J Bus Ethics (2013) 112:585–595
DOI 10.1007/s10551-012-1559-0

Fairness in Financial Markets: The Case of High Frequency Trading

James J. Angel · Douglas McCabe

Published online: 5 December 2012

Abstract Recent concern over "high frequency trading" (HFT) has called into question the fairness of the practice. What does it mean for a financial market to be "fair"? We first examine how high frequency trading is actually used. High frequency traders often implement traditional beneficial strategies such as market making and arbitrage, although computers can also be used for manipulative strategies as well. We then examine different notions of fairness. Procedural fairness can be viewed from the perspective of equal opportunity, in which all market participants are treated alike. The same rules apply to HFT as to other traders. Another approach to fairness is in the equality of outcomes. Many HFT strategies are beneficial to other market participants, so one cannot categorically denounce the practice as unfair. Other strategies, for both high and low frequency trading, are not. It is thus important to distinguish between the technology and the use of the technology to make judgments on fairness.

Keywords Fairness · Justice · High frequency trading · Financial markets · Manipulation

Introduction

Recently, there has been a great deal of concern about the fairness of many features of our financial markets. In particular, a great deal of concern has recently been raised about the use of computers to trade at high frequency in our financial markets. Do these computers give some investors an unfair advantage over other investors?

Fairness is an important consideration in our financial markets. Indeed, the words "fair," "unfair" or "fairness" are mentioned 130 times in the recently passed Dodd-Frank Wall Street Reform and Consumer Protection Act.[1] But what exactly does it mean for markets to be "fair"? Our regulators are also examining this. The U.S. Securities and Exchange Commission recently requested comment on the fairness of various developments in the equity markets. This article examines the notion of financial market fairness in the context of the debate over so-called "high frequency trading" (HFT), the use of computers to trade very quickly and at high speed. Is HFT unfair?

In order to address this question, we must first have an understanding of exactly what investors are doing with HFT. The next section examines HFT and describes in detail what many HFT strategies attempt to do. Many of these activities are actually beneficial to the rest of society. Some are not. We then discuss some of the many different notions of fairness, some basic and some more esoteric. When we want our financial markets to be fair, do we mean good looking, tolerable, or free of moral blemish? Some view fairness as procedural fairness, in that equal rules apply to all participants. Others concentrate on distributional fairness and examine the outputs, rather than the inputs. The final section concludes. Fast computers can be used both fairly and unfairly. It is not the speed of the tool that matters for fairness, but what is done with it.

J. J. Angel (✉) · D. McCabe
McDonough School of Business, Washington, DC 20057, USA
e-mail: angelj@georgetown.edu

D. McCabe
e-mail: mccabed@georgetown.edu

[1] Public Law 111-203, HR 4173 http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=111_cong_public_laws&docid=f:publ203.111.pdf.

High Frequency Trading

The use of computers to trade at high speeds has recently become a controversial topic of public debate. The New York Times (Duhigg 1999) reported in a front page story that

> Powerful computers, some housed right next to the machines that drive marketplaces like the New York Stock Exchange, enable high-frequency traders to transmit millions of orders at lightning speed and, their detractors contend, reap billions at everyone else's expense.

CBS News (2010) recently did a *Sixty Minutes* piece on HFT in which reporter Steve Kroft announced, "There are a lot of people out there who think that the stock market is rigged." Politicians and regulators have taken noticed and asked whether this is fair. The SEC (2010a) requested public comment on many items related to the fairness of high-frequency trading:

> In addition, what standards should the Commission apply in assessing the fairness of the equity markets? For example, is it unfair for market participants to obtain a competitive advantage by investing in technology and human resources that enable them to trade more effectively and profitably than others?[2]

HFT is a major factor in the U.S. equity trading. Brogaard (2010) reports that the HFT traders in his sample were responsible for 68.5 % of trading volume. Although these traders trade in large quantity, their profit margins are generally very thin. In the Brogaard (2010) study, they earned .72 cents for every $100 traded, which is approximately one tenth of a penny per share traded.

In order to determine the fairness of HFT, we must first understand what it is. There are a wide variety of different trading strategies that can be implemented with high frequency technology. In most of these cases, the strategies are not new, but merely old strategies that use fast computer technology, using computer programs called "algorithms." Traders are still trying to solve the same basic problems they have always been trying to solve. Investors desiring to establish or liquidate a position attempt to do so while minimizing their trading costs. Other traders seek to profit from short-term changes in price. Here are some of the strategies used:

Market Making

Buy orders and sell orders do not arrive in markets at exactly the same moment. An investor who wants to trade immediately may not find an acceptable (dare we say "fair"?) price and be forced to wait for an acceptable price. This causes the investor to sit on the risk of the position longer than the investor wants. Most markets rely on some market participants to act as market makers to smooth out trading. The business model of a market maker is similar to that of a shopkeeper or a used car dealer. A used car dealer buys cars at a low trade-in price from customers who want to sell, and then sells at a higher retail price to customers who want to buy. The car dealer does not want to drive the car home and keep it. Similarly, a market maker stands ready to buy at the low bid price from customers who want to sell, and sell at the slightly higher ask or offer price to customers who want to buy. Just as used car dealers make a profit from the difference between the trade-in and retail prices, market makers profit from the difference between the bid and offer prices, known as the bid-ask spread. Their business model is to make a small profit on a large number of trades, and they generally do not like to hold large positions in any given stock. Just as the car dealer does not want to keep a car for a long period of time, market makers do not want to hold stocks in their inventory for long periods of time, either.

