

The Mysterious Ethics of High-Frequency Trading

Ricky Cooper

Michael Davis

Ben Van Vliet

Illinois Institute of Technology

ABSTRACT: The ethics of high frequency trading are obscure, due in part to the complexity of the practice. This article contributes to the existing literature of ethics in financial markets by examining a recent trend in regulation in high frequency trading, the prohibition of deception. We argue that in the financial markets almost any regulation, other than the most basic, tends to create a moral hazard and increase information asymmetry. Since the market's job is, at least in part, price discovery, we argue that simplicity of regulation and restraint in regulation are virtues to a greater extent than in other areas of finance. This article proposes criteria for determining which high-frequency trading strategies should be regulated.

KEY WORDS: High-frequency trading, deception, prudence, efficient markets, finance ethics, ethics of financial services

INTRODUCTION

“By the summer of 2013, the world's financial markets were designed to maximize the number of collisions between ordinary investors and high-frequency traders at the expense of ordinary investors and for the benefit of high-frequency traders, exchanges, Wall Street banks, and online brokerage firms. Around those collisions an entire ecosystem had arisen.”—Michael Lewis, *Flash Boys* (2014, 179)

TODAY, A SPECTER SEEMS TO HAUNT FINANCIAL MARKETS; its name is “high-frequency trading.” Though high-frequency trading is embodied in hundreds of legal businesses employing thousands of people, there are at least three reasons why it might seem a hostile, disembodied, and unintelligible presence in financial markets. First, there are the large profits high-frequency trading consistently generates year after year, suggesting parasitic rent-taking rather than honest earnings. Second, there is the fact that high-frequency traders do not treat securities (financial instruments) the way ordinary participants in financial markets do. When high-frequency traders buy a share in a company, they typically hold it only for a short time, often a few seconds, seldom, if ever, for a day. They care little, if anything, about the underlying value of the companies, the shares of which they buy or sell, much less about the social

responsibility of those companies or even their honesty. High-frequency traders seem more like gamblers than investors. Why then do they make so much money, and make it so consistently? Which brings us to the third reason they seem so spectral. The answer high-frequency traders give to questions about their success is (something like) “by using algorithms to reduce market inefficiencies.” High-frequency trading is algorithmic trading. Hence, the large number of mathematicians, computer scientists, and engineers such firms typically employ, “geeks” who appear to know little of finance and even less about how to run a business. To outsiders, it might well seem like alchemy or some darker art.

Government regulators seem to have decided that the dark art high-frequency traders practice is “deception.” While regulators have long banned certain trading strategies—front running, late trading by mutual funds, cornering the market, squeezing, pumping-and-dumping, market conditioning, wash trading, and so on (see Kyle and Viswanathan 2008, 275-276)—under the capacious heading “market manipulation,” they have only recently defined “deception” (what Kyle and Viswanathan (2008, 277) call “bluffing”) as a distinct form of misconduct in algorithmic trading. (See, for example, US CFTC Final Rule 180.1, and US SEC Rule 240.10b-5, both of which prohibit “manipulative and deceptive devices.”) Under these new regulations, a trading algorithm may no longer be programmed to mislead a competitors’ trading algorithm.

Like poker, the financial markets are competitive domains. And, as in poker, some competitive practices, like bluffing, are “fair,” and others, like playing with a marked deck only one player can interpret, “unfair.” The rationale for this distinction in poker is pretty clear. Bluffing is hard to regulate and adds to the excitement of the game. In contrast, playing against a marked deck takes much of the fun out of poker, an important consideration for those of us, perhaps the great majority of poker players, for whom fun is the chief reason to play. If all know the markings, poker becomes a game of mere chance, a game in which all cards might as well be face up. If only one player knows the markings, the game will be one few would knowingly choose to play unless they were the player who knew the markings. A game is fair if, and only if, all participants, knowing its rules, would choose to play. To cheat (to act unfairly) is to act contrary to the rules of a fair practice (whether a game or other cooperative activity) (Gert 2005, 196-197). Bluffing in poker is not cheating because the rules of poker allow bluffing. Playing with a marked deck that other players do not know about is cheating because the rules of poker do not allow marked decks (and poker is a fair practice).

In a competitive financial market, there are at least four kinds of regulation possible: 1) self-regulation (firm by firm “corporate responsibility”); 2) regulation by venue (the New York Stock Exchange, Chicago Mercantile Exchange, or the like); 3) regulation by industry or federation of venues, such as the Financial Industry Regulatory Authority or similar self-regulatory organization (SRO); and 4) governmental regulation (state, federal, or intergovernmental, such as the SEC or the Organisation for Economic Co-operation and Development).

Whatever the kind of regulation, the questions we wish to consider are the same, and all are evaluative:

1. What problem does high-frequency trading in general, or some of its practices in particular, pose?
2. Will enforcement of a rule like US SEC Rule 240.10b-5 fix the problem, if there is one, without creating problems at least equally bad?
3. How should financial markets be redesigned to fix the problem (assuming there is one and broad rules against deception like US SEC Rule 240.10b-5 are not the way to fix it)?

We propose to contribute to the existing literature of ethics in financial markets (as defined in Boatright 2008, 5), most notably the contributions of Boatright (2015), Ryan, Buchholtz, and Kolb (2010), Nielsen (2010), and Kane (1997), by examining whether the recent trend of regulating deception makes markets better. Though our approach fits with Nielsen's (2010) notion that markets should allow the values of individuals to be expressed without distorting them, the conclusions we draw may not completely agree with his. We argue that, in a purely transactional arena such as the financial markets, almost any regulation (other than the most basic) will create a moral hazard. The Nielsen (2010, 315 and 322) approach assumes that transparency of regulation must reduce moral hazard. However, any proposed change in regulation creates an incentive for the affected parties to seek to influence the final rule. After all, in financial markets, even small regulatory changes can mean big profits since the changes are applied millions of times each day. Once a biased rule is enacted, transparency will not help. The best way to avoid moral hazard is, all else equal, to keep regulation to a minimum. Indeed, as Ryan, Buchholtz, and Kolb (2010, 686) notes, complex financial markets have a bigger problem with information asymmetry than other markets. Since the market's job, in part at least, is price discovery, we argue that simplicity of regulation and restraint are virtues in complex financial markets to a greater extent than in other areas of finance. Ryan, Buchholtz, and Kolb (2010, 686) go on to ask what should be the ethical obligation of those who trade for their own account. As we shall argue, the obligation of the high frequency trader should be to act prudently and put out messages that can be responded to, letting the invisible hand do the rest.

We shall proceed in this way. First, we shall explain how high-frequency trading works, especially the good it contributes to financial markets. Second, we shall consider how high-frequency traders can routinely make profits fairly and how they can cheat. Third, we shall consider what sorts of deception should be tolerated, what sorts expressly prohibited, and what sorts merely discouraged. We conclude that it is a mistake to treat mere deception, or even mere intentional deception, as misconduct in a financial market.

HOW HIGH-FREQUENCY TRADING WORKS

Financial markets are not inherently efficient. They must move toward efficiency through price discovery, that is, as "private information is incorporated into prices

through trading by informed traders” (Cao, Hansch, and Wang 2009, 19). When *all* information is incorporated into the price, the randomness of price changes in a market (its “conditional Shannon entropy”) is maximized (Zhang 1999, 4-5). To the extent that non-randomness exists in price movements, there is an opportunity for a trading firm to develop private information about that non-randomness and create a trading strategy that exploits it. High-frequency traders typically have private information derived from market microstructure data (i.e. short-term information). “Low-frequency traders” typically have private information derived from fundamental price data (i.e. long-term information).¹ The difference between the two types of traders (apart from how long they tend to hold positions, their typical trading strategies, and the like) is the type of data and the analytical methods by which they derive their private information and justify their trading decisions. That a market appears to be efficient, then, is in part a byproduct of the competitive interaction among high-frequency and low-frequency trading strategies (and counter-strategies) removing non-randomness.

