LUISS

Dipartimento di Economia e Finanza

Cattedra Corporate Finance

The Efficient Market Hypothesis and trading AI advancements – An overview

Prof. TARANTINO EMANUELE
Relatore

MANCUSO FRANCESCO
Matr. #243261
Candidato

Anno Accademico 2021/2022

Sommario

1. Introduction

This essay will discuss the Efficient Financial Markets Hypothesis and the possible implications that newly introduced trading techniques based on Artificial Intelligence might have on market efficiency.

First, the Hypothesis itself will be introduced. The first chapter is about the discussions around the definition of efficiency, and how it has evolved over the years. It presents evidence supporting the theoretical and empirical validity of the hypothesis in general, and then goes on to highlight some counterarguments, pitfalls and faults on the efficiency theory as presented in the first section.

Finally, the debate is settled around a more economically sensible version of market efficiency. Chapter 1 is to be taken as a general introduction to the problem and is intended to reflect the fogginess of the debate around it, also due to blurry lines between relevance of statistical findings and contradictory empirical experience. It is concluded that although not absolutely precise, the Efficient Market Hypothesis constitutes a useful approximation of how the market should work in theory.

The second chapter narrows the scope of the essay around its main topic, Artificial Intelligence's impact on market efficiency. It uses more recent studies, and it's contextualized around the resolution of the debate as presented in Chapter 1. Although not descending into the details of how exactly algorithms work, as it would be well beyond the scope of this essay, it provides an intuition of the strategies employed by developers; such an understanding might be useful to frame better the ensuing debate on AI's impact on market efficiency.

The same chapter presents optimistic scenarios for market efficiency first. These views use statistical studies to argue that AI activity improves quality of prices, incorporates information more efficiently, and helps the market function better overall. Pessimistic views are opposed to the optimistic studies in the subsequent section. According to this view, a critical role is played by input data and model quality on whether algorithmic trading, especially in its High Frequency variant, might help the market to function more efficiently and better allocate the economy's capital resources.

The essay will conclude that although the financial world appears way more unstable and fragmented than when the Efficient Market Hypothesis was first formulated, and severe cracks have appeared in its theoretical soundness, also due to Algorithmic Trading, the same factors that might suggest that such a hypothesis is out of fashion could be turned to instead show that the basic foundations of the Hypothesis still resist the passage of time.

2. Efficient Financial Markets: Critics and Proponents

"When the capital development of a country becomes a by-product of the activities of a casino, the job is likely to be ill-done." *- John M. Keynes*

In this first chapter, the "strong-form" of market efficiency is first enunciated in its basic features and implications. Objections to the hypothesis' strong form are then examined in the second section, ranging from theoretical reasoning to more formal proof of market inefficiency. Finally, these views are reconciled in the third section and efficiency evaluation parameters are then enunciated as a basis for further analysis.

A. Brief overview of the Hypothesis in its canonical strong form

The notion of efficient financial markets synthesizes the theory, supported by empirical data and documented observations, that capital markets' main instrument for resource allocation, the price mechanism, is efficient, and so the price of a security at any given point in time "fully reflects" all publicly available information regarding that security (Fama, 1970). This simple statement carries many powerful implications; implications that have been widely discussed - and challenged - by scholars and experts along the theory's long lifespan. To get into this debate, for the sake of clarity, let us first get situated in the financial markets field, and then deconstruct and discuss the arguments posed for and against this fascinating theory.

When talking about capital markets, we must first clarify that, for the purposes of this essay, that expression is to define the aggregate of all those institutions and mechanisms apt to allocate capital across the economy. Their function is channeling funds from surplus units (economic actors that save money, i.e., units whose disposable income exceeds their capital outflows) to deficit units (economic actors whose capital outflows exceed inflows) with productive investment opportunities. When talking about capital markets we do not then differentiate between markets for securities with different maturities, so we do include in our analysis the so called “money markets”, i.e., markets for high-liquidity securities, usually set apart from capital markets in the narrow sense, that would include only markets for less liquid, long-term securities.

The market’s instrument to perform its function is the price mechanism. Onto the workings of this mechanism lies the bulk of the debate around efficient financial markets. Defining the medium through which efficiency is evaluated, as emphasized by E. Fama in his influential 1970 article and subsequent studies, requires referencing the rather neoclassical concept of equilibrium; a market is therefore “efficient” when security prices reach the equilibrium point, in which all available relevant information is accounted for in the outcome of the price determination mechanism (Fama, 1970). The price of a security is reflective of its expected return and its risk (volatility), as per the “two-parameter” asset pricing model (Sharpe, 1964). The asset pricing model used when considering issues of market efficiency is of primary importance, as price determination is deeply intertwined with the definition of efficiency.

Without a consensus on how to measure efficiency, i.e., a model of price determination, the hypothesis is not testable. An asset pricing model is necessary to give the hypothesis context and practical relevance. The hypothesis of market efficiency would remain an abstract concept, and there would be no way of determining whether information is reflected in the market. The need for an asset pricing model to be always employed when testing efficiency generates the “joint hypothesis problem” (Fama, 1991). This problem refers to the fact that when examining empirical results to evaluate efficiency, any lags in the reflection of new information into prices or any correlations between subsequent price changes could be the result of an imperfect pricing method rather than an inefficient market. As the two are inextricable, the question always remains open when considering matters of market efficiency.

How these variables are evaluated and translated into price changes because of new information is the focus of scholars interested in market efficiency. Price determination is a fascinating and complex phenomenon; the market “prices in” risk and expected return because of the collective independent actions of small individual units, in this case the investors. Investors’ actions are assumed to be utility maximizing and rationally dictated by a well-defined “information set” (Fama, 1970). Investors include in the information set all relevant news and notions that can help them evaluate investment opportunities.

The efficient financial markets hypothesis is primarily concerned with relating the changes in investors’ information sets, so the introduction of new information in the market environment, to changes in security prices. The statement that financial markets are efficient, and therefore security prices fully reflect all available information, has a strong empirical content, and its implications have been tested and observed in many instances. The main pillars of the hypothesis’ experimental background are then best described together with the tests to their validity.

The strongest and most controversial theoretical consequences of an efficient market are on the probability distribution of returns across securities and on investment strategies. These are best observed in return predictability tests, or weak form tests, and event studies, or semi-strong form tests (Fama, 1970) (Fama, 1991). Return predictability tests hinge on the theory that security prices follow a “random walk”; according to this theory, a security’s successive changes in price are uncorrelated to each other, and therefore the time series depicting it will follow a random path. Return probabilities are, according to the random walk theory, uniformly distributed across time; therefore, forecasting future returns by using historical stock price data is not possible.

The consensus is that, since external events can be assumed to be random, as they are incorporated in market prices, then securities’ changes in price too can be said to be random. Studies on price fluctuations’ non-correlation are highlighted by Fama in his work. Event studies, on the other hand, assess the speed at which the market prices-in new information. In this field too, empirical research seems to support a strong capability by the market to portray information via prices. An interesting example is provided by stock prices’ reaction to earnings announcements by companies, as it was observed that stocks tended to increase in value gradually from up to 3 or 4 months in advance with respect to the actual positive earnings announcement (Ball et al., 1968) (Fama, 1970). More generally, if specific time coordinates can be identified for a firm-specific event, the market can be observed to adjust corresponding securities' prices quickly and efficiently.