The practice of market making benefits markets because it usually stabilizes short-term prices by eliminating the "air pockets" that would otherwise occur because a customer's buy order usually does not arrive in the market at the same moment as a sell order. The willingness of a market maker to buy and sell at all times provides immediacy to other investors who want to buy and sell. Competition between market makers reduces the spread between bid and ask prices to a competitive level.

Market making was traditionally done by firms acting as NYSE specialists and NASDAQ market makers. The NYSE now calls its market makers "designated market makers" (DMMs). Whereas the bulk of market making used to involve human interaction, both NYSE DMMs and NASDAQ market making firms use high frequency technology to maintain quotes in the markets and to update those quotes. Note that whenever market conditions change, the updating of a quote looks like the cancellation of an order. As market conditions change rapidly, firms following a market making model must enter and cancel large number of orders rapidly.

Here is an example of how market making helps investors. Suppose that the best buy order from a long-term investor who really wants to own the stock is $10.00 and the best sell order from a long-term shareholder who wants to exit their position is $10.10. In other words, there exists a potential buyer who refuses to pay more than $10, and a seller who wouldnot accept less than $10.10. A market maker who has no position in the stock (and who doesnot really want one) is willing to quote a bid

[2] SEC (2010a, pp. 41).

price at which he or she is willing to buy of $10.04 and an offer price at which he or she is willing to sell for $10.06. When another long-term shareholder comes into sell shares at the market bid price, the market maker buys it at $10.04. Later, another would-be long-term investor arrives who is willing to buy at the current offer price, and the market maker sells at $10.06 for a two cent profit. Note that both the buyer and the seller got better prices than they would otherwise have gotten: Without the market maker, the seller would have received only $10.00 and the buyer would have paid $10.10. Furthermore, there has been less volatility in the price as well: Instead of the price bouncing from $10 to $10.10, its range was reduced to $10.04 to $10.06.

Some proprietary trading firms provide intense competition for NYSE DMMs and NASDAQ market makers with HFT strategies that are fundamentally market making strategies. This competition benefits investors by reducing bid-ask spreads and increasing the number of shares (depth) that investors can trade at any given moment without moving the price. This makes markets more "liquid," which means that investors can buy or sell larger quantities without moving the price.

Arbitrage

Many financial instruments are economically related to one another. When the price of one instrument gets out of line from its economic relationship to another instrument, it is possible for an arbitrageur to make money by selling the cheaper asset and purchasing the more expensive one. This puts upward pressure on the cheaper asset and downward pressure on the more expensive asset, pushing them back into their proper alignment.

Here is an example: Suppose that the price of an exchange traded fund (ETF) that contains a portfolio of all 500 stocks in the S&P500 is currently $100.00 bid and $100.01 offered per share. However, the value of the 500 stocks in the ETF portfolio is only $99.87 bid $99.90 offered. Retail investors who purchase the ETF at its current market price would be paying more than the current value of the stocks inside the ETF. In other words, they would be paying too much.

A trader hoping to profit from this discrepancy could just buy all 500 stocks, thinking they are underpriced. This is a risky strategy, because the whole stock market could go down as well as up, so there is a lot of risk involved. Alternatively, the trader could borrow the overvalued shares of the ETF from another investor and sell them for $100 per share. This is commonly called a short sale. However, the overall market may be going up, so this would also be a risky strategy. A safer alternative is to do both at the same time, which is known as arbitrage: short the ETF at $100 while simultaneously buying the 500 constituent stocks at a price of $99.90. In this way, the arbitrageur is hedged against movements in the overall market prices, and is merely betting that the difference, or basis, between the ETF and the constituent stocks will get smaller.

Note that this trading will push the price of the ETF downward and the price of the constituent stocks upward, reducing the difference (known as the "basis") between the ETF and its constituent stocks. Arbitrageurs will keep doing this trade until the basis is less than their transactions costs. Suppose that during this time the overall market has gone up, so that the ETF is now $101 bid/$101.01 offered, and the constituent stocks have also gone up to 100.99 bid/$101.02 offered. The arbitrageur then trades out of the position by purchasing the ETF at $101.01 and selling the constituent stocks at $100.99. The arbitrageur lost (101.01–100.00) = $1.01 on the ETF, but made ($101.02–99.90) = $1.12 on the underlying stocks, for a net profit of 1.12–1.01 = 11 cents.

This type of activity indisputably makes markets more fair to the retail investors who invest through ETFs. It means that the prices they get when they buy or sell ETFs will very closely track the underlying value of the shares inside the ETF. If there was an absence of this type of arbitrage, then the prices of the ETFs could well deviate extremely from the prices of the stocks inside the ETF.