Deregulation of financial markets has led to a multiplication of execution venues (henceforth, “exchanges”) listing the same securities. If two exchanges list the same security, the price of a security on one exchange may differ from that on the other, creating an opportunity for arbitrage. If a trader (a firm operating through a human agent or trading algorithm) can move fast enough, it can buy the security at the lower price on one exchange and sell it at the higher price on the other. Traders can also make money in a more traditional way, that is, by buying a security when offered at a low price on an exchange and holding it until a buyer appears at that exchange willing to buy it at a higher price. High-frequency traders can also make money by supplying liquidity, that is, by posting limit order bids and offers in the exchange order book. If both the bid and offer are “hit” (say, in rapid succession), then the trader keeps the difference between the buy price on “the bid” and the sell price on “the ask.”

Competition has forced most communication between traders and exchanges to be by computers over high-speed networks. Human traders are just not fast enough to compete directly with computers when time matters. That communication by computer typically takes one of three forms:

1. **Add “limit orders”** to the exchange order book. (A limit order is an instruction to buy a security at no more than a specified price, or to sell a security at no less than a specified price. A limit order gives the trader control over the price at which the trade is executed but at the cost of ruling out filling the order at a higher or lower price than specified. Seldom, if ever, do traders send an “unlimited order,” that is, an order to buy or sell a certain security whatever the price. A limit order “rests” in the book until executed or cancelled.)
2. **Cancel limit orders** already in the book.
3. **Trade or execute a marketable order** that is matched at the exchange against resting limit orders. (A marketable order is either a market order to buy at or above an actual bid or a limit order to sell at or below an actual ask.)²

While trades certainly contain price information, they are not alone in containing such information. A limit-order add or cancel also conveys information (see Hasbrouck 1995, 1192; and Cao, Hansch, and Wang 2009, 39).³ That each limit-order add or cancel also conveys price information matters because high-frequency trading often uses limit orders as part of various trading strategies—market-making, complex dispersion, real-time replication, or index arbitrage.^{4,5} Through their limit orders, high-frequency trading strategies provide liquidity to the market (Cooper and Van Vliet 2015, 4-5). Low-frequency traders are beneficiaries of that liquidity—that is, they are happy to find a certain security available at a price they are willing to pay or to find a buyer willing to pay the price they are asking. They need not wait till an appropriate low-frequency trader appears in the market. Many academic studies support the view that, all else equal, high-frequency trading makes markets more liquid, reduces transaction costs, and dampens volatility (see for example Brogaard 2010, 4-5; Jovanovic and Menkveld 2010, 2; Hendershott, Jones, and Menkveld 2011, 3; and Hendershott and Riordan 2013, 1003). As a representative from a major exchange said, “the net benefit is that we have a better market with the participation of [high-frequency traders]” (Fabozzi, Focardi, and Jonas 2011, 29 and 35). High-frequency traders are not necessarily parasites on the market; indeed, they (generally) stand in a symbiotic relation to it.

When an exchange receives a limit-order add, an order-cancel, or an execution (i.e. a marketable order), it broadcasts a message about that event in real-time⁶ to firms that subscribe to its publicly-available data feed. Algorithmic trading strategies that receive this data feed (i.e. *receivers*) respond to it as they are designed to. Of course, such algorithmic receivers cannot (strictly speaking) “know” the intention, plan, or expectations behind the sender’s add, cancel, or execution orders (or the intentions, plans, or expectations of those who use such messages). What algorithmic receivers receive is just data. Such data, though providing information about activity on an exchange, contain no “nonverbal or stylistic [behaviors]...intended to bolster the credibility of the message” (Buller et al. 1998, 3) as might be the case between, say, traders in a pit. Nevertheless, trading firms have both the motive and the means to engage in “algo-sniffing” (see Boatright 2015) and attempt to infer the “intentions” of the algorithms of other market participants from the orders they send—what is called their “market footprint.”

In any competitive activity—sports, politics, business, or the like—no strategy works forever. Successful strategies draw a response from competitors in the form of imitation or counter-strategy. Any opportunity being exploited by one firm will eventually be discovered by other firms, some of which at least will then imitate that strategy, and “the more agents [that] use a specific strategy, the worse its results are” (Huber 2004, 171). (See also Farmer and Lo 1999, 9992.) The algorithmic trading strategy eventually stops working because the market, through the competitive process, arbitrages out that opportunity and the market is left more efficient. Likewise, if a trading firm reaches a conclusion about certain market activity (and therefore the intentions behind competitors’ algorithms), it may attempt to program its own algorithms to respond in an advantageous way. For example, the firm could program its own algorithm to buy in advance of a market rise driven

by competitors' algorithms (for example, purchase 1000 shares of such-and-such a security as soon as firm A seems to have begun buying it). Or the firm could attempt to reverse engineer a competitor's strategy or design a new algorithm that profits at the other's expense. As all firms have an incentive to extend the life of any successful strategy, it is not surprising that firms may seek to obscure their intentions, lest their observable footprint reveal some part of their strategy.

Competitive intelligence and counter-intelligence, trading strategies and counter-strategies, have probably been around since markets began. They are part of the competitive interaction between self-interested traders that overall promotes efficient markets.

ESTABLISHING WHAT'S FAIR AND UNFAIR IN HIGH-FREQUENCY TRADING

The invisible hand works through the competitive interaction of traders, each supplying or demanding securities to buy or sell according to a preferred strategy.⁷ Competition will then tend to promote a socially beneficial outcome overall, provided the financial markets are reasonably voluntary, transparent, informationally efficient, and reliable.

- A financial market is *voluntary* if participation is by free choice, not the result of compulsion or trickery, that is, all costs are internalized.
- A financial market is *transparent* if trade and quote data is disseminated in real-time so that all traders can make informed decisions (Bloomfield and O'Hara 1999, 5).
- A financial market is *informationally efficient* if observed prices contain all relevant information (Roll 1984, 1128).
- A financial market is *reliable* if there is a high probability that it will adequately perform its intended purpose—to process transactions in financial securities, enabling price discovery—in all economic environments (or, at least, in all those having a significant probability of arising). (Reliability is an important proviso because of the large cost that technological failure can impose on both market and society.)

Let us describe markets that have these four characteristics as *effective*. Society (that is, all of us, even those who do not trade or invest) should desire effective financial markets insofar as such markets tend to make life better for everyone, for example, by fostering economic growth or giving us immediate access to our savings.⁸ Effective financial markets tend to minimize the costs associated with capital formation and risk transfer. Effective financial markets also attract investors insofar as such markets justify confidence in the market's overall trustworthiness, including the fairness of prices.