The random walk theory and evidence from event studies suggest an important development of the efficient market hypothesis: an investment's performance in one instance is not linked to any future or past performances. An investment strategy's success in one period does not indicate that it will keep succeeding in following periods too, since price fluctuations are unpredictable. From this it follows that there is no strategy that can realistically and systematically outperform a simple buy-and-hold strategy on average. In other words, in an efficient market there is no investment portfolio available to consistently beat the market (Malkiel, 2003) (Fama, 1970) (Fama, 1991).

Since everyone has access to the same information, assuming acting on the information does not bear excessive costs, any unexploited arbitrage opportunity will be corrected by the market almost instantly; the increase in demand for an underpriced security will drive the price upwards until equilibrium level is attained, while the rise in the supply of any overpriced security will restore the equilibrium by applying downwards pressure on its value. The efficient market is then endowed with perfect arbitrage capabilities for investors, which will eliminate inefficiencies and valuation errors automatically.

Given open and complete access to available information, in an efficient market any investor can profit from just holding on to a diversified portfolio of securities long-term. The only way for an investor to obtain greater than average returns on her investment without accepting greater than average risk would then be to have access to not publicly available information - if market efficiency holds, the only source of extraordinary returns is the existence of monopolized information. The tests for this implication are called "tests for private information", or strong-form tests (Fama, 1970) (Fama, 1991). Tests for private information are difficult to carry out since insider traders have little incentive to make their position public.

Nonetheless, strong-form tests can be carried out by evaluating the performance of professional investment managers relative to the market. Indeed, if investment fund managers had access to non-public information, either of their own production (expertise, effective technical and/or fundamental analysis) or insider information, their funds' returns should be expected to regularly outperform the market. Observations show that this is not the case (Fama, 1970) (Fama, 1991) (Malkiel, 2003). Investment funds that pursued active strategies did not outperform the market on the regular basis that the fees charged by them would lead some to believe.

B. Arguments against and some pitfalls of the strong form Hypothesis

Although the efficient market hypothesis was the centerpiece of financial-economic thought for more than two decades after its initial formulation, to the point that it was dubbed "the best-established empirical fact in economics" (Malkiel, 2003) because of its apparently strong and irrefutable practical roots, many studies and schools of thought challenged the hypothesis and its neoclassical background on the wide fronts of price mechanism and information incorporation efficiency.

A wide span of critiques to the hypothesis in its strong form is put forth by neo-keynesian economists. In fact, through a Keynesian view of the stock market, perfectly rational behavior by investors would completely turn the hypothesis on its head based on its very own assumptions (Andersen, 1983). It is argued that in a setting of perfect information, if prices accurately reflected all publicly available information, then an investor should be in theory able to infer at least some of said relevant information from the price itself. Since this is hardly possible, neo-keynesian financial economists attack the notion of information incorporation efficiency by arguing that prices are noisy signals, in the sense that they portray a wide range of information different from just the relevant news contained in the information set.

The metaphor put forward by Keynes himself in his widely influential *General Theory of Employment, Interest and Money*, that compares the stock market to a beauty contest (Keynes, 1936) is very explicative of this view. According to such view, perfectly rational investors have no incentive for gathering relevant information, as all of it is supposedly believed to be reflected by prices; instead, the market participants are focused on predicting fellow investors' expectations on an increasingly high number of levels, as they will try to predict what the average expectation of the average expectation of returns will be, and so on.

It follows that in an information-efficient market any price displayed for a security will be believed to be an equilibrium price (Andersen, 1983).

This leaves the market price determination mechanism extremely vulnerable - and, in the most extreme Keynesian view, determined by - a series of spiraling self-fulfilling prophecies. In other words, if everyone believes that an information is relevant, whether it is or not, the information will make itself relevant in a self-fulfilling way (Andersen, 1983) by reflecting itself in the "equilibrium" price.

Prices in the stock market must then necessarily be "noisy", or there would be no incentive for private actors to produce information, and the market would "*live in its own dreams*" (Samuelson, 1957, p. 215). It follows that the market is not efficient in incorporating information, or what constitutes "relevant information" in the context of the efficient market hypothesis is not specified clearly enough.

The neo-keynesian critique is set forth on theoretical grounds by pushing the hypothesis' assumptions to their extreme; but issues when it comes to effective valuation of new information, self-fulfilling prophecies, and behavioral considerations on efficiency have been highlighted more practically by other studies as well. Examining these contributions emphasizes some very interesting and recurring objections to market efficiency.

Perhaps the most famous and discussed instance of a recurring, apparently unexploited price inefficiency is the January effect. The January effect refers to unusually high, seasonal, and predictable returns on equity of firms with a small market capitalization in the first weeks of January (Reinganum, 1982). Among them, the most actively traded securities are those that end up experiencing the highest unusual returns. According to the strong form of the efficient market hypothesis, such an occurrence should be immediately corrected by arbitrageurs which would call on those stocks in December and sell in January, rebalancing supply and demand and eliminating the price inefficiency.

Some argue that such an occurrence could be the result of investors selling portfolios composed of underperforming "losers" at the end of the fiscal year for tax purposes (Reinganum, 1982). That would return rather large quantities of equity relative to the small size of the concerned firms to the market; this equity then, usually somewhat illiquid because of the small size of the firms' market capitalization, would benefit from the increased trading rate, receiving a liquidity premium in January. This hypothesis is although not confirmed by empirical proof and admittedly doesn't explain the anomaly in its entirety (Reinganum, 1982). There remains the possibility of at least a portion of the returns being explained by an irrational speculative motive by some investors, pricing itself in as a self-fulfilling prophecy.

Other behavioral explanations to market inefficiencies such as the January effect have been proposed. There may exist a tendency by investors to overreact or underreact to new information being introduced (De Bondt, Thaler, 1985). This theory puts into question the market's efficiency at pricing-in information, as it proposes that small capitalization "losers" are not only more likely to get sold at the end of the fiscal year, but investors also require a risk premium on that equity that is greater than what would be rationally expected.

Indeed, more broadly, it has been found that, even adjusting for risk, the portfolio composed of the bottom fifth of firms by market capitalization realized consistently larger returns than the portfolio comprising the top capitalization firms over a 45-year period (Lustig, Leinbach, 1983). This so-called Small Firm Effect was reconducted to a joint hypothesis problem, with behavioral ramifications.

The small firm effect is said to be due to the CAPM, the pricing method commonly used to derive equilibrium prices, being "misspecified" (Lustig, Leinbach, 1983). It is argued that the CAPM may be less efficient at evaluating small capitalization firms compared to large firms, and that is because not all investors are endowed with the same knowledge of the market; it is impossible, using the classical pricing system, to factor in the opportunity cost bore by investors to gather additional information on the small firms and undertake a safer and more sensible investment. According to this view, if that cost was factored in, the returns wouldn't be abnormal.

Moreover, the overall reliability of the statistical methods used to evaluate market efficiency has been put into question. In fact, persistence of large inefficiencies in the market may be as consistent with empirical proof as market efficiency itself (Summers, 1986). It follows that prices may not be rationally reflective of all available information.

