THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A material political economy

Citation for published version:

MacKenzie, D 2017, 'A material political economy: Automated Trading Desk and price prediction in high-frequency trading', *Social Studies of Science*, vol. 47, no. 2, pp. 172-194.
<https://doi.org/10.1177/0306312716676900>

Digital Object Identifier (DOI):

[10.1177/0306312716676900](https://doi.org/10.1177/0306312716676900)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Social Studies of Science

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Material Political Economy: Automated Trading Desk and Price Prediction in High-Frequency Trading

Donald MacKenzie

August 2016

Author's Address:
School of Social and Political Science
University of Edinburgh
Chrystal Macmillan Building
Edinburgh
EH8 9LD
Scotland

DonaldMacKenziePA@ed.ac.uk

A Material Political Economy: Automated Trading Desk and Price Prediction in High-Frequency Trading

Abstract

This article contains the first detailed historical study of one of the new high-frequency trading (HFT) firms that have transformed many of the world's financial markets. The study, of one of the earliest and most important such firms, Automated Trading Desk (ATD), focuses on how ATD's algorithms predicted share price changes. The article argues that political-economic struggles are integral to the existence of some of the 'pockets' of predictable structure in the otherwise random movements of prices, to the availability of the data that allow algorithms to identify these pockets, and to the capacity of algorithms to use these predictions to trade profitably. The article also examines the role of HFT algorithms such as ATD's in the epochal, fiercely contested shift in US share trading from 'fixed-role' markets towards 'all-to-all' markets.

Key Words

High-frequency trading; HFT; Automated Trading Desk; political economy; prediction; fixed-role markets

Charleston, South Carolina, is not where I would have expected to find roots of the high-technology trading that has reshaped 21st-century finance. The past can seem ever-present here: in the beauty of the cobbled streets and antebellum houses; in the shaded paths of the College of Charleston, Spanish moss swathing its live oaks; in the inescapable reminders that this was once North America's busiest slave port. Nor, if I had happened upon it in the early 1990s, would I have been likely to make much of a makeshift, computer-packed office in a 1950s' cinderblock motel building on Charleston's Wappoo Road, occupied by (as one of them put it) 'kids, barefoot, t-shirts, short pants'. Before a decade was out, however, those kids' computers were trading shares worth over a billion dollars every day (Collier, 2002); at its peak, almost one in every ten shares traded in the United States was bought or sold by their firm (Philips, 2013).

This paper presents a history of that firm, Automated Trading Desk (ATD), within the context of a broader study of the rise of ultrafast, high volume, highly automated 'high-frequency trading' or HFT (the type of trading of which ATD was a pioneer). When ATD began operations in 1989, most financial trading still took place directly among human beings, either over the telephone or face-to-face on crowded trading floors. Even as electronic trading became more widespread in the 1990s and beyond, it was initially not usually fully automated: it was conducted by human beings using computer screens and keyboards. After 2000, however, algorithmic trading gained momentum fast, and now dominates several important markets (including the market in US shares).

This article is a contribution to what has become known as ‘social studies of finance’, the application to financial markets not of economics but of wider social-science disciplines such as anthropology, politics, geography, sociology and science and technology studies (STS). STS-inflected work has been particularly prominent within social studies of finance, and much of that work has examined aspects of the transition sketched in the previous paragraph (especially the growth of electronic trading by human beings: HFT is only gradually becoming a focus).¹ For example, Muniesa has examined the pathways via which markets have moved from face-to-face interaction to the algorithmic matching of supply and demand (Muniesa, 2005 and 2011), and Knorr Cetina and Preda have taught us how to conceptualize interaction in electronic markets among human traders and between those traders and their computer screens (e.g., Knorr Cetina and Bruegger, 2002; Preda, 2013). More broadly, Beunza and Stark, for example, have demonstrated how the material layout of trading rooms influences trading, even when that trading is electronic (Beunza and Stark, 2004), while Poon has shown (e.g. in Poon, 2009) how the US mortgage market was reformatted – hugely consequentially – by algorithmic credit scoring. There is also a growing body of literature in STS on the role of algorithms in economic and social life beyond finance: see, e.g., Gillespie (2014). In the background to much STS-inflected work in social studies of finance is Michel Callon’s application to ‘economics’ – a term that, as discussed below, he uses in a very broad sense – of the actor-network theory postulate that all ‘[s]cientific theories, models, and statements... are

¹ See, for example, Borch, Hansen and Lange (2015), MacKenzie et al. (2012) and Seyfert (in press).

performative, that is, actively engaged in the constitution of the reality that they describe' (Callon, 2007: 318; see also Callon, 1998).