Such arbitrage opportunities can occur between any set of related financial instruments, such as between stocks and their options and futures, between American Depositary Receipts (ADRs) and their foreign ordinary shares, and between ETFs and their constituents. Because the strategy is so simple, such opportunities tend to disappear quickly. For this reason, arbitrageurs need to use high frequency technology to respond as quickly as possible. The arbitrageurs race against each other to take advantage of profitable trading strategies before they disappear.

Pairs Trading and Statistical Arbitrage

There are many financial instruments that are economically related even though there is no strict arbitrage relationship between them. However, their prices do tend to go up and down together, and when their prices diverge there are risky profit opportunities to investors. For example, The Coca Cola Company (KO) and Pepsico (PEP) are both diversified beverage and food companies. As seen from the following chart, their prices tend to go up and down together during a trading day.

When the prices start to diverge, a trader following a "pairs trading" strategy at 10:50 would seek to short KO and purchase PEP and then reverse the position at 1:50. Of course, there is no guarantee that the prices will always converge.[3] Sometimes there is news that affects one firm but not the other. Thus, this strategy will sometimes incur losses. Pairs trading helps to keep the prices of related stocks in their proper alignment. When noise moves the prices apart, the pairs traders move them back into position.

Pairs trading is a form of "statistical arbitrage," or "stat arb." Traders can look beyond pairs of stocks and identify groups of related instruments that tend to move together. For example, they may keep the stock price of a company in line with several of its key suppliers.

News Reaction Strategies

It is no secret that news moves stock prices. When a company makes a major announcement or some other news event occurs, the market consensus about the value of the stock will change. When unexpected good news comes out, the price usually rises. Likewise, the stock usually falls when unexpected bad news comes out. The price gets from the old price to the new price through trading. Investors who process the information in news announcements are part of the information processing in markets that determines prices.

Since the beginning of stock trading, markets have responded quickly to news, and investors have devoted substantial resources to acquiring and processing information. This process helps markets to quickly incorporate all of the available information into a consensus estimate of the value of a financial instrument.

As in arbitrage strategies, speed is of the essence for a news reaction strategy. Investors have always devoted resources to getting information and getting it faster. One of the famous stories in finance lore is the story about how Nathan Rothschild used carrier pigeons to get faster information about the British victory at Waterloo.[4] Today, some practitioners of the news-reaction strategy use computers to scan news feeds for relevant information and to make and implement trading decisions.

Technical and Other Predictive Strategies

Investors have been following "technical" trading strategies since the beginning of stock trading. Technical traders use a variety of techniques that use recent price data in an attempt to discern the future direction of prices. Some of these strategies are trend- following strategies that attempt to identity a forming trend and go along. Other technical strategies attempt to spot reversals that are about to occur. While academics debate the effectiveness of various tools, they are widely used in practice.[5] As rule based approaches, many technical trading systems are readily automated with computers.

Order Discovery Strategies

Closely related to predictive strategies are strategies that attempt to discover the existence of large orders that have not been filled. For example, a trader may "ping" the market with a small order to see if there is any trading interest lurking there. If the order gets filled, the trader may surmise that a large block buy order is in process and that

[3] See Gatev et al. (2006) for a study of pairs trading.

[4] See Ferguson (1998) for more details on the Rothschilds.

[5] For a good discussion of technical analysis, see Lo et al. (2000).

the price will move up as a result. The trader then purchases stock hoping to profit from the price rise.

This is part of a cat-and-mouse game that has been going on between institutional traders and other investors for many years. Large block traders go to great lengths to reduce the price impact of their large trades by breaking them up into smaller trades. However, every time a small piece of a block is traded, the price and volume, but not the trader's identity, become public information, partially revealing some of the traders' information. Other investors attempt to spot patterns that indicate a large block is in the process of execution, which would impact the price. When they spot such a pattern, they trade accordingly. This speeds up the process by which the market incorporates information into the market price. Of course, this also would appear to increases the market impact for the institution that is executing the block. To counteract this effect, institutional traders attempt to mask their trading through a variety of different strategies. Some HFT algos may indeed guess that a large block is being executed, but it is merely a guess based on public information.

Some would call such legitimate order anticipation strategies "predatory" and compare them with illegal front running.[6] However, they do benefit the market in several ways. First, such strategies pay close attention to the market and attempt to trade away perceived mispricings. They thus help the market to produce more accurate prices. Second, by moving the price closer to the new equilibrium price, they can speed up the market's adjustment process.

One false meme that has circulated is the notion that HFT gives some investors an advance look at other investor's orders without the permission of the order submitter. For example, one internet posting (Whitney 2010) stated "It all boils down to this: HFT allows one group of investors to see the data on other people's orders ahead of time and use their supercomputers to buy in front of them."

This is not true. No exchange or trading platform would stay in business very long if it breached client confidentiality and leaked out trading information without the permission of the order submitter. Pipeline Trading was a "dark pool" trading operation that promised to protect the confidentiality of the orders in its system. The firm quickly died after it was sanctioned by the SEC for leaking information about customer orders.[7] Some exchanges do permit controlled display of the information, but with the consent of the customer. For example, the old NYSE permitted floor brokers to share some information about customer orders with other floor brokers in the hope of getting a better quality execution for the customer.