Valuation errors may be difficult to detect by just looking at security prices. It is suggested that there may exist a collective bias towards market efficiency, that makes errors difficult to detect; if informed market participants believe that there are no errors, thus there is no point in looking for them. Errors in price determination would then be eliminated only when widely noticed across the market, and that is where investors would act in the way predicted by the efficient market hypothesis. The so-called "smart money" on which the EMH counts so heavily might then not be that prevalent in the market, overwhelmingly composed by irrational investors (Summers, 1986).

Market efficiency may also be way more difficult to attain than thought before. According to J. S. Jordan [12], if the set of all the possible economies is considered, and we define any economy as a finite set of investors, each of them characterized by a unique set of endowments, and an accompanying vector of unique securities, then the number of economies in which the efficient markets hypothesis is verified is negligible. In other words, an economy in which prices are fully reflective of all available and relevant information needs to satisfy some conditions which are extremely rare.

Said conditions are either (i) one or more investors to be risk neutral, which would equate the expected return on a security to its price, or (ii) all investors to be endowed with an identical amount of risk aversion, or (iii) all investors to be absolutely and constantly risk averse. Jordan's paper concludes that in an economy in which the number of sources of information is larger than the number of assets, equilibrium prices generally don't reflect all available information, except in the aforementioned very special cases, which are numerically insignificant with respect to the whole set of possible economies. Complete market efficiency, according to this view, is extremely rare (Jordan, 1983).

C. The weak-form Hypothesis

All of the above evidence has been produced against the most extreme form of the efficient market hypothesis, the strong form. However, while extremely interesting, studies tackling and taking down complete market efficiency are aiming at a mistaken target. Fama himself, in both of his papers, stated clearly that the strong form hypothesis was an extreme representation of reality, and it should not be taken as an exact representation of the real world (Fama, 1970) (Fama, 1991). The analysis then goes on to propose a weaker and "*more economically sensible*" version of market efficiency.

Indeed, the strong-form hypothesis is attacked mostly on grounds related to its reliance on information circulating freely and without friction among investors, who can then use it without incurring any cost. The weak form efficient market hypothesis instead states that for investors to act on information, the profit to be made by acting so must exceed the marginal costs (Fama, 1991). The critics to the strong form hypothesis then are failing to consider the costs related to operating on the market.

If the costs are considered, then almost every attack on complete market efficiency, while still valid, can ultimately be reconducted to the joint hypothesis problem, i.e., the lack of pricing models able to measure risk-adjusted expected returns also taking into account opportunity costs, transaction costs and the like. Since these costs exist, and are surely positive, then it is apparent that complete market efficiency must be nonetheless false in practice (Fama, 1991).

The weak form of the efficient market hypothesis then proposes to take the strong form efficiency model as a benchmark against which to measure spreads between reality and an efficient market. Admittedly, the market can make mistakes (Malkiel, 2003); there can be lags in the incorporation of information in prices due to investor underreaction, small but consistent autocorrelations in subsequent price changes have been observed in the very short term (Fama, 1991), and inefficiencies may arise and go unnoticed for some time.

Nonetheless, proponents of weak form market efficiency point out that the great difficulty in obtaining consistently greater than average risk adjusted returns has not been challenged, as well as the objective rarity of arbitrage opportunities. Empirical evidence in support of the weak form hypothesis remains abundant and solid (Fama, 1991) (Malkiel, 2003). Moreover, event studies show that the market does indeed adjust prices to new information rather quickly, and the existence of insider trading proves that private information is required to obtain significantly greater risk-adjusted returns than average (Malkiel, 2003).

The weak form hypothesis then states that in practice, markets are not perfectly efficient but still remarkably effective in portraying relevant information through the price mechanism.

Errors in market valuation and bubbles, although unequivocally ruling out complete market efficiency, still do not create portfolio trading opportunities capable of beating the market (Fama, 1991). Moreover, proponents of weak form efficiency urge to differentiate between statistical significance and economic significance of some instances of market inefficiency (Malkiel, 2003). The weak form hypothesis admits that prices are not fully transparent. This creates incentive for informed investors (smart money) to invest in producing quality information to make sensible investment decisions. Their actions in turn influence price transparency positively.

Reported persistent pricing errors may not be corrected because of their very modest size, or because of the costs associated with acting on information. Markets are ultimately taken as a social phenomenon, and as such prone to social dynamics such as bandwagon effects and various irrationalities linked to varying risk perception, linked pro-cyclically to business conditions and other exogenous variables. Moreover, being the stock market at the center of a great media attention, inefficiencies, and market failures such as bubbles may tend to be overemphasized (Malkiel, 2003).

In conclusion, the strong-form market efficiency still is valid as a benchmark against which to compare the actual conditions in capital markets. While there surely are real inefficiencies and autocorrelations in returns in the short and very short term, in the long term the market can be observed to be way more reliable on average (Malkiel, 2003). It is useful to evaluate how degrees of efficiency can vary when different time spans and exogenous conditions are considered.

Ultimately, three necessary conditions are outlined by Fama for a market to be efficient, i.e., where information "*obviously*" reflects all available information (Fama, 1970). These are to be taken as the three main parameters against which compare efficiency in empirical terms: (i) no transaction costs in trading securities, (ii) all relevant information is costlessly available to all market participants, and (iii) all agree on the implications of said information for current price and future price distributions of securities (Fama, 1970).

3. Market efficiency and machines

"The machine is only a tool after all, which can help humanity progress faster by taking some of the burdens of calculations and interpretations off its back. The task of the human brain remains what it has always been; that of discovering new data to be analyzed, and of devising new concepts to be tested." *- Isaac Asimov*

In this second chapter, the notion of High Frequency Trading is very briefly introduced. Its impact on market efficiency is then discussed, first considering the evidence of its positive impact on overall market efficiency through liquidity provision and information incorporation improvement, as argued by the cited studies; then, reasons for concern on the quality of HFT's contribution to market efficiency are presented, namely, the decoupling between informational efficiency and market efficiency, model risk and data quality.

A. A brief overview of algorithmic trading

Algorithmic trading is a now widespread method of operating in financial markets by large institutions. It involves the usage of very advanced artificial intelligence to not only carry out trades on the market automatically and at great speed, but also to devise investment strategies and often carry out their entire execution without human intervention. It is characterized by a very high turnover of capital and a low average return per trade (Aldridge, 2010) (McGowan, 2010).

The possible applications of Artificial intelligence in the financial industry are many. Today, much of the trading happening on the markets is automated (McGowan, 2010). Algorithms play a vital role in the functioning of the capital markets and are deeply embedded in much of their dynamics in the short and medium term (Yadav, 2015). What their effect could be in the very long term, however, is still up for debate. Because of their unrivaled efficiency at collecting information wherever available, and the volume of trades throughput, algorithms applied to financial securities trading are a cornerstone in the debate on market efficiency.

Artificial Intelligence's competitive advantage in the financial industry is based on its efficiency at collecting and processing information, and the speed at which it acts according to it. Large institutional investors use sophisticated systems of integrated software to engage in a wide range of activities, such as market probing and web scraping, and feed the information to the trading algorithm which finally formulates and executes an investment strategy according to said information (Yadav, 2015) (Zhang et al., 2016).