STS-inflected social studies of finance could, however, usefully engage more deeply with what one might call the 'political economy' of the financial system.² The term is polysemic. It can connote approaches as different as Marxism and the analysis of political decision-making using rational choice theory; sometimes, it is used simply as a synonym for 'economics'. Here, I intend 'political economy' to flag three closely related issues. First, financial firms are economic enterprises. They make profits, in recent decades rewarding their senior staff – and, less consistently, their shareholders – very handsomely.³ They are also at risk of financial failure (a fate to which Automated Trading Desk almost succumbed three times in its short history). Second, the capacity of financial firms to make profits often depends upon systematic forms of advantage and disadvantage, such as occupation of (or exclusion from) central roles as intermediaries in financial markets. Third, how those markets are organized is a political matter: sometimes explicitly so, when political actors seek to change markets or to preserve their current organization; always implicitly so, because of the consequences of how the

² Clearly, there has been some concern with issues of political economy, for example in Muniesa's history of the automation of the Paris Bourse (see, e.g., Muniesa, 2005). I am arguing, nevertheless, for a greater, more explicit focus on political economy.

³ ATD's remuneration policy was more similar to that of a Silicon Valley start-up than a bank. All employees received stock and/or options, and some secretaries 'retired as millionaires' (Whitcomb interview 2). In ATD's periodic lean times, its chief executives sometimes themselves took no bonus, while awarding bonuses to others.

financial system is organized for the distribution of income and wealth and for economic life more generally.

From the viewpoint of STS, there is no adequate theoretical treatment of political economy to take 'off the shelf'. For example, Fligstein's 'markets as politics' approach (for which see, e.g., Fligstein, 1996 and 2001) helpfully focuses on struggles between incumbent firms and their challengers, and on the influence of the state on the outcomes of such struggles. Even it, though, largely neglects the themes central to STS-inflected social studies of finance, above all the STS emphasis on materiality: on artefacts and technical systems, on human bodies, on mathematical models not just as sets of equations but as material computational procedures (see MacKenzie, 2009).

So how might a 'material political economy' of high-frequency trading and other forms of automated trading be developed? This article sketches such a political economy, focusing on a thoroughly material question: how do HFT algorithms predict price movements?⁴ The question is deeper than it might appear. The efficient market hypothesis of financial economics suggests that, at any point, all currently publicly known information has already been incorporated into prices, so future price movements can be influenced only by future new information, which (if it is genuinely new) cannot

⁴ Algorithms can trade profitably without predicting prices, for example by identifying 'arbitrage' opportunities (in which a financial instrument can be bought more cheaply on one trading venue than it can be sold for on another). However, price prediction was central to ATD's algorithms and remains crucial for HFT more widely.

be predicted. That, in effect, means that price movements in an efficient market are an inherently unpredictable random walk. ('Efficiency' is a model, not an empirical reality, but much of the time the US stock market seems reasonably efficient, in the sense that predicting prices or 'beating the market' is very hard.) For profitable price prediction to be possible, pockets of 'structure' must therefore exist within the random movements of prices. ATD found at least two pockets of powerfully predictive structure; both were created by conflictual 'political economy' processes.

The existence of pockets of 'structure' is a necessary but not sufficient condition for successful predictive trading. Trading algorithms must also be able to identify concrete instances of potentially profitable structure quickly enough to exploit them. That is a matter of what data are available, to whom (or to what), and when. These too are questions of political economy. An analogy used by one of my ATD interviewees nicely captures the role of data in automated price prediction: 'It's as if someone puts certain game pieces on the table, and it's like, okay, I've got these pieces of data, let me look at all the different ways I might use them'.