[6] See Arnuk and Saluzzi (2009).

[7] See D'Antona (2012) for more information about the Pipeline Trading scandal.

Manipulative Strategies

There are a number of manipulative trading strategies that attempt to move prices away from their real value to profit from the manufactured discrepancy. Most of them are low frequency strategies but some of them can also be implemented using high frequency technology. Such manipulative strategies are generally illegal violations of the anti-fraud provisions of the Securities Exchange Act of 1934.[8]

Front running

Front running occurs when a market participant discovers that another investor is about to make a large transaction and then "runs in front" of the trade. For example, suppose a broker receives an order to buy five million shares of stock. This large order will probably take several hours if not days to complete. With the knowledge that the large order will push the price up, the broker then buys stock for his own account before executing the client's order. This pushes up the price that the client ends up paying for the shares.

Several critics allege that the use of high frequency technology can be used to front run other orders. So-called "predatory" algorithms, or "algos," figure out that a large order is in the process of execution and jump in front of it. This is one of the predictive strategies discussed above. While it is clear that brokers who front run their own customers are violating their ethical duties to their customers, it is not clear that there is anything wrong with investors using information that is publicly available to everyone to make their trading decisions. But is it fair if some investors have access to faster computers than others? We will return to this question below.

Order Triggering Strategies

A classic manipulation is a "bear raid" in which the raider enters a short sale order large enough to push the price down. Other investors may view the drop in price as an indication that somebody knows something and follow suit, pushing the price down even more. If the price falls enough, it may trigger further sales from stop orders and liquidated margin accounts that will depress prices further. The manipulator then buys stock back at a lower price to "cover" the short and thus make a profit. Modern perpetrators of this type of illegal manipulation may use computers to search out situations when such an abusive activity may be most profitable,

[8] Section 9 of the Securities Exchange Act of 1934 bans various practices such as wash sales. Section 10b more generally bans "any manipulative or deceptive device or contrivance" as defined by the SEC. The entire text of the law can be found at http://www.sec.gov/about/laws/sea34.pdf.

such as when there are very few visible orders on the buy side of the market.

Spoofing

A trader wishing to buy may place a sell order to trick other investors into trading. Sometimes computerized traders can be tricked into changing their quotes. Here is an example. Suppose that a stock is bid $10.00 and offered at $10.02 with a large quantity available for sale at $10.02, and an investor wants to buy 5,000 shares. The investor suspects that the large size offered at $10.02 is being offered by market making algorithms that will also offer substantial size at $10.01 if that were the best offer, perhaps along with institutional sell algorithms that are trying to unload a block with a passive trading strategy of matching the offer price. By placing sell orders at $10.01, the investor may induce these quote matching algorithms to come down in price to $10.01. The new offer is now $10.01 with substantial size shown at the offer. The investor then cancels his own sell order at $10.01 and immediately places a large buy order at $10.01, thus saving $.01 on the purchase price of the transaction.[9]

Wash Sales

A wash sale is a fictitious sale that is reported for the purpose of making it appear that there is more trading activity in the stock than there really is. For example, a manipulator may put in simultaneous buy and sell orders through different accounts to make it look like there is trading interest in the stock. The manipulator does this at higher and higher prices to make it appear that the market is more liquid than it really is and that there is upward price pressure on the stock. The intent is to attract still more trading interest in the stock from gullible investors.

Quote Stuffing

It is possible to use high speed computers to send thousands of orders into a stock exchange and then cancel them immediately. The alleged intent is not to trade, but to slow down other traders whose computers are slowed down by all the message traffic. Such intentional "quote stuffing" is clearly an abusive practice that should be punished. However, it is not clear that all instances of high cancellation rates are intentional. Some might be the unintentional consequence of poorly designed software or the complex interactions of different computer systems. Regardless of the cause, excess cancellations are a form of pollution that imposes costs on everyone who has to deal with the massive quantities of data generated by these instances. Whether the product of scienter or carelessness, authorities should take action to reduce this quote pollution and to punish intentional manipulators.

Other Issues

The Need for Speed

Since many of these trading strategies are fairly simple, there is a large amount of competition to implement them. Indeed, this competition helps the markets by providing more market making capacity and by making sure that arbitrageurs keep prices in their proper relative alignment. However, the simplicity of many of these strategies means that good trading opportunities often disappear quickly. The first trader to take advantage of an opportunity wins. The second trader loses, even if her or his order was only one millionth of a second slower in arriving at the exchange. This means that there is an arms race for speed between the different competitors. Traders employing these strategies need to invest in fast computers and fast data connections to the exchanges.