There are many nuances to the types of techniques implemented by algorithmic traders. They vary depending on their learning style and speed, as well as on the type of information that they employ to come up with investment strategies. The algorithm may act according to a pre-assembled data set to use as a reference or may be relying on a trial-and-error learning style, calibrated by a payoff system. Algorithms with different learning styles are optimal for different trading environments.

Artificial intelligence relying on pre-assembled instruction sets is characterized by a set of "rigid" rules programmed by the developer. Its learning styles can be broadly differentiated between Supervised Learning techniques (SL) and Unsupervised Learning techniques (UL). SL algorithms not only rely on a rigid set of rules to devise strategies, but also act on equally limited data environments, selected by the user. In the case of SL, both the training dataset (the dataset on which the algorithm's rules are tested and developed) and the input data are known. SL algorithms are best suited to input a defined set of labelled data (like a time series) and make reasonably precise predictions, such as mean reversion.

UL algorithms, although still relying on rigid programming, i.e., having a known training dataset, interact with unfiltered input data and operate without human intervention. UL algorithms are used for tasks like cluster analysis; they are "fed" with enormous amounts of unfiltered data and used to find hidden patterns, classes, and correlations among them. UL algorithms can deal with much more dynamic environments than their SL counterparts.

Artificial intelligence that does not rely on a precise set of rules is said to be employing a "Reinforcement Learning" technique. Through RL, an artificial intelligence learns how to carry out an activity by interacting with a real, dynamic environment with the goal of maximizing a payoff. The programmer then just defines the payoff metric, and the AI will develop its own set of decisions or routine to satisfy the payoff metric, through a process of trial-and-error. Reinforcement Learning requires high levels of computing power and programming expertise to be successfully carried out.

Algorithms are capable of operating at a wide range of speeds, or frequencies. The frequency at which an algorithm works varies according to the trading strategy employed by its developers, its machine learning technique, and the role that it plays in the market.

Low-Frequency Trading algorithms (LFT) work to devise investment strategies in the medium and long term. They help investors make more informed decisions, thanks to their superior efficiency in manipulating data, and can make investment decisions themselves at times. LFT operates at a slower speed, i.e., at a speed that a human can potentially keep up with; LFT algorithms are usually used to design strategies that take up weeks or even months and years. They contribute to a higher standard of efficiency by improving human traders' work, taking most heavy-duty computational tasks off their back, and decreasing brokerage costs, transaction costs and bid-ask spreads consistently by acting as market makers. SL and UL techniques are best suited to train algorithms that can carry out these tasks.

On the other hand, High-Frequency Trading algorithms (HFT) have redesigned the concept of market efficiency. HFT's revolutionary impact on market dynamics revolves around its capability of generating consistent returns by trading at extreme speed on microscopic price changes of a security in very short timespans. These returns are often accounting for mere fractions of a cent, but HFT firms benefit from the computer generating a large enough volume of tiny positive trades to add up to consistent profit (McGowan, 2010).

Thanks to its speed in submitting orders, the algorithm can generate returns while minimizing risks, by leaving the market quickly in case of downturns. The need for these machines to operate at this rate makes supervision by human traders very difficult and their maintenance costly (Yadav, 2015) (McGowan, 2010). HFT algorithms are based on cutting-edge technology, neural networks capable of adapting instantly to forecasted market conditions, carrying out immensely complex stochastic calculations to predict market movements a handful of seconds ahead.

Although HFT mostly uses Reinforcement Learning techniques, best suited to respond to an ever-changing environment, RL and other machine learning techniques are not mutually exclusive. UL techniques especially are very effective when used in combination with RL at high speed.

Much work has been done to determine whether algorithms' undeniable contribution to the market's functioning is for the best, or if their speed and efficiency might end up being too much, or too little, in the optics of market efficiency, i.e., if their superior informational efficiency will transfer in overall better market health and quality. The following analysis, although referring to algorithms working at any frequency, poses particular emphasis on the implications on market efficiency by HFT algorithms.

B. Implications of automated trading on market efficiency: the case for optimism

Recalling the conditions at which, in the words of Eugene Fama, "*the current price of a security obviously "fully reflects" all available information*" (Fama, 1970), as described at the end of the last chapter, i.e.: (i) no transaction costs in trading securities, (ii) all relevant information is costlessly available to all market participants, and (iii) all agree on the implications of said information for current price and future price distributions of securities, it is apparent how algorithmic trading constitutes a strong push towards getting closer to those conditions. The strategies that market participants implement by using artificial intelligence significantly enrich the set of information reflected in market prices, thus increasing market efficiency (Aldridge, 2010).

The most straightforward and wide positive effect on market efficiency from algorithmic trading, and especially HFT, is that it greatly increases market liquidity, exerting strong downwards pressure on transaction costs by minimizing latency times and reducing bid-ask spreads (Litzenberger et al., 2012) (Aldridge, 2010).

Algorithms today have important market making functions (Yadav, 2015), so they stand ready to buy and sell securities at any time. Algorithms are incredibly efficient as market makers; their superior speed compared to traditional market-making methods (human brokers) allows them to always keep securities available for trades instantly, thus providing other investors with very low agency and brokerage costs, providing liquidity, and favoring information price-in (Frino et al., 2016).

Fama himself wrote in his 1970 paper: *"With modern computers, it is hard to believe that a more competitive and economical system would not be feasible. It does not seem technologically impossible to replace the entire floor of the N.Y.S.E. with a computer, fed by many remote consoles, that kept all the books now kept by the specialists, that could easily make the entire book on any stock available to anybody (so that interested individuals could then compete to "make a market" in a stock) and that carried out transactions automatically."* (Fama, 1970). HFT algorithms' positive effects on liquidity help decrease market risk and protect investors from negative downturns by quickly closing the position, if necessary, thus cutting losses very short (Yadav, 2015). Decreasing liquidity risk and low transaction costs have been shown to increase firm value and improve market efficiency (Hatch et al., 2021).

Algorithms at any frequency are so effective as market makers because of their unparalleled efficiency at collecting and processing information; however, HFT makes the most intensive use of large amounts of data to keep ahead of competitors. Let us briefly discuss some of the strategies employed by AIs to gather relevant information on the market.

As mentioned, many strategies involve extensive web-scraping (Zhao, 2017) of public forums, news sites, and social media, to extract relevant keywords of which the algorithm then analyzes the tone and sentiment. Accurate sentiment analysis allows the machine to match a relevant event with the sentiment that it generates in the market and trade accordingly (Zhang et al., 2016). Another source of information for algorithms is the practice known as "pinging" (Yadav, 2015). The algorithm may send out dozens of "decoy" orders out in the market, to detect any large, potentially market-moving orders, to then cancel the "feeler" trades and alter its position ahead of the subsequent associated market trend. Pinging is especially popular among HFT AIs using Reinforcement Learning techniques, as it involves substantial speed in processing data and placing orders.

For example, the algorithm might detect widespread negative sentiment in the market following some bad news about a company's performance. It may then "feel" that large sell orders are coming up from major investors, driving the stock's price down. The algorithm would then immediately sell any related securities it holds to avoid potential losses; all of this happens in mere fractions of a second (Hatch et al., 2021) (Yadav, 2015).