The political-economic point is that not all 'game pieces' *are* put on the table. For example, very useful for prediction is a market's 'order book' (a file, now always electronic, of the bids and offers that have not yet been executed). For instance – to put it simply – if there are many more bids to buy a particular stock than offers to sell it, then its price is more likely to rise than to fall. However, up until the 2000s, access to the order books of the world's most important share-trading venue, the New York Stock Exchange (NYSE), was the fiercely guarded prerogative of the NYSE's 'specialists'. Each stock

traded on the NYSE had a designated specialist, a trader who, sometimes aided by clerks, maintained the order book for it and matched bids to buy and offers to sell, or if there was no match could trade himself (nearly all specialists were men). The specialist's privileged role brought with it the responsibility to ensure orderly, continuous trading – which could mean having to accept temporary losses – but specialists earned very healthy overall profits (see, e.g., Coughenour and Harris, 2003).

Which 'game pieces' should be on the table is typically political in at least two senses: first, it is contested; second, government regulatory bodies (especially the chief regulator of the US stock market, the Securities and Exchange Commission or SEC) increasingly intervene in it. Similarly 'political' in these two senses are the actions available or not available to a human trader or trading algorithm, for instance if the latter predicts a price move. For example, by 1990 Automated Trading Desk's algorithms could send orders via a dedicated telephone line from Charleston to computer terminals in the trading-floor booths of the NYSE specialists. Once there, orders could in principle have been executed entirely automatically and immediately, but often weren't. Specialists could set up their systems so that an order could not be executed without them manually pressing 'enter' on the terminal's keyboard. That might seem the most minor of material details, but it gave specialists 'a ton of discretion', as interviewee DM put it, to delay execution until a more convenient or – for them – profitable moment. (Single letters are used for ATD interviewees, and double letters for those in my wider dataset: see Table 1.)

The capacities of ATD's algorithms to predict and to act were thus matters of political economy that went far beyond Charleston. They involved wider conflicts over how US stock trading should be organized (and, as we shall see, also over what it should be legal for futures markets to trade). In ATD's first decade, the organization of the financial system was a matter external to ATD, an environment to which its algorithms had to be adapted. Actor-network theory, however, reminds us that the 'micro' (here, the capacities of ATD's algorithms) and the 'macro' (the overall organization of US share trading) are not fixed in their relative scales: the micro can become macro, and vice versa (Callon and Latour, 1981).

And so it was with ATD. The actions of algorithms began to have effects on the organization of the US financial system. ATD's algorithms – along with similar algorithms deployed by other nascent HFT firms – played a causal role in a shift from 'fixed-role' markets towards 'all-to-all' markets. (I take the term 'fixed-role' from Aspers, 2007, but employ it differently.)⁵ The New York Stock Exchange, as late as the 1990s, was an example of a fixed-role market. Even a sophisticated financial firm such as ATD could not trade directly on it, unless it bought an NYSE membership, which was very expensive. To get its orders into the NYSE's trading rooms, ATD therefore

⁵ For Aspers (2007), a fixed-role market is one in which some actors are only ever sellers and others only buyers. In the fixed-role financial markets discussed here, any actor can both buy and sell, but how they can do so is constrained. Note also that an 'all-to-all' market may still involve intermediaries, especially marketmakers (see below). However, marketmakers no longer have special privileges or responsibilities, and there are no formal barriers to any market participant acting as a marketmaker.

had to pay steep fees to a firm that was a member. Within those trading rooms, 'specialists' had the distinct role, privileges and responsibilities that I have just sketched.⁶ An 'all-to-all' market, in contrast, is one in which any participant can trade with any other participant, and there are no fixed, privileged roles. No financial market is completely all-to-all (in many, indeed, fixed roles remain largely intact), but US share trading has moved substantially towards all-to-all, and ATD was centrally involved in that shift.