Co-location

One of the more controversial features of HFT is that some traders actually place their computers in stock exchange data centers so that they can trade faster. If their computer is closer to the exchange computer, then their order has a better chance of getting in first and winning the race with the competitors. The speed of light thus matters. It takes about five millionths of a second for light to travel one mile. Thus, if the computer that is doing the trading is in an office one mile away from the stock exchange data center, its orders will arrive five millionths of a second slower than the same computer that is co-located in the exchange data center. An investor sending in an order from the opposite coast 3,000 miles away would experience a delay of approximately one one-hundredth of a second.

Does "co-location" give traders an unfair head start? To a certain extent, traders have always invested heavily to get closer to the scene of trading. Physical proximity has always mattered. A generation ago, traders paid extra to get a seat on an exchange that would give them direct access to the exchange. Now they just pay extra to put their computer in the exchange data center. Similarly, brokerage firms rented office space right next to the exchange in the olden days (or even in the exchange building itself) so that

[9] Recently, the Financial Industry Regulatory Authority, FINRA, fined Trillium Brokerage Services LLC for engaging in such activity. See http://www.finra.org/Newsroom/NewsReleases/2010/P121951 and http://www.finra.org/web/groups/industry/@ip/@enf/@ad/documents/industry/p122044.pdf for details.

their runners could carry orders from the brokerage office to the trading floor faster.

Risk

Another criticism of HFT is that the combination of many different high speed traders may impose additional risk on the market and cause excessive volatility. What if the various computer programs somehow "misfire"? Indeed, this could cause serious harm to other investors. The so called "flash crash" of May 6, 2010 demonstrates the ability of our computerized markets to misfire. Although the actual causes of the May 6 event are still being debated, the event itself does demonstrate how our computerized market structure can malfunction.[10] Are the high frequency traders making money at the expense of imposing unacceptable risk on other market participants? If so, the risks can be reduced by putting electronic safeguards into place to stabilize the market when it misfires. Soon after the May 6, 2010 event, the U.S. began to impose trading halts on stocks whose prices moved more than 10 % in 5 min.

Notions of Fairness

Our brains appear to be hard wired to prefer fair outcomes. Tabibnia et al. (2008) report brain imaging studies which show that achieving a fair outcome in an ultimatum game activates the same pleasure centers in the brain as other pleasurable activities. Unfair outcomes activate other parts of the brain. Thus, we are programed to prefer fairness. But what exactly is it?

As Boatright (2010) observes, "Fairness is a notoriously complex moral concept that has a wide range of application and standards." The word "fair" can mean a variety of different things in different contexts. One place to start is in the common dictionary definition of "fair." Dictionaries carry literally dozens of definitions of fair, and we shall not go over all of them here. As a noun, the *Oxford English Dictionary* (1993) gives its first definition of fair as "a periodical gathering of buyers and sellers often with shows and entertainments, in a place and at a time ordained by charter or statute or by ancient custom." Note the essentially commercial nature of this definition, as a gathering of buyers and sellers. It is tempting to contemplate whether this use of the word is related to the bringing together numerous buyers and sellers so that the competition arrives at "fair" prices.

Other definitions of fair include a measure of attractiveness: Again from the OED: "Beautiful to the eye; of pleasing form or appearance; good-looking." Thus fairness is also an attractive quality, not only of people, but of markets. However, markets should not only be more than just good looking but also fair in other respects as well.

The word "fair" can also connote mediocrity rather than attractiveness, or as the OED puts it "of tolerable though not highly excellent quality." When Congress called for "fair and orderly" markets, were they setting the rather low standard of tolerable?

Indeed, it is not until the 10th definition of fairness as a noun that the OED gets to the moral context:

> 10. a. Of conduct, actions, arguments, methods: Free from bias, fraud, or injustice; equitable, legitimate. Hence of persons: Equitable; not taking undue advantage; disposed to concede every reasonable claim. Of objects: That may be legitimately aimed at; often in fair game, fig.; fair wage

This definition gets to the common usage such as fair game, fair play, fair, and square.

Discussions of fairness and justice go back to antiquity. In book 5 of Aristotle's (1908) Nicomachean Ethics, he speaks of justice as fairness and points out that there are differences of opinion on how goods should be distributed among "unequals." Rawls (1958) also builds upon the concept of justice as fairness. A fair outcome is one that would be chosen by people in the "original position" who are behind a "veil of ignorance" in that they do not know what role in society they play. They would thus choose the process or outcome that provides the most benefit to the least advantaged. Leventhal (1977) advocates examining both the fairness of procedure as well as of distribution.

Economists have also tried to define fairness in more narrow settings.[11] Kahneman et al. (1986a, b) demonstrated that subjects are willing to give up payoffs to punish those who treat them unfairly. Fehr and Schmidt (1999) model fairness as "inequity aversion" in which people are willing to give up some payoff to prevent an inequitable outcome.

Pava et al. (1999) differentiates between simple fairness, in which "one person should not achieve a gain by simply imposing an equivalent loss on another" and complex fairness in which "One person should not achieve a relatively large gain by imposing a relatively small loss on another."

This notion of not causing harm to others has found its way into the U.S. law. Even though the U.S. Congress uses the word fair numerous times in the Dodd-Frank bill, it did

[10] The event was allegedly set off by a large low frequency mutual fund that put in a very large sell order. The ensuing chaos caused data integrity problems that led many HFT firms to turn off their computers because they did not have confidence in the data they were receiving from the exchanges. This caused a lack of arbitrage, leading to crazy prices for many stocks and especially for ETFs. See SEC (2010b) for more details.