On the level of the individual trader, effective markets are facilitated by prudent behavior. Originally prudent behavior was defined by the Prudent Man Rule, an outgrowth of the landmark legal case, *Harvard College v. Amory*. The rule has traditionally applied to intermediaries and fiduciaries—money managers

and brokers—agents whom others trust with their money. The original statement mandated that such agents “observe how men of prudence ... manage their own affairs ... considering the probable income, as well as the probable safety of the capital to be invested” (Massachusetts 1830). The standard has undergone many changes over the years (see, for example, Fleming 1977; Del Guercio 1996; and Rosenburgh and Spieler 2009). Nonetheless, as Longstreth (1986, 7) points out, the recurring theme is that prudence demands adherence to processes that reliably produce strategies with desirable characteristics, including monitoring results in light of the strategy’s purpose, managing risk, and minimizing the possibility of large losses. Insofar as high-frequency trading is an interdisciplinary endeavor, prudence in high-frequency trading is (like corporate social responsibility) an organizational virtue (Davis, Kumiega, and Van Vliet 2013, 864-866). But, we can also think of prudence in a more general sense, one covering non-agents (those who trade for themselves). Investors in a proprietary trading firm typically want to be prudent with their own capital (that is, to act as they would, when at their rational best, wish to act, whatever anyone else does). Though not a moral virtue, prudence is a virtue, that is, a disposition any reasonable agent must, all else equal, wish to have (Gert 2005, 294-295).

Cooper, Ong, and Van Vliet (2015) have proposed a framework that can serve as a proxy for prudence in algorithmic trading. That framework borrows concepts from the literature of quality control. It requires that an algorithmic trading strategy satisfy three criteria:

1. The strategy must operate in statistical control at all times with respect to its critical characteristics⁹, and have in place real-time monitoring to ensure containment of the trading strategy in the event it operates outside of expectations (see Cooper and Van Vliet 2012, 57).
2. The strategy must have acceptable loss behavior (i.e. reasonable risk).
3. The strategy must be able consistently to generate sufficient revenue in excess of its costs with acceptable certainty over a period of time acceptable to the investors being served¹⁰ (see also Kumiega, Neururer, and Van Vliet 2014).

If these three criteria are met, the algorithmic trading strategy is unlikely to harm the firm, the trader, or other market participants inappropriately whether it succeeds or fails; it possesses favorable risk-return characteristics relative to its operating costs. Since “the greatest concerns about HFT are the risks it poses for both firms and the markets” (Boatright 2015), this framework is an explicit attempt to incorporate two principles of Nielsen (2010, 309)—avoiding or preventing harm to others, and leverage-proportionality and prudence (see also Ryan, Buchholtz, and Kolb 2010). Cooper, Ong, and Van Vliet (2015, 64) argue that this framework is a sufficient standard, one that, if applied by reasonable people or at least certified by a firm of reasonable people, must satisfy any reasonable definition of prudence. Furthermore, we argue now that a marketplace of prudent strategies is essential for an effective market. Such a marketplace would have strategies eliminating market inefficiencies without the fear of out-of-control strategies or improperly assessed risks.

But, prudence is only a necessary condition for effectiveness. It is not sufficient inasmuch as it does not guarantee fairness. After all, even playing poker with a marked deck against those who do not know it is marked can be prudent (that is, an instance of acting according to standards the player would—at his rational best—want to follow even if no one else did the same, indeed, especially if no one else did the same). A market that is unfair will soon be known to be and will so have trouble attracting sufficient participants to be effective. This leads us to the idea that attempting to make money from trading, through the use of prudent strategies—and not utilizing unfair deception—is almost the definition of promoting an effective market. With this in mind, we need to delve more carefully into the exact nature of *fairness* and *deception*.

The debate over the fairness of certain high-frequency trading strategies seems to hinge on the ability of one trading algorithm to mislead another, that is, the ability of one trading algorithm “to take advantage of the order placement strategy of another not-so-smart computer with the possible effect that prices are driven away from equilibrium” (Stoll 2006, 169). So, if we are to analyze what makes for a fair algorithmic trading strategy, what we need is not just a general analysis of market fairness, as in Angel and McCabe (2013), but also an analysis of deception in particular, something like the analysis that allowed us to see why bluffing should be fair in poker but playing with a marked deck should not be.

The *Oxford English Dictionary* defines *deception* as causing “[someone] to believe something that is not true, typically in order to gain some personal advantage.” However, using terms like “belief” with respect to computers or algorithms as the literature on deception in high-frequency trading often does, invites confusion. Certainly, neither computers nor the algorithms that direct them have beliefs about the actions of competitors in the way that, say, a human trader in a trading pit might. Computers interpret messages according to algorithms programmed into them. The ability of computers to “think” (i.e. perceive, believe, or the like) is both controversial and well beyond the scope of this article. No one seems to claim that algorithms can literally think. We therefore use the less anthropomorphic notions of *interpretation* and *conclusion* (instead of perception and belief) to understand deception between algorithms.

According to Buller, et al. (1994, 369), deception may arise from one or more of three types of communication from a sender: falsification (deception by lying); concealment (deception by omission); and equivocation (deception through misleading ambiguity). *Attempted* deception “occurs when a sender knowingly transmits a message intended to foster a false belief or conclusion in the receiver” (Buller et al. 1998, 3). *Actual* deception occurs when the attempt succeeds and the intended false belief or conclusion occurs or when the false belief occurs as a result of an action not intended to deceive. We can deceive without intending to, for example, when I smile after drawing a card, looking to my poker mates as if I drew just the card I need when in fact I smiled because the useless jack reminded me of my son.

Some argue that a limit order is supposed to represent an intention to trade and that, therefore, actual deception occurs, for example, if the sender does not, in fact, intend to trade, but rather intends to cancel the order before it can be executed.

We can model this argument as: algorithm *A* sends order *X* to the market to get algorithm *B* to infer *p* (that a certain sender will sell certain securities at a certain price), while its controller justifiably has concluded $\sim p$ (that the sale will not occur) because the order will be cancelled before the securities can be sold. We might say, then, that *X* is deceptive (i.e. entails a proposition justifiably believed to be untrue by those having control of its content). Of course, a limit order cannot itself be false (since it is an order, not a proposition). Indeed, even the corresponding proposition may not be false (though the controller justifiably believes it is). Another market participant *C* can also respond to *X*. The communication in question (the offer) is not private. *C* may happen to have better computers or faster algorithms than expected, rendering *A*'s strategy with respect to *B* irrelevant. Thus, if *C* responds to the limit order causing it to be executed before it can be cancelled, *p* can be true regardless of the sender's justifiably concluding $\sim p$ —which is to say that an algorithmic trading strategy cannot “know” $\sim p$. Insofar as the sender's belief (or, rather, its conclusion) is in principle subject to disappointment in some such way as this, the sender's deceptive intention is more like the bluffer's deceptive intention than like playing poker with a marked deck. The sender is gambling as long as other market participants *can* (successfully) respond to the limit order (thus making *p* true).

In the sense of Kane (1997, 55), then, messages to add or delete limit orders cannot be condemned as disinformation. Kane requires “deliberate misleading statements about what particular data mean.” Messages that can be responded to are not statements about what data mean (propositions). They are only data, any meaning arising only in the mind (or algorithmic logic) of the receiver typically unknown to the sender. The sender only knows whether or not it is willing to accept the risk of the order being accepted. From that, the sender infers that it can or cannot execute the strategy repeatedly to generate profits. The acceptance or rejection of a limit order add message is no more or less deceptive than the add message itself. Such messages may be intended to exploit information asymmetries, but eliminating asymmetries is precisely one role of markets. Traders who enter this competition do so voluntarily. Longer-term investors not involved in this competition benefit from the lower cost of liquidity and lower volatility that the messages produce (see for example Brogaard 2010; Fabozzi, Focardi, and Jonas 2011; Hendershott, Jones, and Menkveld 2011; Hendershott and Riordan 2013).

