the US markets today.² Not just US, after a brief lull following the financial crisis, algorithmic high frequency funds are on the ascendant in almost all global trading venues.³ This is partly because legislations in Western markets, like Regulation National Market System (RegNMS) in the US, and Markets in Financial Instruments Directive (MiFID) in Europe, make the transition to algorithmic trade an inherently lucrative proposition.⁴

¹ Adam Brandenburger and H. Jerome Keisler (2006). “An impossibility theorem on beliefs in games.” *Studia Logica*, 84(2), 211–240. Protagonist names changed.

² Evelyn Chang. “Just 10% of trading is regular stock picking, JPMorgan estimates.” *Cnbc.com*, June 13, 2017. Accessed: November 03, 2017.

³ Robin Wigglesworth, “The quickening evolution of trading — in charts” *Ft.com*, April 11, 2017. Accessed: November 03, 2017.

⁴ David Easley, Marcos Lopez De Prado and Maureen O’Hara (2013). *High Frequency Trading*. Risk Books, London.

A less noticed but equally important development has been the growing maturity of the algorithms at use. Trading algorithms had their birth in the portfolio insurance movement of the 1970s and 80s.⁵ Yet the algorithms of yore were very simple species: elementary sets of instructions to automate repetitive human trading tasks. Over the years, as research in computer science accelerated, so did the sophistication of algorithms used in trade. Many algorithms today virtually learn on their own: fed with vast reams of historical data, they first spot consistently profitable trading opportunities; these are then strenuously tested with more historical and simulated data, and if the discovered pattern stands scrutiny, the algorithms are released into the market. Ask a human overseer of algorithmic trading to define concretely the trading rule that is being implemented, and he will turn a blank. All that the human can tell you, really, is the dataset used in training and the P&L on that dataset.⁶

2. Romancing High-Frequency

From the early days of portfolio insurance, researchers have been interested in understanding the impact of algorithms on markets. In fact, each new algorithmic innovation has meant dozens of new papers in important journals extending older models to account for new facts. The latest in this line of work is the study of high frequency trading. Early algorithms were simply routine automations – their primary advantage was that they did not make “human mistakes” when the same task had to be repeated mechanically, innumerable times. Since the early 2000s, however, an added feature of algorithms has been the fiendish pace at which they work. A number of technological innovations – from the power of the chips that do the calculations to the infrastructure that conveys market signals – have contributed to this jump to high-frequency. The state-of-the-art research in finance, at present, tries to understand the many effects of this fast paced trading environment in markets that still harbor many slow, lumbering legacy traders.

Yet, right from the beginning, there have been dissenting murmurs in research circles: maybe algorithms represent a completely new paradigm.⁷ Maybe tweaks on older models do not convey the full power of the algorithmic vision. Maybe we are missing the forest for the trees.

2. Some Economic Theory

Most models used to study algorithmic or high-frequency trading work under the aegis of classical non-cooperative game theory or its asymmetric information variant. Given market participants and their strategies, the aim is to uncover the Nash equilibrium or one of its

⁵ Anora M. Gaudiano. “Here’s one key factor that amplified the 1987 stock-market crash.” Marketwatch.com, October 19, 2017. Accessed: November 03, 2017.

⁶ Bryant Urstadt. “Trading Shares in Milliseconds.” MIT Technology Review, December 21, 2009. Accessed: November 03, 2017.

⁷ Herbert A. Simon (1969). *The Sciences of the Artificial*. MIT Press, Cambridge.

refinements. An equilibrium represents a stable point of the game – no market participant wants to deviate once he finds himself in equilibrium. However, in a series of influential papers beginning in the 1960s, Robert Aumann, John Harsanyi, Reinhard Selten and their many collaborators began to lay bare the many pitfalls of naively adopting the classical approach. Aumann’s work, in particular, established epistemology as a legitimate concern of game theoretic reasoning. All these early pioneers, including Nash, went on to win the Nobel memorial prize in Economics.

What this line of research contended, broadly, was that any actual game-play needed to be preceded by reasoning about game-play – how a game played out in reality depended crucially on how players reasoned about other players’ plays and payoffs. Each player had to build a mental model of the game before she began play, and had to update the model as the game proceeded. In equilibrium, the mental models of the players about each other, and about the game, had to be consistent.

Publicly available information about the game and game-play – called common knowledge in game theoretic parlance – reduces the need for such game related reasoning because a player can directly use the public information. In classical non-cooperative game theory, everything about the game and its players, except for the final equilibrium strategy, is common knowledge. So reasoning in such games is confined to deriving equilibrium strategies. In asymmetric information games on the other hand, some players do not know their own payoffs while other players do; all else, however, is common knowledge. Any reasoning in such games is confined to the asymmetry in information and concomitant equilibrium strategies. For most of the history of financial market research, classical and asymmetric information games have been deemed enough to understand the behavior of market participants. Such models were earlier used to understand markets with exclusive human traders, and now they have been extended for markets with algorithmic trading.

3. Connecting the Dots

If there is a single characteristic of algorithms that is becoming more and more pronounced, as algorithmic trading matures, it is not their speed – which of course is going up – but the way the algorithms uncover trading opportunities. In the early days, this was human dictated: look for arbitrage opportunities that come from speed advantage, or high Sharpe ratio, or positive alpha, or whatever other designated criteria we humans would decide. In fact, this is exactly how human traders would train themselves! For this is what MBA and other advanced finance programs teach. And this teaching represents a lot of common knowledge. Most human market participants have this inbuilt corpus of common knowledge before they begin serious market trading. Most algorithms had this corpus too – at least till recently.

When algorithms spot profitable trading opportunities “on their own”, without human intervention, they no longer share the exact human corpus of common knowledge. And what is

more, because humans have no way of deciphering the algorithmic logic – beyond the fact that it makes profits or losses – we can never be sure what common knowledge an algorithm has inferred from data. Common knowledge allows humans to bypass the ordeal of building complicated mental models of trading games. Going back to the example at the beginning of this article, because common knowledge makes belief models unnecessary, we do not have to wrestle with the impact of such impossible beliefs in trading situations. For an algorithm learning on its own, however, the picture is unclear. At the very least, its common knowledge corpus is likely to be different from a human trader. This implies that asymmetric information game models, prevalent in the literature at present, may not be enough to represent the complexity of trading situations involving algorithms.

4. The Challenges

These are still early days for research in algorithmic trading, but one thing seems fairly certain: fruitful progress in the field shall come only through meaningful collaborations between researchers in finance, game theory and computer science. A promising approach seems to be the study of adaptive strategies. Recent work suggests that any market with adaptive algorithms – i.e. algorithms that try to repeat past successes and avoid past mistakes – converges to a correlated equilibrium.⁸ Correlated equilibria result when players have access to a shared public signal, in addition to private signals, and they are supersets of Nash equilibria.

For Indian financial markets, the message seems to be mixed. Algorithmic trading has been growing fast in India, but Indian market regulators need to be wary of prescriptions derived from complicated market trading models – such models are still very much works in progress. On the other hand, this represents a great opportunity for the Indian finance research community. Given India's historic strengths in computer science and programming, it is but natural to expect interesting fundamental research on algorithmic trading from Indian academics.

⁸ Sergiu Hart, Andreu Mas-Colell (2013). Simple Adaptive Strategies. World Scientific, Singapore.