[11] See Rabin (1993) and Konow (2003) for more complete surveys.

not define the term. However, Congress did come up with a narrow definition of unfair, but only for consumer financial products:

> (A) the act or practice causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers; and
> (B) such substantial injury is not outweighed by countervailing benefits to consumers or to competition.[12]

Note that like Pava's definition, this definition focuses on the tradeoff between the harm to others and the benefits to society.

Heath (2010) addresses fairness in financial markets in the context of public finance, currency, insider trading, and other areas. He focuses on the comparative treatment of different groups and views fairness as "to be treated similarly to others with respect to a rule, agreement, or recognized expectation."

Shefrin and Statman (1993) view fairness as a "claim to entitlements" in various dimensions. They identify seven dimensions of fairness in financial markets:

- *Freedom from coercion* Participants are not free to participate or not participate in a transaction.
- *Freedom from misrepresentation* Fraud is not involved.
- *Equal information* All participants have access to the same information, so there is no insider trading. This notion of fairness is enshrined in our securities laws and regulatory apparatus, which seek to create a level playing field by forcing disclosure of relevant information.
- *Equal processing power* There is no disparity in the ability of participants to process information. This follows closely from the equal information criterion. What good is information to someone who can't understand it?
- *Freedom from impulse* Participants are protected from their own irrational impulses. For example, prohibitions on drugs or cooling off periods that allow someone to cancel a transaction ex post.
- *Efficient prices* Prices reflect all the information available in the market.
- *Equal bargaining power* There is no gross disparity in the power relationships between the participants.

Is HFT Fair?

We now seek to analyze HFT in the context of these various definitions of fairness. This task is complicated, however, by the diversity of different HFT trading strategies. Clearly, some manipulative uses of HFT are unfair under any definition. For example, order triggering strategies that manipulate prices away from their real value to trigger others to trade steal from those other traders and provide no benefit to other investors. Intentional quote stuffing is the moral equivalent of intentional pollution. But what about other practices?

First, let us examine a common denominator of many HFT strategies, the co-location of their computers in stock exchange data centers. From the perspective of procedural fairness, there is nothing wrong with co-location so long as it is available to everyone on the same terms. If a stock exchange were to prohibit some participants from access to co-location, then it would be unfair.

In the context of distributional fairness, co-location raises the question of the unequal endowments of society. Many, if not most, investors lack the resources to purchase expensive computers and rent space inside stock exchange data centers. Furthermore, this is not the only inequality in our financial markets in that many investors have resources to acquire one type of edge or another in the markets. Is it fair that some investment firms can hire the best and brightest analysts that money can buy? Is it fair that Warren Buffett has more skill than most investors?

Are non-co-located investors at a disadvantage because others can execute their orders a few millionths of a second faster than they can? Clearly, the co-locators think they have an advantage worth paying for. The only investors at a disadvantage would be those who are competing in the same types of speed sensitive trades as the HFT traders, such as arbitrage and market making. Long-term investors are not attempting to profit from split-second imbalances. Indeed, they benefit from the liquidity provided by market makers and the quality of prices enforced by arbitrageurs. Market making and arbitrage have always required rather large investments in capital, trading technology, and skilled employees, so it is hard to argue that adding computers suddenly changes the fairness or the morality of these activities.

A Rawlsian veil of ignorance provides a useful insight. Traders in the "original condition" would not know what kind of investor they would be: poor non-participant, retail investor, institutional investor, HFT user, regulator, etc. What is the best outcome for the worst of these? In other words, would these investors choose a world that permitted co-location even though they did not know what kind of investor they were? We believe they would, as the arbitrageurs and market makers improve the quality of the market for everyone. Although the market makers and arbitrageurs compete vigorously with each other, no one has to trade with them. Even the non-investor gains from the benefits that properly functioning capital provide to society, as well functioning capital markets foster an efficient allocation of capital to productive uses.

[12] Subtitle C, Sect. 1031 of the Dodd–Frank law. For another attempt at implementing fairness in law, see Ledvinka (1979) for a statistical approach.

The multiple dimensions of Shefrin and Statman (1993) also provide a useful framework for analyzing the fairness of co-location and HFT in general. Clearly, HFT would be considered fair under the dimension of freedom from misrepresentation, as there is not fraud involved. Freedom from impulse is not an issue, as HFTs are not taking advantage of drunks or others unable to control themselves. Likewise, there does not appear to be any inequality in the bargaining power of the co-located firms relative to other traders in getting access to stock exchange co-location facilities, and there is no coercion involved with co-location.

The other dimensions present more of a gray area and demonstrate why there has been such a public outcry. The co-location of HFT computers in exchange data centers strikes some as a violation of the equal information dimension, as their close location to the exchange gives them access to data faster than others located further away. Indeed, earlier versions of the Nasdaq trading system in the 1990s were designed to provide quotation information at exactly the same time to investors all over the country by building in delays based on location. Now the reverse is true, and firms that pay extra can buy a speed and thus information advantage.