In algorithmic trading, all deception is *interpretive*. Interpretive deception occurs when a receiver arrives at a false conclusion with inadequate evidence or inadequate certainty about the sender's strategy. For a trading algorithm to be deceived it must interpret messages and arrive at a false conclusion about a sender's strategy. This is the non-intentionalist approach described in, for example, Johnston (1988, 65-78), Barnes (1997, 95-100), and Mele (2001, 6-24). A receiver's false interpretation can arise with or without the intention of the sender (depending on the rules and background customs of the market). For example, suppose trading algorithm *A* identifies an arbitrage opportunity and sends a limit order to the exchange. Trading algorithm *B* receives this message and responds by sending a marketable order to execute against *A*'s limit order. But, before *B*'s response to *A*'s message reaches the exchange, *A* finds that the arbitrage opportunity no longer exists and cancels its limit

order. In this case, *A*'s cancel order has *unintentionally* led to *B*'s mistake (sending a marketable order to fill a limit order that no longer exists). Algorithm *B* has been misled into acting as if it had an opportunity to profit when that opportunity had already ceased to exist.

Notice that *B*'s false conclusion is not accidental. Rather, it is a consequence of its trading strategy and the desire of the trading firm to program algorithms that profit by transacting against opportunities like that *A*'s limit order too briefly represented. It is the trading strategy that explains the conclusion-formation process, not the intention of the sender (or the person controlling the sender). *B*'s false conclusion is not a naive mistake, but a consequence of the algorithm's design, the limits of technology (such as a certain computer's speed), and the intention of those who designed the strategy (presumably, to make money in a certain environment) (see Dewese-Boyd 2012). Notice also that not all competitors need be misled by *A*'s order but only those who were set up to respond to messages of that sort. Deception occurred in this market (in part at least) because the designers of the algorithm chose to risk certain errors.

That is not to say that high-frequency traders cannot *unfairly* induce deception¹¹ (either intentionally or unintentionally), only to make a point about how deception occurs. Sometimes trading firms tolerate the risk of error because the cost of eliminating it is too high, whether it is the cost of increasing response speed or some other factor affecting overall profit.

We will specifically address two forms of deception in the next section, one unfair and one fair, in which the intention of a sender's order may seem to matter but in fact does not. Most other forms of deception resemble one of these two forms and should be treated accordingly.

TWO FORMS OF DECEPTION

One form of deception occurs when an algorithm adds a limit order to a "lit exchange" (i.e. a transparent venue) while simultaneously placing an opposite limit order in a "dark pool" (a non-transparent venue). In this form of deception, the price information in one of the two limit orders is concealed from the market. Thus, information about supply and demand in lit exchanges does not convey the efficient price (the price an efficient market would set). Because dark pools are extremely selective not only about who may transact in them but also about who may know about what is going on in them, they are necessarily inconsistent with what we have called "an effective market." Anyone knowing that they must "play" against a strategy involving a dark pool into which they cannot see would have a good reason not to play. In this respect at least, dark pools are to financial markets what marked decks are to poker. Absent special considerations, the rules of any fair market will prohibit them.

Another form of deception occurs through a rapid sequence of limit-order adds and cancels. The result of such a rapid sequence of orders may be that some trading algorithms arrive at a false conclusion through their inability to interpret the data correctly or quickly enough to respond in an advantageous way. An algorithm so misled is said to have been "spoofed" (see Angel and McCabe 2013, 590; Boatright 2015).

We might expect that algorithms so spoofed will soon be replaced by algorithms that are less susceptible to spoofing.¹² Over time, this form of deception may become increasingly unsuccessful against increasingly intelligent competitors (see Hespanha, Ateskan, and Kizilocak 2000, 2). This is the sort of change in algorithms that tends to promote society's interest in effective financial markets—but at the expense of research that might have produced a different distribution of wealth had regulation ruled out certain deceptive strategies to begin with.

Davis, Kumiega, and Van Vliet (2013, 853) have called the evolutionary process in algorithmic trading “quality arbitrage.” Trading firms that build better algorithmic trading strategies (including ones that are less prone to being misled) survive at the expense of firms that build less good algorithmic trading strategies. A continuous cycle of quality arbitrage is, all else equal, the best way to promote society's interest. Such a cycle increases the information processing ability of financial markets and, by increasing the reliability of their trading systems, decreases systemic risk, such as risk of market disruption. Yet, this evolutionary process cannot occur if regulators make the receiver's correct interpretation of market orders the responsibility of the sender—as the prohibition of spoofing would. The (net) social cost of the regulation may, then, be higher than the (net) social cost of algorithmic defense against deception. Thus, spoofing seems to resemble bluffing in poker, deceptive in a way consistent with attracting players to the game.

In the past, human traders in a trading pit or behind a computer screen who misinterpreted the actions of other traders typically changed their trading strategy without worrying about the intentions of the other traders. The concept of spoofing never arose.¹³ In algorithmic trading, similar change in response to misinterpretation is also always possible. But, now there are significant costs associated with changing a trading strategy. Redesigning an algorithmic trading strategy requires those whose current strategy misled their algorithms to invest in the research and development of new algorithms and perhaps new technology (faster computers, better cables, and so on). The success of such investments is not guaranteed. Once researched and developed, a new algorithmic trading strategy may fail to interpret market data in a way that increases the firm's profits enough to justify its cost.

Because of this riskiness, misinterpretation in algorithmic trading invites a different response. Some trading firms may find it cheaper, and the outcome more certain, to have regulations (whether venue-based, industrial, or governmental) handicap smarter, faster, or trickier competitors (much as boxing and wrestling have weight classes to avoid disadvantaging smaller competitors). But, the point of competition in financial markets is to make as much money as possible within the rules of the game, and the reason no trading strategy works forever is always because the conclusion-forming algorithm is no longer as good as needed. Insofar as regulation makes the development of better algorithms unnecessary, the market will cease to be an evolutionary ecosystem.

Of course, stopping evolution is not necessarily a bad thing. The attempt to eradicate “invasive species” or to protect “endangered ecosystems” is precisely an attempt to stop evolution—even though those attempts are relatively uncontroversial. There are two reasons not to attempt to do the same in financial markets.

First, it is not clear that, in the long run at least, either financial markets in particular or society in general would benefit from protecting the algorithms in question. The net benefit is an empirical hypothesis about which we do not have enough information to draw a firm conclusion. We do, however, have a theoretical argument making quality arbitrage seem likely, all else equal, to produce the best results in the long run. Except where we have strong empirical evidence to the contrary for a particular algorithm or kind of algorithm, letting evolution run its course in financial markets seems a better way to move toward an effective market.

Second, given the present state of knowledge, the proposal to treat deception as misconduct is, in effect, a proposal to shift costs from one kind of market participant to another, an inherently political undertaking. Prohibiting deception would mean that some market participants would be allowed to keep certain trading strategies simply because they are unwilling to incur the costs necessary to improve them (i.e. develop algorithms more likely to form correct conclusions). With some strategies granted such protection, the market will necessarily contain strategies that reduce (rather than increase) the market's efficiency. These strategies will profit by a means other than removing inefficiencies. Strategies that tend to form incorrect conclusions among honest people *ought* to disappear, and typically will disappear if the competitive process is allowed to proceed unfettered. Why should some market participants be immune to a process that, all else equal, promotes society's interest?