It is apparent from just these two examples of automated trading strategy that algorithms' activity immensely decreases time lags in incorporation of new information in prices thanks to superior speed and liquidity provision (Ke et al., 2020). Superior computational capabilities and economies of scale reduce the marginal costs of gathering and acting on information dramatically (Aldridge, 2010). Recalling condition (ii) to benchmark against a hypothetical completely efficient market, it may not be an overstep to infer that, especially in the short run, algorithms take financial markets multiple notches up on the efficiency scale for the efficient pricing-in of information (Frino et al., 2016).

Computers' ability to produce so much information has implications on arbitrage as well. Algorithms are capable of scanning multiple markets at once looking for price discrepancies, and thus trade away even razor-thin inefficiency margins (Yadav, 2015). Algorithms employ a range of different arbitrage detection techniques, such as statistical arbitrage, i.e., basically detecting any slight anomaly or pattern in the variance of a security's price using staggering volumes of data, mostly mined from records of historical returns of a security. It rests on the basic assumption that whenever a variance anomaly occurs the price will (probably) mean-revert to the pattern (Aldridge, 2010).

Another arbitraging technique widely used by mainly HFT firms is event arbitrage, which differently from statistical arbitrage relies on data extraction from real-time feeds (news sites, forums, and the like). Event arbitrage works on even shorter time spans than statistical arbitrage and aims to exploit the tiniest irrationalities which follow ad-hoc relevant events such as monetary policy announcements or big economic news, by trading extremely quickly on the forecasted price momentum following the announcement (Aldridge, 2010).

The scope of algorithms' capability of removing arbitrage by trading is much wider than that of traditional "smart money". Thanks to extensive information availability, an algorithm can execute a "sweep" strategy, i.e., set a price, it buys all the available supply of a security that's sold at that price in every available market (Yadav, 2015). Once it determines that there is an arbitrage opportunity, the computer can thus eliminate it from all the markets at once.
The presence of algorithms and their computational and informational efficiency seems to overwhelmingly corroborate the efficient market hypothesis' assumption that arbitrage opportunities cannot exist for long and will self-destruct (Malkiel, 2003).

When constructing an algorithm, especially for trading at high frequency, developers have the interest of making it as precise as possible in evaluating a security's "correct" price given relevant information, to further minimize the lower bound in a spread's significance for profit (Hendershott et al., 2011). Therefore, algorithms are incentivized to use many different pricing methods at once on a security (Yadav, 2015). Algorithms have a better performance in collecting as well as evaluating new information. (Frino et al., 2016). This makes algorithms not only more efficient at determining the "equilibrium" price given all relevant information, but also mitigates the joint-hypothesis problem (Hendershott et al., 2011).

The probability of efficiency being masked by an appropriate pricing method is reduced when valuation of a security is carried out by many methods at once, all fine-tuned to be optimal according to the individual situation, often optimized by the minute. HFT has been found to decrease adverse selection problems and enhance the overall informative quality of prices (Hendershott et al., 2011). Algorithms in general price securities more efficiently than human traders, limiting overreaction or underreaction to news (Ke et al., 2020).

In conclusion, if the relationship between informational efficiency and overall market efficiency is indeed unequivocally positive, information systems applied to high-frequency security trading do play a major role in making the market more efficient, by greatly increasing market liquidity, thus lowering transaction costs (Hatch et al., 2021); moreover, they improve the incorporation of information in prices in the short run, in speed, scope and precision, through extensive production of relevant information, at a significantly lower cost with respect to human traders (Ke et al., 2020).

In addition, HFT algorithms perform arbitrage operations on a near microscopic scale, smoothing out even the smallest price divergences across markets (McGowan, 2010). All these factors considered, computers could indeed achieve, as foreshadowed by Fama, near-perfect efficiency in capital markets, or at least bring about a great increase in the markets' degree of efficiency and quality (Litzenberger et al., 2012) (Frino et al., 2016).

C. The case for pessimism

Although it seems apparent from available evidence that AI is a driving force for efficiency in capital markets especially in the short run, an alternative view is indeed argued for in relevant literature and challenges that optimistic opinion on theoretical and practical grounds. In fact, algorithmic trading could be said to have stripped even the weak form hypothesis from its empirical roots, and, in the long run, to be the driver for market instability rather than efficiency.

Diverse hints point in that direction: first, HFT itself builds its profitability from the "*persistence of the underlying tradable phenomena"* (Aldridge, 2010). Speed and low transaction costs, together with the low marginal cost of producing relevant information, allow HFT firms to exploit the recurrent autocorrelations in subsequent returns already identified by Fama: *"The lagged price change can be used to explain about .36% of the variation in the current price change"* (Fama, 1970).

The Author then goes on to point out that this result is not economically significant, as correlations are so small and repentine that it would be extremely difficult to exploit them, also considering transaction costs: "*In particular, it is unlikely that the small absolute levels of serial correlation that are always observed can be used as the basis of substantially profitable trading systems."* (Fama, 1970); but time has proved that HFT would instead be able to rely on precisely these "*small absolute levels of serial correlation*" to generate consistent profits (Aldridge, 2010).

In this key, it could be argued that HFT finds its reason for existence precisely in the failure of the efficient market hypothesis, and this also has consequences on the corollary of the unpredictability of price changes that necessarily follows from the assumption itself.

Even though the random walk assumption, according to which price changes are as random as the news that cause them (Fama, 1970) was already redimensioned by Fama himself in his second article on the efficient market hypothesis, in which it is highlighted that returns are, in some measure, predictable from relevant fundamentals (Fama, 1991), the joint-hypothesis problem is still brought up against the findings' direct challenge to the efficient market hypothesis' validity.

In the Author's view, predictability in returns could then be a by-product of an incorrect pricing approach (Fama, 1991); however, as mentioned above, the wide range of information used simultaneously by algorithms when evaluating trading opportunities, while not completely vanquishing these doubts, makes pricing inaccuracies less likely (Litzenberger et al., 2012) (Ke et al., 2020).

Ultimately, there is experimental evidence to date that it is possible to effectively predict price changes in the market: Skabar and Cloete [20] show that a neural network using data from historical returns could effectively generate consistent risk-adjusted returns on the observed price time-series of the NASDAQ, while it did not perform as well when applied on actual random price time-series (Skabar, Cloete, 2001).

Techniques used by algorithms to act so efficiently as market makers rely on price change prediction methods like momentum forecasting (Litzenberger et al. 2012), linked to event arbitrage as its means of execution (Aldridge, 2010); and the most advanced algorithms can forecast market movements up to four seconds ahead (Yadav, 2015), hence the seemingly flawless market making.

All the information mining, and all the latency minimizing, liquidity-providing techniques serve the goal of generating greater than average risk-adjusted returns, often outperforming the naive buy-and-hold strategy, and factually put into question the integrity of even the weak-form efficient market hypothesis. A goal that HFT firms seemingly meet on the regular (Yadav, 2015).