Although it is not the article's main theme, the history of ATD also allows me briefly to address twin misleading interpretations of the performativity of economics. First, it is not always understood that for Callon 'economics' does not refer only to the academic discipline, and certainly not just to economic theory. Economics includes 'operating methods, calculation tools, and technical instruments... algorithms...', including those that have no basis in academic economics (Callon, 2007: 334). Second – and obviously relatedly – performativity has been misinterpreted as in effect a version of the discredited 'linear model' of innovation, in which technologists simply 'apply' science. For example, in an otherwise insightful analysis of the global financial crisis, Ewald Engelen and colleagues view performativity as implying 'some kind of rationalist application of, or formatting by, prior theoretization ... [a] rationalist grand plan'. They then contrast performativity with 'bricolage ... the creative and resourceful use of materials at hand' (Engelen et al., 2011: 51).

⁶ The classic sociological study of the trading rooms of the NYSE in their prime is Abolafia (1996).

If a single proposition is core to the STS view of innovation, it is surely that *all* innovation is bricolage. Lévi-Strauss's (1966) distinction between the myth-making bricoleur and the rationalist scientist is untenable, as Barnes pointed out over forty years ago (Barnes, 1974: 58-59). While it is perfectly reasonable to view HFT's history as involving performativity – as sketched in the conclusion below, algorithmic practices have reshaped markets so that they have become sociotechnical environments honed to the demands of those practices⁷ – but, in the case of ATD (and, indeed, the other major HFT firms represented in my interview sample), it is a history of bricolage, not of 'prior theoretization'.

ATD's co-founder, David Whitcomb, was an academic economist, a specialist in 'market microstructure' theory (see, e.g., Cohen, Maier, Schwartz and Whitcomb, 1981). In principle, ATD could have implemented that theory in its algorithms: the classic focus of microstructure theory is marketmakers (traders, such as the NYSE's specialists, whose role is to 'provide liquidity', constantly quoting a 'bid' price at which they will buy the asset being traded and a higher 'offer' price at which they will set it), and ATD was primarily an automated marketmaker, albeit not an official one.⁸ However, ATD's

⁷ Algorithms are not representations, and so their performativity does not generally take the form of a representation becoming 'more true' by being employed in economic practice (as in MacKenzie, 2006). Nevertheless, algorithmic prediction is still an epistemic matter, because (as Martha Poon put it to me) 'being right' – making accurate predictions – is an epistemic quality.

⁸ ATD's systems often simply posted either a bid or an offer, not both simultaneously, as an official marketmaker normally has to.

algorithms were *not* implementations of theoretical models of marketmaking. Whitcomb's immediate inspiration was more mundane: his previous efforts to predict the outcomes of horseraces. More broadly, ATD survived its three brushes with bankruptcy precisely because it was a collective bricoleur, prepared to make 'creative and resourceful use of materials at hand'.

This article has three main sections, in addition to this introduction, a discussion of data sources, and the conclusion. The first main section, 'A capsule history of Automated Trading Desk', sketches the overall history of ATD until 2007 (when it was acquired by Citigroup for \$680 million, and its focus shifted away from HFT). The next two sections then delve in greater depth into how ATD's algorithms predicted prices. Those two sections flesh out the 'material political economy' of price prediction that this paper proposes, and the second of them (along with the conclusion) also sketches how ATD's algorithms contributed to the shift in US share trading from fixed-role markets towards all-to-all markets.

Data sources

Investigating high-frequency trading empirically poses considerable challenges. With the exception of a small number of datasets made available to financial economists by exchanges, and datasets privately accessible to researchers working for market regulators, there are no direct traces in publicly-available data of whether any given bid or offer is placed or any particular trade consummated by an HFT firm, so quantitative research often has to employ proxies for HFT activity such as the extent of electronic

message traffic on an exchange's system. However, HFT is actually a less secretive domain than the more lurid popular accounts of it might suggest. Although accumulating this total took five years, in aggregate I have conducted 65 interviews with 51 founders, employees or former employees of HFT firms.

The chief difficulty of the interviews was enquiring in adequate depth into the material practices of HFT, which is a natural goal for a researcher who comes to the social studies of finance from STS. HFT firms fiercely protect the intellectual property crystallized in their algorithms, and over-intrusive questioning would have led interviews to be terminated. The solution found was iterative. In early interviews, I began to build up a sense of the predictive techniques and other practices that were common knowledge in the industry, and known to be common knowledge by all experienced practitioners. It was then possible to ask in more detail about matters such as the pockets of 'structure' on which these techniques rely. That in turn led to a broadening of the research: precisely because the material practices of HFT often turned out to depend upon matters of political economy, it was necessary to research not just HFT firms but also the exchanges and other trading venues involved, the influence of regulation, and so on (see Table 1 for the full dataset of interviewees).