Certainly, the presence of trading algorithms in a financial market that are more prone than necessary to misinterpret order data poses a systemic risk. Such strategies, in the event of a market panic or other extreme event, seem more likely to respond incorrectly and, therefore, contribute to a cascading failure across the financial system. We think the following hypothesis is, in the present state of our knowledge, at least plausible: *Market risk, typically understood as standard deviation of log returns, is a function of uncertainty around prices and the safety and reliability of the market system.* Rather than being protected, trading algorithms (and technologies) prone to misinterpret order data ought, all else equal, to be driven from the market. Such algorithms add unnecessary risk.

But, it may be objected, something like this defense of deception in financial markets could be made for hacking or even burglary. They too fuel an evolutionary process, forcing us to spend ever greater sums on security. Yet, their undeniable contribution to that evolutionary process does not seem to be a convincing argument in favor of legalizing either. Indeed, the expense of protecting against hackers and burglars seems a good reason to keep such activity illegal. The less of it, the better. We would all prefer a world where it is unnecessary to lock our windows or protect our computer files with passwords. Unless spoofing, and similar deceptive practices, differ from hacking and burglary in some significant way, there must be something wrong with the evolutionary argument just offered.

Our response is that as long as the competitive activity of suppliers of liquidity (that is, high-frequency traders) promotes effective financial markets overall, society (and regulators) should, all else equal, not treat high-frequency traders like hackers or burglars. Society should allow quality arbitrage to work unimpeded. If one trading algorithm, high frequency or low, makes money by exploiting a weakness

in another trading algorithm, that should be of no concern, because the net result of such exploitation will, in time, be deeper, more efficient, and more reliable markets in which to execute low-frequency trades—unless, of course, low-frequency traders are somehow morally entitled to that part of high-frequency trading profits attributable to deception in general or deliberate attempts to mislead in particular.

THE MORAL ENTITLEMENT OF LOW-FREQUENCY TRADERS

Debate over the morality of high-frequency trading seems to presuppose that algorithmic trading strategies are “implicit moral agents” (see Moor 2006, 19). Such strategies are not human, but nevertheless what they do can be evaluated as the actions of those who designed them, set them in motion, or oversee their operation. Among the moral rules that apply to such humans is, “Do not deceive.” Of course, that moral rule is merely presumptive. It applies all else equal (Gert 2005, 76-78). There certainly are exceptions to it. That is why bluffing (a form of deception) is morally permissible in poker. So, though we must agree that, *all else equal*, low-frequency traders (like everyone else) are entitled not to be misled, or at least that high-frequency traders should not mislead them intentionally, it does not follow that they are so entitled *all things considered*. There are at least three reasons to treat many deceptive practices of high-frequency traders as we treat bluffing in poker.

1. **Consent.** Low-frequency traders enter a financial market freely. In doing so, they consent to its official rules (all those rules that are public and supposed to govern the market). Low-frequency traders may stay out of the market on any day they find the official rules so disadvantageous that they do not wish to play the market. The large presence of low-frequency traders in a market is therefore strong evidence that they have consented to its official rules. Insofar as those rules allow for deception, low-frequency traders have consented as well to risk being deceived, just as they would consent to being bluffed if they joined an ordinary poker game. Their consent brings deception in the market under an exception to “Do not deceive.” All else is not equal.

Of course, this argument applies to deception that is officially allowed. It is does not apply to all deception, just as allowing bluffing in poker does not mean that playing with a marked deck is also allowed. The presence of low-frequency traders is proof of consent to the rules of the market only insofar they know those rules and have a reasonable understanding of what they mean. We do have reason to believe that low-frequency traders give their informed consent to the official rules of today’s financial markets even if they would prefer somewhat different rules. They enter the market every day even though they know increasingly more about high-frequency trading and its overall effect on financial markets. The next reason explains why they would, even if fully informed, do so even if deceptive practices like deliberate spoofing were officially allowed.

2. **The invisible hand.** The invisible hand now operates through an evolving “ecosystem” of algorithms (as explained above). Many agree that having more and more of those algorithms providing efficiency services in the market, especially if they differ from each other, tends to promote market effectiveness (see for example

Farmer and Skouras 2013, 334). As with any other business, so in financial markets: supply and demand tend to find a proper balance—including, in the case of markets, enough diversity in high-frequency strategies. Success and failure among such strategies (and the technologies that enable them) tend to improve both the ability of the market to process information and the reliability of the global trading system.

As long as these strategies promote market effectiveness, high-frequency trading clearly supports society's interest overall. Therefore, we should not declare anything about such trading *inherently* immoral. Angel and McCabe (2013, 585) support this view. They analyze what it means for a market to be fair and conclude, "one cannot categorically denounce the practice [high-frequency trading] as unfair."

Nevertheless, algorithmic trading strategies are under scrutiny by both the media and regulators. The intent behind these strategies is, of course, generally supposed to be to make money and so, it is no exaggeration to say, the premise of the invisible hand itself seems to be up for debate. The debate begins with good intentions—there is "concern over the practices of one group of players and the resulting costs to other players (Kearns, Kulesza, and Nevmyvaka 2010, 50)." Angel and McCabe (2013, 585) claim that, "...it is not the speed of the tool that matters for fairness, but what is done with it." Essentially, the fear is that the consistent returns some high-frequency trading strategies generate represent some kind of systematic defrauding of other market participants. The free-market rebuttal is that, all else equal, such returns must represent *fair* profits earned by providing efficiency-promoting services that have high value in the marketplace (Menkveld 2013, 715). It is up to those who claim otherwise to show that all else is not equal in some significant way.

3. Parity between high-frequency traders and low-frequency traders. In any free market, the interplay between supply and demand is a competitive price negotiation. Suppliers seek to charge the highest possible price, while demanders seek to pay the lowest possible price. Each side seeks competitive advantage in the negotiation through their respective sources of private information. This competition may lead to deceptive practices (deliberate attempts to mislead as well as accidental, negligent, or reckless deception).

In the negotiation of the cost of liquidity, one form of deception available to liquidity providers is to cancel their limit orders when they infer the arrival of orders driven by fundamentals (the strategies of low-frequency traders). Now, algorithmic trading strategies have no occult ability to see the future arrival of fundamentally-driven orders. They may be able, however, to infer such arrivals from the footprint of low-frequency traders' execution strategies, and this footprint includes activity on other exchanges. If this footprint contains information from which high-frequency traders can infer future activity on the same exchange or another, then the footprint reveals a short-term inefficiency. High-frequency traders must be free to act on that information or the market will be slightly less efficient. So, should liquidity providers be required to ignore the market footprint of low-frequency traders? If so, then price changes will not incorporate all information. Should the limit-order book contain short-term price information because its availability serves the interests of low-frequency traders without benefit to society? The expectation that limit-order books make short-term information freely available to long-term traders, though

the cost of discovering it falls on high-frequency traders, seems both inconsistent with an effective market and (all else equal) unfair.

Traders must earn profits by competing effectively. But, in the competition to avoid paying for liquidity and earn short-term profits, low-frequency traders often lose out to faster traders who are better at extracting profits from short-term information. This fact certainly suggests that the markets are rigged in favor of high-frequency traders (and are therefore unfair). Further, low-frequency traders may seem good (since they are “investors”), while high-frequency traders seem bad (“mere speculators”), based on nothing more than the type of inefficiency they seek to profit from. This is the dubious appeal of the deception-as-immoral argument. Why declare one category of trading strategy immoral simply because some traders are unwilling or unable to incur the costs necessary to build and operate trading strategies that are able to compete successfully against them (for short-term profits)? The appropriate criterion for regulation is not the welfare of any group of traders as such but the welfare of society overall.