Apart from the theoretical challenges to market efficiency posed by AI, Yesha Yadav [26] poses a series of convincing arguments in the sense that the extensive usage of artificial intelligence in the market could in fact hinder market efficiency from an allocative perspective, considering that it could end up having a reverse effect on information incorporation and price transparency than previously argued (Yadav, 2015).

The Author observes that an increase in informational efficiency might not be necessarily linked to a proportional increase in overall market efficiency. This view revolves around the notion of "model risk", a new type of systemic risk that proliferates in the market the more algorithms are involved in trading (Yadav, 2015). The introduction of model risk in the analysis of Artificial Intelligence in general, and especially HFT's impact on market efficiency could be determinant, as it could potentially decouple the traditionally positive relationship between informational efficiency and market quality as proposed by optimists and examined in the previous sections.

Model risk can arise from several sources. Firstly, it must be acknowledged that models are not, by definition, a perfect depiction of reality (Yadav, 2015). Moreover, algorithms rely on a few pre-set strategies and assumptions that tell the computer how to evaluate a security's pricing, when and what to trade, and for how long (Aldridge, 2010) (Yadav, 2015). Essentially, the algorithm relies on pre-set strategies as metaphorical "rails" to follow and carry out successful trades.

These models are extremely complicated and expensive to develop. They require many man-hours and expertise to fine-tune and optimize, and often require continuous maintenance to monitor performance and prevent the threat of reverse engineering (Yadav, 2015). Programmers then face steep trade-offs between keeping these models general enough to guarantee a decent useful life and thus a reasonable margin of cost-effectiveness, and yet precise enough to be in the conditions to output reliable predictions (Yadav, 2015).

Incorrect assumptions and overly generalized approaches increase the probability of trading errors and information loss (Yadav, 2015). Algorithms might sacrifice information that human traders usually would consider in the name of speed and efficiency, as having more constraints to consider might mean losing precious milliseconds for the computer, with the subsequent loss in profit (Yadav, 2015) (Aldridge, 2010).

Model risk might then reverse the causality between HFT intensity and informational efficiency. According to this view, the decrease in time lags for information incorporation is not due to a raw increase in informational efficiency, but to a loss in complexity and completeness of prices (Yadav, 2015). Hence excessive reduction of noise and perceived increase in price transparency could harm the price mechanism's reliability when portraying relevant information (Yadav, 2015).

Another important issue about adequacy of the HFT industry is the complexity and expensiveness of developing HFT algorithms, which greatly increases the incentive for copying software developed by rival firms.

Reverse engineering is a real threat that HFT firms must deal with on a very frequent basis and considering that algorithms must be changed and updated by the hour if not by the minute (Yadav, 2015), such frequent rearrangement of complex systems raises the probability of programming errors (Aldridge, 2010).

Thus, if a faulty algorithm routine is reverse engineered, the fault could spread like wildfire before any human supervisor is able to spot and correct it. This could possibly have serious effects on systemic risk, short- and long-term market volatility, and ultimately market efficiency (Yadav, 2015).

Model risk is at the root of another relevant tradeoff when evaluating HFT's impact on the soundness of the financial system; stochastic mathematical methods for predicting returns are in fact extremely efficient in the short run but lose any effectiveness in the long run (Aldridge, 2010) (Yadav, 2015). Programmers and investors are incentivized to create algorithms to trade in the areas in which they are most effective, hence the bulk of HFT happens, by definition, in the short-term market (McGowan, 2010) (Yadav, 2015).

Most of the advantages in informational efficiency are indeed registered in the very short term (Ke et al., 2020). The proliferation of systemic risk linked to faulty models could then be hardly detectable in the short run, but make itself apparent later, adding excess volatility in the long run (Yadav, 2015).

This argument raises doubts on the quality of what is regarded as HFT's main contribution to market efficiency, i.e., liquidity provision. Algorithms are hardwired to cut losses short by immediately leaving the market, i.e., closing the position, when they forecast an imminent negative price change (Aldridge, 2010). This is beneficial to HFT users but can be highly detrimental to market efficiency (Yadav, 2015).

Indeed, it has been observed empirically that HFT algorithms' liquidity providing activity is highly procyclical, i.e., works well in "good times", periods of market growth and activity, but breaks down during “bad times”, like market downturns, or worse, crashes, depriving the already struggling markets of the liquidity they have been accustomed to and widely rely on (Yadav, 2015).

On May 6, 2010, the EU sovereign debt crisis was sending its shockwaves all the way to the U.S. financial markets. Hedging instruments against Greek sovereign debt default were in high demand, and euro to dollar exchange rates were in decline. Negative sentiment was prevalent in the market; liquidity was falling, and volatility indexes were unusually high even for such a bearish market (SEC, 2010).

At 14:32, U.S. time, a risk management system belonging to a large mutual fund complex, sensing a surge in volatility, initiated a “sell algorithm” to execute a trade involving the sale of futures contracts against a large equity position it held. The net value of contracts to be sold was around $4.1 billion (SEC, 2010). The sell algorithm was not programmed to consider price or time when executing the sale, but to only target trading volume. It would have taken around five hours to execute the trade, if the algorithm were made aware of price and time, as it had happened in previous instances; on May 6, it was completed in 20 minutes (SEC, 2010).

This very large order was crammed in a market already under stress. HFT algorithms, following their usual trading style, built up large temporary positions on the futures, while simultaneously selling the same contracts to offset their position in the opposite direction; HFT accounted for more than 33% of the total trading volume of that security. At the same time, arbitrageurs transmitted the pressure from the derivatives market to the equity markets by buying futures and selling stocks indexes at the same time (SEC, 2010). In response to the large trading volume, the sell algorithm increased the futures selling rate even more. All of this happened in the span of approximately 3 minutes (from 14:41 to 14:44) (SEC, 2010).

This caused a widespread liquidity crisis across equity and futures markets alike. On May 6, 2010, the Dow Jones dropped more than 900 points in the span of minutes, to rebound only a short while later (Yadav, 2015). Lacking any buyers for the futures they wanted to sell, HFT algorithms began frantically buying and selling the securities to each other; arbitrageurs kept on applying sell pressure on the securities markets. Both markets lost 3% of their value in the 14:41 - 14:44 timespan (SEC, 2010). Liquidity in the E-MINI was around 1% of that morning’s level at that time (SEC, 2010); this vicious circle would make it lose an additional 1.5% in 15 seconds. Similarly, the stock indexes only retained about 25% of their liquidity. It was 14:45.

As market participants had time to get up to speed on what was happening, also thanks to the triggering of a “stop logic function” switch that allowed the markets to unwind, the markets resumed functioning properly and returned to price levels consistent with rational behavior (SEC, 2010). The catastrophic fall, and the equally fast rebound that followed on May 6 were called the “Flash crash” (Yadav, 2015) (SEC, 2010).

The “Flash Crash” is an example of how model risk has real and heavy implications on market risk, and overall financial stability. The SEC report from September 30, 2010, assesses that “*under stressed market conditions, the automated execution of a large sell order can trigger extreme price movements, especially if the automated execution algorithm does not take prices into account. Moreover, the interaction between automated execution programs and algorithmic trading strategies can quickly erode liquidity and result in disorderly markets”* (SEC, 2010). Algorithmic traders exacerbated the liquidity crisis by leaving the market instead of keeping it liquid, overloaded the system with irrational orders, and spread increased volatility beyond the scope of just the futures market by chasing arbitrage.