If one assumes that anyone who is not co-located could locate one mile away, this gives the co-located party five millionth's of a second speed advantage. Critics charge that they use this speed advantage to take advantage of other slower investors. Likewise, HFT firms have access to large amounts of computer processing power that are unaffordable to many investors. Does this give them an unfair advantage? Defenders of HFT point out that this speed of computing and location in exchange data centers are available to anyone who is willing to pay for it, a procedural fairness defense. The five millionths of a second advantage that co-location provides, which really only matters in the competition with other high speed traders with whom they compete. To most investors, the five millionths of a second delay in getting their order in is not going to make any difference—they are not racing with other arbitrageurs or market makers, and those other investors benefit on average from the services of the arbitrageurs and market makers. Such non-HFT investors willingly accept much longer delays in getting their orders processed, including delays in their own computers and communication lines, their own human reaction time, their choosing to trade from extremely remote locations, or delays in their brokerage firm's computers.

The dimension of efficient prices that reflect all information available in the market is a similar area of controversy. The same points and counterpoints apply: Some claim that the five millionths of a second time advantage gives the co-located firms an informational advantage such that they see the real price while the rest of the world sees stale prices. Such firms can jump in and trade first when the price changes. Again, this only matters to those who are competing for the same type of arbitrage and market making trades,

A more important issue is that of the impact on the quality of prices. Is the market more volatile as a result? Does market quality suffer? This is an empirical question, upon which there is much empirical work currently under way. Brogaard (2010) finds that HFT activities are not detrimental to the market and even provide benefits. Clearly different trading strategies will have different impacts on market quality, as the beneficial players improve market quality and the abusive ones degrade it. The U.K.'s Government's Office for Science (2012) generally found computerized trading beneficial but warned of harmful practices as well.

We next examine HFT through the lens of simple fairness found in Pava et al. (1999). Do HFT investors gain by imposing equivalent losses on others? It would be unfair if the activities of high frequency traders impose substantial losses on other investors or otherwise disrupt the market in a manner disproportionate to the benefits they provide. Some would argue that trading is a zero-sum game and therefore the gains earned by HFT traders are earned at the expense of others.

The notion that HFT traders make money in down as well as up markets strikes some as an unfair allocation of both risk and return. We disagree with the notion that trading is zero-sum, as both parties freely enter into a trade expecting to benefit. As financial products are risky, each side may feel that it is better off because it has managed its risk better as well as the price. For example, an investor who is eager to sell may not want to take the risk that the stock may fall in the near future, and thus sells to a market maker at the market maker's bid price. The investor could have placed a limit order to sell at a higher price, but would have then had to wait for a buyer to come into the market willing to pay that higher price. This runs the risk that no such buyer would arrive. Thus, the cost of accepting the lower bid price from the market maker is really a form of insurance against future losses for the seller. Thus, rather than imposing losses on the buyer by paying a lower price, the market maker is actually selling a benefit that the buyer is paying for. The market maker is selling the insurance-like product of immediacy. This is a service that is in demand in both rising and falling markets, and there is nothing wrong with earning a profit in such a situation.

Once again, however, there needs to be a distinction between the beneficial uses of high speed computer technology, which can be used in ways that are beneficial to markets, and the abusive means as well. Users of HFT for order ignition, intentional quote stuffing, or other forms of manipulation are clearly imposing harms on others.

Conclusions

Although the words "fair" and "fairness" describe many separate concepts, there are two common themes in most discussions of fairness. The first is that of procedural fairness, or equal application of the rules. To the extent that any trader is permitted to buy a computer and co-locate it in an exchange data center, there is nothing particular unfair about high-frequency trading from a procedural perspective. The exchanges have plenty of space in their data centers and offer co-location service to all comers at published prices.

The other dimension of fairness is distributive and is concerned with equality of outcome. That some traders earn large profits while other people stand in unemployment lines brings up many of the issues about the inherent fairness or unfairness of a market economy. Our society tolerates such inequality in a market economy because market incentives are a tremendous incentive for efficient production of good and services. However, financial markets have always suffered from this inequality of endowments problem. Some investors start with more resources than others, and some investors invest more to generate a competitive advantage. As for HFT, this does no more to perpetuate the inherent unfairness—in the sense of equality of outcomes—of life than many other features of our capital markets.

Related to the distributive notion of fairness is the question of benefit or harm to others. Many HFT strategies do not impose harm on others, and thus they are not unfair in the sense of harming others. Indeed, many of these HFT strategies provide benefits for other market participants such as reduced trading costs and prices that accurately reflect related instruments.

Some electronic traders may use high-speed technology to engage in traditional manipulative strategies that seek to profit by moving prices away from their fundamental values. These are clearly unfair because they seek to gain by imposing harm on others. It is thus the use of the technology, rather than the technology itself, that determines fairness or unfairness.