I was led to my first ATD interviewee, A, by a contact who had worked for one of the trading venues I was researching: Island (see below). Interviewee A then put me in touch with ATD's co-founder, David Whitcomb, who agreed to three interviews, and introduced me to a second early ATD employee, interviewee B. The latter allowed me to interview him over nearly

two full days. Finally, my broader HFT research led me to a third early ATD staff member (interviewee C).

Oral-history interviewing has two well known perils. First, interviewees may wish to present themselves and their enterprises in an unduly favourable light. Second, their recollection of the details of events decades in the past is bound to be incomplete. My interviewees, however, did not ‘spin’ ATD’s history as entirely successful: they were open about its failures. The problem of partial recall was alleviated by Whitcomb giving me access to letters he had written in ATD’s early years to its shareholders (he was its chief executive and then chair). These letters were not for public consumption (ATD was a privately-held company, not publicly listed), and were frank about ATD’s difficulties: most of ATD’s shareholders were friends or family of Whitcomb, or ATD employees. The letters clarified exact chronology, provided crucial information on the trading experiments discussed below, and (because they often detailed ATD’s profits or losses) helped me understand its history as an economic enterprise. Interviewee B also showed me a large collection of documents, covering primarily technological matters. Another check on the plausibility of what I was told by my ATD interviewees was the wider interview sample (see Table 1), many of whom worked for similar, albeit later-established, HFT firms, and some of whom had knowledge of ATD’s trading because they were involved in running trading venues in which ATD was active.

A capsule history of Automated Trading Desk

Whitcomb's initial plan for ATD was not what has become known as high-frequency trading but the provision of what would now be called 'execution algorithms' that institutional investors could use to break up big orders into small parts and execute them automatically. In the 1980s, he was hired as a consultant by the pioneering screen-based share-trading system Instinet. First established in 1969, Instinet was intended to allow institutional investors to trade directly with each other, circumventing expensive fixed-role markets such as the New York Stock Exchange. Take-up was, however, disappointing, and Whitcomb proposed to Instinet that it should provide its clients with execution algorithms to use on the system (Whitcomb, interviews 1 and 2).

Instinet did not adopt Whitcomb's suggestion. However, a former student of his, James Hawkes (who taught statistics at the College of Charleston, and ran a small firm, Quant Systems, which sold software for statistical analysis), traded stock options, and had installed a satellite dish on his garage to receive a stock-price datafeed. Hawkes mentioned this to Whitcomb, and the two men decided on a joint venture, which they christened Automated Trading Desk. Whitcomb raised most of the new firm's capital of \$100,000 (see, e.g., Whitcomb, 1989a); Hawkes provided the programmers, two College of Charleston students who wrote statistical software for him.

Whitcomb already had in mind execution algorithms that had predictive capacities and could thus optimize how they sent orders to market. He and Hawkes had earlier collaborated on a regression-equation model to predict

the outcomes of horseraces. Bookmakers' large 'vigs' or 'takes' (the profits they earn by setting odds unfavourable to the gambler) meant the model did not earn Hawkes and Whitcomb money, but the equation displayed some predictive power (Whitcomb interviews 1 and 2).

Could an analogous model be developed to predict stock prices?

Whitcomb set to work, designing the relatively simple regression-style model described in the next section, faxing instructions and formulae to Charleston to be turned into code by Hawkes's programmers; Whitcomb, who taught finance at Rutgers University, lived in New York. (I write 'regression-style' because, although Whitcomb's model had the mathematical form of a regression equation, in its first version its coefficients were simply Whitcomb's informed guesses rather than estimated statistically.) Around the equation, Whitcomb designed and the programmers coded the components of a full automated trading system: a module to process incoming market data; a 'pricing' module implementing Whitcomb's equation; a module that tracked the system's accumulated trading position in each stock, and adjusted its trading strategy accordingly; a decision module that calculated how best to trade based on the existing trading position and the pricing module's predictions; a module that dispatched the resultant orders and if necessary cancelled existing orders; a module that calculated in real time the profits or losses being made; and so on – altogether some 80,000 lines of software (interviewee B).