In the negotiation of the cost of liquidity, low-frequency traders also use deception—prudently attempting to conceal the true nature of their demand, usually by the practice known as *iceberging*. Iceberging occurs when an executing broker, or an algorithmic execution strategy, conceals the true size of a large order to obtain a better average price (a price better than would be achieved if the full order were sent in all at once). This form of deception (usually the dividing of a large order into small pieces offered over an extended period) creates an inefficient price because the price of the security does not incorporate all the information about actual demand. The iceberg uses deception (by concealment) to profit at the expense of high-frequency traders.

Since iceberging is deceptive, it should count as market misconduct (that is, assuming that the SEC’s new anti-deception regulation is interpreted literally).¹⁴ Yet, no one (not even high-frequency traders) considers iceberging to be misconduct, much less immoral. Taking large positions stealthily in order to conceal one’s true intentions is part of being a good trader. In fact, the SEC (2014) mandates this form of deception insofar as “brokers are legally required to seek the best execution reasonably available for their customers’ orders.” The regulations that now prohibit spoofing as deception do not, it seems, prohibit iceberging (though it too is deception). That inconsistency demonstrates that deception as such is not the issue between high-frequency and low-frequency traders.

Of course, a common criticism of high-frequency traders is that they only provide liquidity during normal markets and that, during extreme events, they tend to withdraw that liquidity. This argument rests on fact but on fact having nothing special to do with the morality of high-frequency trading. Such liquidity vacuums happened in markets consisting entirely of human trader too, even ones without day-traders or market-makers. Extreme volatility in a market always discourages participation. Unless we are willing to say that high-frequency traders should be responsible for providing loss insurance to low-frequency traders, the question should be how fast high-frequency trading helps to bring the market back toward an equilibrium price.¹⁵

To draw a rather blunt conclusion, the marketplace exists to provide an effective place for investors to buy and sell financial instruments. It does not exist for any particular type of trading firm to make a profit. Trading firms must either add value by facilitating the interactions of traders with the market, or by making the market more effective in performing its role. Ultimately, these roles may be performed by a mix of high frequency and low frequency firms or entirely by one type of firm. However, no firm has a moral right to have its continued viability guaranteed. The behavior of all types of firms should be evaluated in a uniform manner.

WHAT PRACTICES OF HIGH FREQUENCY TRADING TO ALLOW: A PROPOSAL

As we have seen, deception cannot be a proxy for manipulation or other undesirable activity in algorithmic trading. In a financial market that is free, deception is not as such immoral nor is it in society's interest to exclude it. The only appropriate standard for exclusion is whether an algorithmic trading strategy is consistent with society's interest in effective financial markets. Given that standard, we offer three simple criteria for any algorithmic trading strategy to be allowable. We think these three criteria to be individually necessary and, perhaps, jointly sufficient, to foster an effective (and therefore fair) marketplace:

1. **The trading strategy should be prudent** (in the broad sense explained above). Prudence ensures that the risks inherent in the strategy are thoroughly investigated, and that investors and external market participants will not be unduly harmed should the system operate outside of expected behavior.
2. **The trading strategy should not block price discovery.** We have stated this second criterion weakly and in the negative because we cannot say that, to be consistent with society's interest, a trading strategy must have a good intention or even that the orders it sends must add good information. If an order can be responded to, then it is the market's job to process the information in that order (and all other orders) to arrive at an equilibrium price. It is not the job of the regulators, governmental or private, to decide which information is fit for market consumption (provided criteria 1 and 3 are also met). If that were the regulator's job, then the price would not include all information otherwise available but merely information the regulator decided was of high enough quality. In general, regulators are not in a good position to make such decisions, especially in a domain where evolution is rapid. The market is a better judge. In order to maintain an effective market in algorithmic trading, the only requirement is that one market participant may not interfere with the ability of other market participants to add to their private information. So, we can say that quote-stuffing (as in Angel and McCabe 2013, 590) violates this criterion and, therefore, should be prohibited.¹⁶ Quote-stuffing strategies may add limit orders and thus add information; but, as the strategy seeks

to derive profit from preventing others from adding their information, its net effect is to decrease the amount of information in the price.

3. **The trading strategy should not circumvent transparent price discovery.**

Trading strategies that circumvent transparent price discovery, for example, through the use of dark pools or hidden orders, should not be part of a market's official rules. Such trading strategies conceal information needed for a market to be effective—and, insofar as such strategies are known to operate, tend to discourage participation by those likely to be left in the dark.

Operationalizing these three criteria means developing industry-wide standards, private or governmental, for control of dark pools and quote-stuffing. But “control” does not necessarily mean “prohibition.” One of the challenges of rule-based oversight is that “in order for the regulation ... to function properly ... the technology involved must be reasonably stable and well understood” (Unger 1994, 190). But, the economic environment in general, and trading technologies of financial markets in particular, are rarely stable—and never for long. They typically change under the pressure of competition and competition seems to be a given. Regulators may therefore sometimes do better to tolerate a certain otherwise undesirable practice, or merely discourage it, rather than prohibit it outright, if the outright prohibition would discourage quality arbitrage.

Consider, for example, the practice of rapidly posting and cancelling orders intended merely to obtain information about demand or supply (a form of spoofing). This practice is deceptive and does not directly contribute to liquidity. Should it be prohibited? Perhaps not. When a market starts moving rapidly, liquidity providers must be free to update their bids and asks rapidly based on replication or hedging costs. Unnecessarily long resting times would systematically burden liquidity providers, who would then be forced to widen their spreads to compensate for the greater risk. Spreads that are wider than necessary do not, all else equal, promote society's interest. The ability to cancel and re-post limit orders is essential to the function of high-frequency traders and therefore should be their prerogative, all else equal. If an exchange is considering proscribing strategies that merely add and cancel without ever transacting, a reasonable alternative to flat proscription is to fine (or otherwise burden) those who exceed some standard order-to-fill ratio¹⁷. While order-to-fill ratios—like a flat prohibition—also tend to reduce market efficiency, they are much less invasive than a flat prohibition would be and may be part of maintaining an effective market.

Regulations that prohibit spoofing may end up being used to unfairly disadvantage liquidity providers who happen to cancel a large percentage of their orders because of a perfectly proper trading strategy (one serving society's interest in effective markets). Such regulations typically depend on arbitrary thresholds and circumstantial evidence as proxies for intent. As liquidity providers seek to avoid those thresholds, liquidity and market efficiency are likely to suffer significantly, especially during times of market stress when providing liquidity is vital. A firm's orders involving a particular security may (in isolation) look like spoofing¹⁸. There may, nonetheless, be a legitimate reason for adding and cancelling quickly. But, to prove that legitimacy, the firm might have to reveal a trading strategy, and no trading firm wants to reveal a

valuable trade secret, making the investment in the strategy less profitable. Thus, the evidence of wrongdoing, however circumstantial, will typically lie unchallenged on the side of the prosecutor. The defendant firm will often prefer to plead guilty, pay the penalty, and continue to use the strategy (perhaps adjusted slightly to avoid similar charges), rather than try to prove its innocence (though it is innocent).