The “flash crash” also raises questions on such algorithms’ ability to tackle unexpected market crises, such as the large, unforeseen negative occurrences known as “black swans” [6]. The tradeoff at the basis of model risk also makes algorithmic HF routines ill-equipped to deal with this type of events. Endowing an artificial intelligence with the ability to predict such occurrences would require enormous investments in time and capital, let alone huge amounts of computing power, just to develop.

Such an AI would then need to be as fast as other competing algorithms to function and be profitable on the market. It is apparent then that it makes little economic sense to provide HFT algorithms with such capability. An algorithm-intensive capital market such as the present one is therefore greatly exposed to such risks (Yadav, 2015).

Moreover, the high barriers to entry in the HFT industry effectively make these methods available to a selected group of large investors with a lot of capital to invest in the development and maintenance of the whole apparatus required to make an HFT firm function. The small dimension of this circle implies that the greater than average risk adjusted returns coming from ultra-fast arbitrage are effectively exploited by a small minority that will then systematically outperform the Buy and hold.

A study by Gianluca Virgilio [23] shows that a small number of traders with high-frequency tactics were able to consistently outperform the low-frequency traders and the buy and hold strategy in the long term. If it is assumed that a significant part of the market cannot exploit this technology the market then results in a fractured entity which would not reflect information into prices efficiently, just like a broken glass or a shattered lens; thus, further damaging the efficient markets hypothesis position (Virgilio, 2015).

In conclusion, when considering model risk, the direct causality identified up until now between informational efficiency and overall market efficiency could be decoupled. Indeed, the very same characteristics that would make HFT a force for efficiency could in the end be a threat to the financial system's stability and soundness.

D. The challenge to "smart money", the question of information quality, and other issues

Traditionally, throughout market efficiency literature, a major role is played by informed investors. The market relies on informed investors to produce high quality relevant information to trade on, and thus infuse it in the market (Fama, 1970) (Malkiel, 2003) (Yadav, 2015); they have an incentive at producing private information by the possibility of enjoying first mover advantage on potentially profitable trades, or by profiting from trading away arbitrage opportunities (Yadav, 2015). Informed trading constitutes the backbone of financial markets efficiency.

HFT constitutes a challenge for informed traders. Algorithms detect profitable market-making opportunities through order anticipation techniques (Yadav, 2015). By sending out "probe" trades, an algorithm can obtain intelligence on whether a large order or series of orders is taking place on a security, or scan other markets and imitate any relevant positioning by informed traders (Aldridge, 2010).

Similarly, HTF algorithms have the ability of detecting and exploiting arbitrage more quickly and on a much larger scope than traditional informed traders. The computer will then beat informed investors to the trade thus reaping a substantial amount of their payoff, disincentivizing private production of information (Yadav, 2015).

Informed traders are therefore competing against an opponent that can systematically free ride on the information they produce and beat them to any trade they think of doing. They are thus less interested in actively participating in the market. The gradual phase-out of traditional "smart money" could expose financial markets to instability and excess volatility in the long run due to proliferation of model risk (Yadav, 2015).

As mentioned above, there is little to no incentive for developers to invest in creating algorithms capable of calculating long-term market dynamics. The phasing out of human informed traders could therefore further concentrate trading in the very short term (Yadav, 2015), increasing long-run volatility, market procyclicality, model risk buildup, and ultimately financial stability (see the Flash Crash).

The question of the quality of the information that algorithms base their decisions on is too a pressing one. Indeed, computers cannot critically evaluate data, but only check them against their models, which are, as mentioned, subject to model risk (Yadav, 2015). The Flash crash of 2010 has shown how repentine and potentially disastrous the consequences on information loss can be on the price mechanism.

Information loss refers to the decrease in precision of prices' portrayal of information caused by widespread usage of models to decode information and draft trading decisions (Yadav, 2015).

Algorithms must evaluate, model, and devise an effective strategy based on enormous quantities of data in the span of microseconds. Accurate depiction of all that information on prices is impossible, considering that the computer will make a significant skimming of the relevant data, to minimize latency times and risk (Yadav, 2015). The reduction in time lags and noisy signals in prices could therefore not be just thanks to HFT's superior processing speed, but also to loss of complexity and relevant information.

Price transparency could suffer from excessive use of models as well. As mentioned above, HFT firms are engaged in a constant "arms race" to prevent reverse engineering and optimize their routines to the millisecond. Developers have overwhelming incentives to keep their models secret; that, together with the inherent complexity of said models, makes knowing exactly what information is used to determine prices, and how, almost impossible (Yadav, 2015).

The question of information quality and model risk brings back the risk of the market being enveloped in a spiral of self-fulfilling prophecies in the case of overwhelming automation of the informed investment component of the market. Indeed, even if just one faulty routine or low-quality data input brings an algorithm to the decision of closing a large position on a stock, other computers, hardwired to be highly reactive to large orders, will pick up on the incoming large sellout and sell their position as well to cut losses short. Other algorithms based on similar models to the faulty one, and possibly identical machines reverse engineered from those same algorithms or similar ones, are likely to amplify that effect.

It follows that the entire market would pick up on the new price as if it was the equilibrium one, in extremely small timespans. Therefore, model risk and excessive noise reduction in prices for the sake of incorporation speed, could lock the market in a spiral of self-fulfilling prophecies, where algorithms in a constant race to anticipate each other on an increasingly high order strip prices from any meaning.

We could recall this scenario as exactly the one envisioned by Keynes and Samuelson when critiquing the notion of perfect price transparency in the strong form efficient market hypothesis. Total informational efficiency, if based on low-quality data, which would be further impoverished by information loss, could indeed establish a market that "lives in its own dreams" and determines prices by self-fulfilling prophecies, with the bulk of its activity concentrated in the very short run, with high levels of volatility, market risk and systemic risk.

4. Conclusions

HFT's impact on market efficiency must be evaluated along two axes. First, it must be established if the theory that underpins the concept of market efficiency, is still conceptually valid considering technological breakthroughs; then, one must examine and interpret the empirical evidence available to examine market efficiency in recent times.

In this case, the same arguments used to further the hypothesis that HFT improves efficiency in financial markets, and does not collide excessively with the underlying assumptions, can also be used to affirm the contrary, i.e., artificial intelligence might prove detrimental to market efficiency in the long run.

The drawing of either of these two conclusions depends on the relationship between informational efficiency, i.e., the sheer speed and volume of the information being incorporated in prices, and overall market quality, understood as not only transparency, but precision, accuracy, and stability of the price system as a whole.

If indeed it is established, as some of the evidence brought to attention in this paper suggests, that there is a direct and positive relationship between informational efficiency and market quality, then HFT's contribution to market efficiency is straightforwardly positive, as it increases market liquidity, therefore mitigating market risk; enriches information sets, thus adding depth to prices, while improving their transparency as well; all while decreasing time lags and eliminating arbitrage faster than ever before.

On the other hand, if informational efficiency and market quality are decoupled, and one considers the issues of model risk and data quality, then the advantages of automation in securities trading might not be so apparent. Excessive model optimization can lead to pricing errors, information loss, self-fulfilling prophecies, fire sales, excess procyclicality and systemic risk buildup.