ATD was unable to find a mainstream investment-management firm willing to risk being the first to use an execution algorithm, perhaps particularly one as sophisticated as ATD's system would make possible: 'many

institutions ... said that's a very interesting idea, David. I'd be interested in being your second customer' (Whitcomb interview 2). In the summer of 1989, however, a trading simulation convinced Whitcomb that ATD could make money trading on its own behalf, not just by helping clients trade. Whitcomb's accountant introduced him to a man (whom he does not wish to have named) who ran what has subsequently become a well-known hedge fund. It provided the capital for experiments in automated trading, described in the next section, that were loss-making – the simulation 'overstated predicted trading profits by 200%', says Whitcomb (email to author, 24 March 2016) – but demonstrated that Whitcomb's model had predictive power. That persuaded a leading investment bank, with all the advantages that came from its central role in the financial system, to employ ATD and its model in trading on the New York Stock Exchange (again, Whitcomb has requested that I do not name the bank).

ATD's failure to find clients and its loss-making experiments meant that it earned almost no money in 1989 or 1990 (Whitcomb, 1990a & b). It survived only because Whitcomb and his wife invested additional capital and Hawkes's two programmers (who became central to ATD, and eventually took on leading management roles in the firm) agreed to receive half their pay in newly-issued ATD shares, along with existing shares contributed by Hawkes. However, the collaboration with the investment bank made ATD profitable. It moved into offices in a strip mall in Mount Pleasant, across the Cooper River from Charleston, and essentially a suburb of the latter. 'We're on the top floor of a two story brick building (important in low-lying, hurricane-prone Charleston)', Whitcomb told ATD's shareholders. In September 1989, ATD

had been shut down for two weeks by Hurricane Hugo, a giant storm that made landfall just north of Charleston (Whitcomb, 1995; Whitcomb, 1989b).

ATD was nevertheless soon in financial trouble again. The profitability of its NYSE trading gradually declined, and in 1994 ‘we ... made practically no money’ (Whitcomb, 1995). Salvation came from bricolage: ATD began automated trading in the second of the US’s two main stockmarkets, Nasdaq, employing a predictive technique, described below, that owed little to its earlier model (and nothing to economics in an academic sense). ATD learned the technique from successful – but, as Whitcomb puts it, ‘despised’ – human traders. (Why they were ‘despised’ is explained below.)

Thus began ATD’s *Wunderjahren*, from 1995 to 2001. As the fixed-role structure of US share trading began to erode, the opportunities for automated trading grew. ATD expanded to 23 employees by October 1999 (Wipperfurth, 1999) and in the early 2000s to around 70 (see Table 2). Around ATD’s original pricing module were added new trading algorithms; there were twenty such algorithms by the end of the 1990s (Wipperfurth, 1999). Those algorithms were often written by staff members who were simultaneously traders and programmers, with the simpler scripting languages Perl and Python made available to those not fluent in the C or C++ of ATD’s core systems. In ATD’s early days in the old motel, its trading could be conducted by a single computer, with – as interviewee B told me – the various software modules communicating ‘within one computer’s memory’ (although multiple machines were needed for redundancy and to cover the full range of stocks that ATD traded). By the early 2000s, ATD’s trading needed a large-scale

integrated network of ‘distributed multiple programs across multiple machines’, the latter numbering in the hundreds (interviewee B).

Simultaneously, the sophistication of ATD’s automated price prediction was increased. As the variables available as inputs into prediction increased (see this article’s conclusion) from an initial handful employed in its first model to around a hundred (Schmerken, 2011), so the model was elaborated. A new, mathematically different, model was also built. Instead of directly estimating, using a regression equation, a predicted price of the stock in question (as the original model did), the new model employed discriminant analysis – first developed by the British statistician R.A. Fisher (1936) – to estimate whether the price of a stock was most likely to go up, go down, or remain unchanged. The new model also generated an estimate of the level of certainty of its prediction. The certainty level influenced, for example, the choice between placing a new ‘marketmaking’ order that could not be executed immediately (which would be done if the certainty level was ‘Normal’) or the more expensive option of executing against another firm’s existing order, which required ‘High Certainty’ (Whitcomb interview 3; interviewee B’s papers, Pricing Engine file).