Our proposal differs from earlier proposed solutions to the perceived unfairness of deception in high-frequency trading in one important way: we argue that transparency is a better response than flat prohibition as a way to protect against the bad effects on society of deception in high-frequency trading. Even a small regulatory intrusion may create an opportunity for considerable profits since the regulation may affect millions of transactions each day. The moral hazard here is the creation of a competitive arena where regulatory manipulation is more profitable than competition. Minimal regulation of the kind we propose seems to reduce this risk while facilitating transparency and the alignment of public and private regulatory incentives (see Kane 2010, 64-68). The kind of regulation we propose recognizes automated financial markets as a purely transactional environment differing from most other areas of finance where the ethical problems primarily concern the relationship of certain agents to other (human) agents and society.

CONCLUSION

In free financial markets, the intent of algorithmic trading strategies is generally to make money. The invisible hand works through the interaction of trading algorithms competing for profits by supplying liquidity or demanding it and by processing price information. Quality arbitrage between trading algorithms generally promotes society's interest in effective financial markets. Of course, certain market participants—especially those with small research budgets, little imagination, and so on—stand to gain if regulators constrain smarter, faster competitors. This has always been true and regulatory capture is certainly a way for such firms to seek competitive advantage. In financial markets, competition without deception is both hard to arrange and not necessarily in society's interest.

The focus of regulation of financial markets, whether private or governmental, should be promoting effective financial markets. Such markets allow quality arbitrage to improve their efficiency and reliability. Regulation should seek to assure voluntariness, transparency, and so on. Any additional regulation will merely invite regulatory arbitrage, regulatory capture, and less innovation, all without benefit to society. Attempts to regulate the intentions of high-frequency traders will generally result in markets that are less efficient and less reliable in ways that, in the long run at least, harm both society and low-frequency traders.

ACKNOWLEDGEMENTS

We are grateful to *Business Ethics Quarterly's* editor in chief, Denis Arnold, two of his reviewers, and John Boatright for their encouragement and insightful comments and suggestions.

NOTES

1. We use “long term” and “short term” for convenience, recognizing that the information in question lies on continuum that affects the price of a security across all time frames.

2. For example, if the best offer price for a stock is 100, and a buyer enters a market order, then the market order will be matched immediately against the best offer price. Likewise, if a buyer enters a limit order to buy with a limit price greater than 100, that order will also be matched immediately and not enter the limit-order book.

3. Several studies show that investors’ order choices—between limit and market orders—are essential to understanding how new information is incorporated into prices (e.g. Anand, Chakravarty, and Martell 2005; Bloomfield, O’Hara, and Saar 2005).

4. For simplicity, we assume that high-frequency traders do not place marketable orders. To the extent they do, these orders either represent a small portion of the orders they send or are not part of their high-frequency strategy.

5. High-frequency traders place limit orders primarily because the inefficiencies they seek to exploit generate small average returns per trade. They hope to earn the spread as compensation for the service they provide (liquidity) and the risk of adverse selection. It costs money to operate an infrastructure that provides liquidity, and the risk is that low-frequency (or fundamentally informed) traders put positions to them.

6. Or, more properly, near real-time.

7. This is to say that traders pursue their interest knowing that other traders are doing the same. They form expectations about what other traders are going to do and how other traders will respond to their actions.

8. Among other reasons society may foster free markets is that the results yield tend to be fair or that other social arrangements tend to cause more social discord.

9. A critical characteristic of an automated trading system is any measurable pre- or post-trade output variable that is indicative of acceptable or unacceptable performance of the system or portend its failure. Critical technological characteristics—response time, packet loss, jitter—are typically described as quality of service (QoS) (Bernat, Burns, and Llamosi 2001, 308). Critical strategic characteristics are also important—for example, profitability, ratio of winners to losers, number of orders sent to the exchange (see Bilson, Kumiega, and Van Vliet 2010, 42). Once identified, expectations about the future performance of critical characteristics can be established. A specification defines the desired value of a critical characteristic, as well as its allowable variation from that expected value. Thus, each characteristic has a reference distribution of tolerable performance.

10. Cooper, Ong, and Van Vliet (2015, 59) quantifies this measure of capability as in equation (1).

$$GC_p(n) = \frac{\mu_n - c}{\mu_n - Q_n(\alpha)} > 1 \quad (1)$$

Once control is established, if the distribution of average payoffs μ has a location and shape sufficient to cover its costs c (i.e. $\mu_n - c > 0$) over a sample period n with some specified level of confidence (i.e. some probability $1 - \alpha$), then the trading strategy is capable. As a numerical example, if the mean expected return of a trading strategy is 10% on trading capital per n -day period, the cost is 5% per period, and the left tail value of $Q_n(\alpha)$ is 6%, then $(10\% - 5\%) / (10\% - 6\%) = 1.25$ and the trading strategy is, therefore, capable according to equation (1).

11. Whaley (1982, 183) argues that both falsification and concealment are always present in any intent to deceive. In financial markets, however, no message is false (though many are intentionally misleading or otherwise seek to conceal—and may in practice lead to acceptance of false propositions).

12. Typically trading algorithms that lose money do so because they misinterpret market data in some way.

13. Before electronic trading, a trader who tried to spoof (or just wanted to take back an offer made by mistake) would either have to stand by the offer or be driven from the exchange. Spoofing was impractical.

14. To say that an act is “deceptive” is, please note, not to say it is morally wrong, all things considered, much less that is inappropriate in every circumstance. Bluffing, though deceptive, is altogether appropriate in poker.

15. In this respect, participants in the flash crash should not be condemned for the speed with which the market fell, but rather praised for the speed with which the market rebounded to the correct price. This brief disorder in the market is certainly a much better outcome than the crash of 1987, which bankrupted many firms and had other bad long-term effects.

16. Angel and McCabe (2013, 590) defines “quote stuffing” as an attempt to overwhelm the network bandwidth of an exchange by sending and cancelling unusually large numbers of trade requests (say, thousands per second), thereby blocking other market participants’ access to the market.

17. Exchanges have the ability to limit quote-stuffing and spoofing by fining firms based on a message-to-fill ratio, say 50 to 1. Firms that fail to execute at least one trade for every fifty messages (i.e. orders) sent to the exchange server are in violation. Some exchanges assess these fines, some don’t. We do not know how effective this form of spoof-prevention is.

18. For example, an inter-market arbitrage strategy may frequently update bids and offers (i.e. send add and cancel orders) quickly based upon well-researched statistical relationships between multiple securities. As prices in some securities on some exchanges change, the system responds by sending add and cancel orders in other securities on other markets for reasons not discernable to competitors or regulators.