References

Aristotle. (1908). *Nichomachean ethics* (W. D. Ross Trans.). Oxford: Clarendon Press. Accessed 22 Dec 2010, from http://www.ilt.columbia.edu/publications/artistotle.html.

Arnuk, S., & Saluzzi, J. (1999). Latency arbitrage: The real power behind predatory high frequency trading. Themis Trading LLC white paper. http://www.themistrading.com/article_files/0000/0519/THEMIS_TRADING_White_Paper_–_Latency_Arbitrage_–_December_4__2009.pdf. Accessed 2 Dec 2012.

Boatright, J. R. (2010). In John R. Boatright (Ed.). *Ethics in finance, in finance ethics: critical issues in theory and practice (Robert W. Kolb Series)* (Chapt. 1, pp. 3–20). Chichester: Wiley.

Brogaard, J. (2010). High frequency trading and its impact on market quality. In *5th Annual Conference on Empirical Legal Studies Paper*. http://www.futuresindustry.org/ptg/downloads/HFT_Trading.pdf. Accessed 2 Dec 2012.

CBS News. (2010). How speed traders are changing Wall Street: Steve Kroft gets a rare look inside the secretive world of 'high-frequency trading', *Sixty minutes*. Accessed 22 Dec 2010, from http://www.cbsnews.com/stories/2010/10/07/60minutes/main6936075.shtml.

D'Antona, J. (2012, May 10). Pipeline trading ceases operations; Top executive steps down, *Traders Magazine Online News*. Accessed 31 Oct 2012, from http://www.tradersmagazine.com/news/pipeline-aritas-biancamano-trading-110008-1.html.

Duhigg, C. (1999, July 24). Stock traders find speed pays, in milliseconds, *New York Times*, page A1. Accessed 22 Dec 2010, from http://www.nytimes.com/2009/07/24/business/24trading.html?_r=1&ref=todayspaper.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*(3), 817–868.

Ferguson, N. (1998). *The world's banker: A history of the house of Rothschild*. London: Weidenfeld & Nicolson.

Gatev, E., Goetzmann, W. N., & Rouwenhorst, K. G. (2006). Pairs trading: Performance of a relative value arbitrage rule. *Review of Financial Studies, 19*, 797–827.

Government Office for Science. (2012). Foresight: The future of computer trading in financial markets (2012). Final project report, The Government Office for Science, London. Accessed 31 Oct 2012, from http://www.bis.gov.uk/assets/foresight/docs/computer-trading/12-1086-future-of-computer-trading-in-financial-markets-report.pdf.

Heath, E. (2010). Fairness in financial markets. In John R. Boatright (Eds.), *Finance ethics: Critical issues in theory and practice (Robert W. Kolb Series)* (Chapt. 9, pp. 163–178). Hoboken: Wiley.

Kahneman, D., Knetsch, J. L., & Thaler, R. (1986a). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review, 76*, 728–741.

Kahneman, D., Knetsch, J. L., & Thaler, R. (1986b). Fairness and the assumptions of economics. *Journal of Business, 59*, S285–S300.

Konow, J. (2003). Which is the fairest one of all? A positive analysis of justice theories. *Journal of Economic Literature, 41*(4), 1188–1239.

Ledvinka, J. (1979). The statistical definition of fairness in the federal selection guidelines and its implications for minority employment. *Personnel Psychology, 32*, 551–562. doi:10.1111/j.1744-6570.1979.tb02153.x.

Leventhal, G. S. (1977). What should be done with equity theory? new approaches to the study of fairness in social relationships. In K. J. Gergen (Eds.), *Social exchange theory*. New York: Wiley.

Lo, A., Mamaysky, H., & Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *Journal of Finance, 55*(4), 1705–1770.

Pava, M. L., Pava, J., & Hochman, J. A. (1999). Fairness as a constraint in the real estate market. *Journal of Business Ethics, 19*, 91–97.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review, 83*, 1281–1302.

Rawls, J. (1958). Justice as fairness. *The Philosophical Review, 67*(2), 164–194.

Shefrin, H., & Statman, M. (1993). Ethics, fairness and efficiency in financial markets. *Financial Analysts Journal, 49*(6), 21–29.

Simpson, J., & Wiener, E. (1993). *Oxford english dictionary*. New York: Oxford University Press.

Tabibnia, G., Satpute, A. B., & Lieberman, M. D. (2008). The sunny side of fairness. *Psychological Science, 19*(4), 339–347.

U.S. Securities and Exchange Commission. (2010a). *Concept Release on Equity Market Structure*, Release No. 34-61358; File No. S7-02-10. Accessed 22 Dec 2010, from http://www.sec.gov/rules/concept/2010/34-61358.pdf.

U.S. Securities and Exchange Commission. (2010b). *Findings regarding the market events of May 6, 2010: report of the staffs of the CFTC and SEC to the joint advisory committee on emerging regulatory issues*. Accessed 22 Dec 2010, from http://sec.gov/news/studies/2010/marketevents-report.pdf.

Whitney, M. (2010). High-frequency trading as high-tech robbery, *Counterpunch*. Accessed 22 Dec 2010, from http://www.counterpunch.org/whitney04162010.html.