At a time when ATD’s competitors were still almost always human beings, with their inevitably slow reaction times, trading with keyboard and mouse, the simple fact that its trading was computerized gave it a considerable advantage. In 1995, ATD had at most a handful of competitors in large-scale automated marketmaking in US shares. Even by 2001, firms such as Chicago’s Getco and Jump, and Kansas City’s Tradebot (all three of which were to become formidable competitors to ATD) were still in the early stages

of their move from HFT in futures to shares. In the first quarter of 2001, for example, ATD traded around 55 million shares a day, and made an average profit of almost 0.9 cents per share (calculated from the figures reported to ATD's shareholders in Whitcomb, 2001), twenty times what HFT interviewees in my larger dataset regard as a healthy profit rate nowadays. ATD built a \$35 million, hurricane-resistant headquarters (designed in modernist, campus-like style, with a reflecting pool and landscaping) in Mount Pleasant, choosing a postal address, 11 eWall Street, that ambitiously echoed that of the – still at core largely manual – New York Stock Exchange.

ATD's success was a matter of local pride, with South Carolina's governor attending the new building's groundbreaking ceremony. When *The State* sent reporter Joe Collier to Mount Pleasant in 2002 to report on ATD, his story was headlined, 'Lowcountry Firm Among High-Tech's Best' (the Lowcountry is South Carolina's coastal region from the border with Georgia northwards). Collier was struck by the incongruously expensive cars in what was still a strip mall's carpark. He met a 21 year-old College of Charleston student who, in two years working for ATD, had earned enough to buy 'a home and a Porsche Boxter' (Collier, 2002).

Even as Collier wrote his story, however, all was again not well with ATD. In 2003, the firm 'recognized a loss of \$16.0 million' (Swanson and Whitcomb, 2004), a huge sum for what was still only a medium-sized enterprise. There were several reasons for the renewed crisis, but – paradoxically – one was a change in the pricing of US stocks that Whitcomb had advocated energetically. The traditional minimum price increment had been an eighth of a dollar, reduced to a sixteenth (6.25 cents) in 1997. These

large 'tick sizes' meant a 'spread' between the prices of the highest bid and lowest offer for a stock of at least 6.25 cents, keeping marketmakers' profits high and imposing non-trivial costs on other traders and investors. There was fierce Wall Street resistance to reducing the minimum increment, but following Congressional testimony by Whitcomb and others, and pressure from Congress on the SEC, the latter mandated 'decimalization': the transition to pricing in dollars and cents, a transition that was completed by April 2001. Whitcomb had expected spreads to stabilize at around 2.5 or 3 cents; almost immediately, they collapsed, for most heavily-traded stocks, to a single cent (Whitcomb interviews 2 and 3).

It was a decisive moment in the shift from human to algorithmic trading: with spreads of a single cent between the highest bid and lowest offer, marketmaking by human beings was no longer economically viable. Yet the shift hurt ATD. It was, after all, itself a marketmaker, even if unofficial, and smaller spreads reduce marketmakers' revenues. Furthermore, there had previously been profit opportunities in the difference between the coarse pricing grids of mainstream venues such as Nasdaq and the finer grids of new venues, especially Island (see below), whose minimum increment was $1/256^{\text{th}}$ of a dollar.