REFERENCES

- Anand, A., Chakravarty, S., and Martell, T. 2005. “Empirical Evidence on the Evolution of Liquidity: Choice of Market Versus Limit Orders by Informed and Uninformed Traders.” *Journal of Financial Markets* 8: 288–308.
- Angel, J. J., and McCabe, D. 2013. “Fairness in Financial Markets: The Case of High Frequency Trading.” *Journal of Business Ethics* 112: 585–595.
- Barnes, A. 1997. *Seeing Through Self-Deception*. New York: Cambridge University Press.
- Bernat, G., Burns, A., and Llamosi, A. 2001. “Weakly Hard Real-Time Systems.” *IEEE Transactions on Computers* 50: 308–321.
- Bilson, J., Kumiega, A., and Van Vliet, B. 2010. “Trading Model Uncertainty and Statistical Process Control.” *Journal of Trading* 5: 39–50.
- Bloomfield, R., and O’Hara, M. 1999. “Market Transparency: Who Wins and Who Loses?” *Review of Financial Studies* 12: 5–35.
- Bloomfield, R., O’Hara, M., and Saar, G. 2005. “The ‘Make or Take’ Decision in an Electronic Market: Evidence on the Evolution of Liquidity.” *Journal of Financial Economics* 75: 165–199.
- Boatright, J. R. 2008. *Ethics in Finance*, 2nd ed. Oxford: Blackwell Publishing.
- . 2015. *Ethics in Finance*, 3rd ed. Forthcoming.
- Brogaard, J. A. 2010. “High Frequency Trading and Its Impact on Market Quality.” Northwestern University working paper. Accessed on December 10, 2015. https://secure.fia.org/ptg-downloads/hft_trading.pdf.
- Buller, D. B., Burgoon, J. K., White, C. H., and Ebesu, A. S. 1994. “Interpersonal Deception VII Behavioral Profiles of Falsification, Equivocation, and Concealment.” *Journal of Language and Social Psychology* 13: 366–395.
- Buller, D. B., Burgoon, J. K., Buslig, A., and Roiger, J. 1998. “Interpersonal Deception Theory: Examining Deception from a Communication Perspective.” Arizona University. U.S. Army Research Institute for the Behavioral and Social Sciences, DTIC Research Note. Accessed on December 10, 2015. http://www.au.af.mil/au/awc/awcgate/army/ari98-16_interpersonal-deception.pdf
- Cao, C., Hansch, O., and Wang, X. 2009. “The Information Content of an Open Limit-Order Book.” *Journal of Futures Markets* 29: 16–41.
- Cooper, R., Ong, M., and Van Vliet, B. 2015. “Multi-scale Capability: A Better Approach to Performance Measurement in Algorithmic Finance.” *Algorithmic Finance* 4: 53–68.

- Cooper, R., and Van Vliet, B. 2012. "Whole Distribution Statistical Process Control for High Frequency Trading." *Journal of Trading* 7: 57–68.
- Davis, M., Kumiega, A., and Van Vliet, B. 2013. "Ethics, Finance, and Automation: A Preliminary Survey of Problems in High Frequency Trading." *Science and Engineering Ethics* 19: 851–874.
- Del Guercio, D. 1996. "The Distorting Effect of the Prudent-Man Laws on Institutional Equity Investments." *Journal of Financial Economics* 40: 31–62.
- Deweese-Boyd, I. 2012. "Self-Deception." *The Stanford Encyclopedia of Philosophy*, edited by Zalta, E. N. Accessed on December 10, 2015. plato.stanford.edu/archives/spr2012/entries/self-deception/
- Fabozzi, F., Focardi, S. M., and Jonas, C. 2011. "High-Frequency Trading: Methodologies and Market Impact." *Review of Futures Markets* 9: 7–38.
- Farmer, J. D., and Lo, A.W. 1999. "Frontiers of Finance: Evolution and Efficient Markets." *Proceedings of the National Academy of Sciences* 96: 9991–9992.
- Farmer, J. D., and Skouras, S. 2013. "An Ecological Perspective on the Future of Computer Trading." *Quantitative Finance* 13: 325–346.
- Fleming, A. 1977. "Prudent Investments: The Varying Standards of Prudence." *Real Property, Probate and Trust Journal*, 12: 243–255.
- Gert, B. 2005. *Morality: Its Nature and Justification*, Revised Ed. Oxford: Oxford University Press.
- Hasbrouck, J. 1995. "One Security, Many Markets: Determining the Contributions to Price Discovery." *Journal of Finance* 50: 1175–1199.
- Hendershott, T., Jones, C. M., and Menkveld, A. J. 2011. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66: 1–33.
- Hendershott, T., and Riordan, R. 2013. "Algorithmic Trading and The Market for Liquidity." *Journal of Financial and Quantitative Analysis* 48: 1001–1024.
- Hespanha, J. P., Ateskan, Y. S., and Kizilocak, H. H. 2000. "Deception in Non-Cooperative Games with Partial Information." In *Proceedings of the 2nd DARPA-JFACC Symposium on Advances in Enterprise Control*.
- Huber, J. 2004. "A Contribution to Solving an Old Puzzle: Why Different Trading Strategies Persist in Competitive Markets." *Journal of Academy of Business and Economics* 3: 171–190.
- Johnston, M. 1988. "Self-Deception and the Nature of Mind." In *Perspectives on Self-Deception*, edited by McLaughlin, B. and Rorty, A. O., 63–91. Berkeley: University of California Press.
- Jovanovic, B., and Menkveld, A. J. 2010. "Middlemen in Limit-Order Markets." Accessed on December 10, 2015. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1624329
- Kane, E. J. 1997. "Ethical Foundations of Financial Regulation." *Journal of Financial Services Research* 12: 51–74.
- Kearns, M., Kulesza, A., and Nevmyvaka, Y. 2010. "Empirical Limitations on High Frequency Trading Profitability." *Journal of Trading* 5: 50–62.
- Kumiega, A., Neururer, T., and Van Vliet, B. 2014. "Trading System Capability." *Quantitative Finance* 14: 383–392.
- Kyle, A. S., and Viswanathan, S. 2008. "How to Define Illegal Price Manipulation." *American Economic Review: Papers & Proceedings*. Accessed on December 10, 2015. <http://w4.stern.nyu.edu/finance/docs/pdfs/Seminars/081m-kyleviswanathan.pdf>
- Lewis, M. 2014. *Flash Boys: A Wall Street Revolt*. New York: W. W. Norton & Company.

- Longstreth, B. 1986. *Modern Investment Management and the Prudent Man Rule*. Oxford: Oxford University Press.
- Massachusetts, State of, Harvard College v. Armory. 26 Mass (9 Pick) 446. 1830.
- Mele, A. 2001. *Self-Deception Unmasked*. Princeton, NJ: Princeton University Press.
- Menkveld, A. J. 2013. "High Frequency Trading and the New Market Makers." *Journal of Financial Markets* 16: 712–740.
- Moor, J. H. 2006. "The Nature, Importance, and Difficulty of Machine Ethics." *Intelligent Systems* 21: 18–21.
- Nielsen, R. P. 2010. "High-Leverage Finance Capitalism, the Economic Crisis, Structurally Related Ethics Issues, and Potential Reforms." *Business Ethics Quarterly* 20: 299–330.
- Roll, R. 1984. "A Simple Implicit Measure of the Effective Bid–Ask Spread in an Efficient Market." *Journal of Finance* 39: 1127–1139.
- Rosenburgh, M., and Spieler, A. C. 2009. "21st Century Pensions: The Risk, the Hedge and the Duty to Consider." *Journal of International Business and Law* 8: 45.
- Ryan, L. V., Buchholtz, A. K., and Kolb, R. W. 2010. "New Directions in Corporate Governance and Finance: Implications for Business Ethics Research." *Business Ethics Quarterly*, 20: 673–694.
- Stoll, H. R. 2006. "Electronic Trading in Stock Markets." *Journal of Economic Perspectives* 20: 153–174.
- Unger, S. H. 1994. *Controlling Technology: Ethics and the Responsible Engineer*. New York: John Wiley & Sons.
- U.S. Securities and Exchange Commission (SEC). 2014. "Best Execution." Accessed on December 10, 2015. www.sec.gov/answers/bestex.htm
- Whaley, B. 1982. "Toward a General Theory of Deception." *Journal of Strategic Studies* 5: 178–192.
- Zhang, Y. C. 1999. "Toward a Theory of Marginally Efficient Markets." *Physica A: Statistical Mechanics and Its Applications* 269: 30–44.