With the markets increasingly complex and fragmented, then, does it still make sense to talk about "market efficiency"? Has the hiatus between short-term money markets and long-term capital markets widened too much for the efficient market hypothesis to survive?

The answer is no. The efficient market hypothesis was not deemed “the best-established empirical fact in economics” for nothing. Keynes and Samuelson’s views, although captivating, also due to their high relevance in the face of recent developments in technology, are useful to keep in mind a “negative” benchmark to market efficiency, just like the strong form market efficiency hypothesis constitutes a positive one.

Even if the advent of HFT might have dealt some theoretical and empirical blows to the classical notion of market efficiency, the old theory still survives behind the joint-hypothesis problem. As widely mentioned above, trading AIs are extremely difficult and expensive to develop, maintain, and operate. The costs associated with just keeping these computers active might deflate the significance of the observed risk-adjusted returns obtained on the market by these firms. The great practical and economical difficulty in outperforming the Buy and hold still exists.

Being a relatively young field employing cutting-edge technology, the HFT industry is arguably still in the rising phase of the technology adoption bell curve [9]. Although undeniably influential, we still don’t know where this technology might peak, and any of its implications might be overstated by the hype surrounding it.

Moreover, being a young industry, the HFT environment is subject to consistent disruption by constant innovation. There are real incentives for developers to actively prevent reverse engineering and imitation of their programs. This creates other steep costs for HFT operators and might mitigate the risk of fire sales in the long run.

Time will tell if Keynes was right after all, when he compared the stock market to a casino, or if Eugene Fama’s predictions on automated systems being ultimately a strong driving force for efficiency in financial markets were correct instead.

5. References

1. Aldridge, I. (2010). *High Frequency Trading*. John Wiley & Sons, Inc., Hoboken, New Jersey.

2. Alex Frino, Tina Prodromou, George H.K. Wang, P. Joakim Westerholm, Hui Zheng,An empirical analysis of algorithmic trading around earnings announcements,Pacific-Basin Finance Journal,Volume 45,2017,Pages 34-51,ISSN 0927538X, https://doi.org/10.1016/j.pacfin.2016.05.008. (https://www.sciencedirect.com/science/article/pii/S0927538X16300531)

3. Andersen, T. M. (1983). Some Implications of the Efficient Capital Market Hypothesis. *Journal of Post Keynesian Economics*, *6*(2), 281–294. http://www.jstor.org/stable/4537813

4. BBC, (n. d.). *Trader Made Error in "Flash Crash", Citigroup says*. https://www.bbc.com/news/business-61303217

5. Brian C. Hatch, Shane A. Johnson, Qin Emma Wang, Jun Zhang,Algorithmic trading and firm value,Journal of Banking & Finance,Volume 125,2021,106090,ISSN 0378-4266,https://doi.org/10.1016/j.jbankfin.2021.106090.(https://www.sciencedirect.com/science/article/pii/S0378426621000480)

6. Corporate Finance Institute, (n. d.). *Black Swan Events*. https://corporatefinanceinstitute.com/resources/knowledge/finance/black-swan-event/

7. Eugene F. Fama. (1991). Efficient Capital Markets: II. *The Journal of Finance*, *46*(5), 1575–1617. https://doi.org/10.2307/2328565

8. Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, *25*(2), 383–417. https://doi.org/10.2307/2325486
9. Gartner Research Methodologies, (n. d.). *Gartner Hype Cycle*. https://www.gartner.com/en/research/methodologies/gartner-hype-cycle
10. HENDERSHOTT, T., JONES, C. M., & MENKVELD, A. J. (2011). Does Algorithmic Trading Improve Liquidity? *The Journal of Finance*, *66*(1), 1–33. http://www.jstor.org/stable/29789770
11. Ivan L. Lustig, & Philip A. Leinbach. (1983). The Small Firm Effect. *Financial Analysts Journal*, *39*(3), 46–49. http://www.jstor.org/stable/4478645
12. Jordan, J. S. (1983). On the Efficient Markets Hypothesis. *Econometrica*, *51*(5), 1325–1343. https://doi.org/10.2307/1912277
13. Ke, Yun, and Yanan Zhang. "Does High-Frequency Trading Reduce Market Underreaction to Earnings News?" Finance research letters. 34 (2020)
14. Keynes, J. M. (2013). Teoria generale dell'occupazione, dell'interesse e della moneta. Italia: UTET.
15. Litzenberger, R., Castura, J., & Gorelick, R. (2012). The Impacts of Automation and High Frequency Trading on Market Quality. *Annual Review of Financial Economics*, *4*, 59–98. http://www.jstor.org/stable/42940378
16. Malkiel, B. G. (2003). The Efficient Market Hypothesis and Its Critics. *The Journal of Economic Perspectives*, *17*(1), 59–82. http://www.jstor.org/stable/3216840
17. Marc R. Reinganum,The anomalous stock market behavior of small firms in January: Empirical tests for tax-loss selling effects,Journal of Financial Economics,Volume 12, Issue 1,1983,Pages 89-104,ISSN 0304-

405X,https://doi.org/10.1016/0304-405X(83)90029
(https://www.sciencedirect.com/science/article/pii/0304405X83900296)

18. Michael J. McGowan, The Rise of Computerized High Frequency Trading: Use and Controversy, 9 *Duke Law & Technology Review* 1-25 (2010)
19. Ray Ball and Phillip Brown. "An Empirical Evaluation of Accounting Income Numbers." Journal of Accounting Research, 6 (Autumn, 1968), 159-78.
20. Skabar, Andrew & Cloete, Ian. (2002). Neural Networks, Financial Trading and the Efficient Markets Hypothesis. ACSC '02: Proceedings of the Twenty-fifth Australasian Conference on Computer Science. 4. 10.1145/563857.563829.
21. Summers, L. H. (1986). Does the Stock Market Rationally Reflect Fundamental Values? *The Journal of Finance*, *41*(3), 591–601. https://doi.org/10.2307/2328487
22. U.S. Securities and Exchange Commission, U.S. Commodity Futures Trading Commission. (2010). *Findings Regarding the Market Events of May 6, 2010; Report of the staff of CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.* https://www.sec.gov/news/studies/2010/marketevents-report.pdf
23. Virgilio, Gianluca. (2015). High-frequency trading and the efficient market hypothesis.
24. Werner F. M. De Bondt, & Thaler, R. (1985). Does the Stock Market Overreact? *The Journal of Finance*, *40*(3), 793–805. https://doi.org/10.2307/2327804
25. William F. Sharpe. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." Journal of Finance, 19 (September, 1964), 425-42.

26. Yesha Yadav, How Algorithmic Trading Undermines Efficiency in Capital Markets, 68 Vanderbilt Law Review 1607 (2015) Available at: https://scholarship.law.vanderbilt.edu/vlr/vol68/iss6/3

27. Zhang, J. L., Härdle, W. K., Chen, C. Y., & Bommes, E. (2016). Distillation of News Flow Into Analysis of Stock Reactions. *Journal of Business & Economic Statistics*, *34*(4), 547–563. http://www.jstor.org/stable/44166589