Another reason ATD's trading revenues declined from mid-2001 onwards was that increasing numbers of other firms had also begun high-frequency trading. ATD soon came to realize that its systems, though faster than the fastest human, were most likely slower than those of some of these new competitors. The firm had been a happy workplace. Although HFT was demanding and often stressful, it had the enjoyable excitement of a high-

stakes, immersive game. ATD's co-founder James Hawkes, who stepped aside after around a year to concentrate on his other firm, Quant Systems, told Charleston's *Post and Courier* two decades later why he regretted leaving: 'I gave up my dream ... to play a game for money that I could use a computer to play with' (Kearney, 2012). ATD's young trader-programmers would 'take ... out frustration and stress' at the end of the trading day by using their powerful, interconnected computers to play 'shooter games, Counterstrike and Doom and Quake' before 'head[ing] back to their office and ... crunch on numbers and work on code'. They would uncomplainingly take on (and even regard as 'fun ... unlike today') the material bricolage required to build a pioneering trading system: 'these were days when I was on a ladder running cable through the ceiling ... days ... splicing cables' (interviewee B; interviewee C). By 2003, however, the fun seemed to have evaporated, as interviewee B's files reveal: 'From the troops ... Overworked ... No bonuses/need more money ... No end in sight ... Unrealistic expectations/demands' (Minutes for weekly Team-Leader meeting, September 3, 2003).

ATD responded by cutting costs and raising capital (it sold and leased back its new headquarters). It stepped up efforts to improve its pricing engine, and set up 'a taskforce to attack latency' (Whitcomb interview 2): to eliminate delays in its systems. Considerable speed-up was achieved, but ATD's war on latency also caused collateral damage. Updating an algorithm, a young trader-programmer accidentally interchanged a plus sign and a minus sign. 'Unfortunately, the ... error was in the interpretation of inventory [the program's holdings of shares]'. Rather than keeping inventory safely close to

zero – as was the intention – the program therefore increased it in ‘geometric progression... [it] doubled and redoubled and redoubled and redoubled’. It took only 52 seconds for the trader to ‘realize ... something was terribly wrong and he pressed the red button [to stop trading]. By then we had lost \$3 million’ (Whitcomb interview 2). ATD was sharply aware of the risks of errors in programs and systems – the firm had suffered an almost crippling loss in the late 1980s, when a flaw in Instinet’s systems allowed what Whitcomb had believed to be a trading simulator to make real trades – and ATD’s systems had had risk controls that would have stopped the geometric progression before it caused serious damage. Those controls, however, had been removed in the efforts to reduce delays.

Most important to ATD’s survival in the 2000s was that once again it found a ‘creative and resourceful use of materials at hand’ (Engelen et al., 2011: 51). It deployed the modelling and technological expertise built up in its HFT to become one of the first of a new generation of high-technology ‘wholesalers’: firms that act as marketmakers for orders from retail investors. Unlike professional market participants, members of the general public rarely possess information not already known to the market at large. Even with one-cent spreads, therefore, a wholesaler whose costs are low enough (because, like ATD, its marketmaking is automated) can afford to pay retail brokerages to send it their customers’ orders, and still make a steady profit executing those orders.

Although these retail orders had to be processed quickly by ordinary human standards, and automated price prediction was still needed in order for algorithms to decide whether to fulfil an order internally or send it on to the

public markets, the demand for speed was much less than in HFT in those markets: ‘you literally have hundreds of milliseconds or maybe up to a second to respond’ (interviewee A). Wholesaling not only enabled ATD to survive the decline in its HFT revenues but persuaded Citigroup to buy the firm: ATD’s ‘wholesale marketmaking arm ... is what Citi acquired’, says interviewee A. For nine years Citi kept ATD’s Mount Pleasant offices open, but in May 2016 Citi sold ATD’s wholesaling business to the Chicago hedge fund Citadel, which closed it, merging it into its own wholesaling activities.

Predicting prices in fixed-role markets

Let me now return to the central thread in ATD’s history as an independent firm, evident from its first trading through to its successful transformation into a wholesaler: price prediction. As noted, the immediate antecedent to Whitcomb’s efforts to predict stock prices was his and Hawkes’s horserace predictions, and the two efforts had the same mathematical form: that of a linear regression equation, in which the values of a number of ‘independent’ or ‘predictor’ variables are employed to predict the value of a ‘dependent’ variable.

In the case of horseracing, the variable to be predicted was ‘the horse’s speed at that distance’, and the predictor variables included ‘the weight the horse was carrying today relative to the other horses, the jockey’s previous winning percentage, the horse’s speed in the past relative to other horses, some dummy variables for the kind of race’, all of which were publicly available information (Whitcomb interview 2). In the case of share trading, the

dependent variable to be predicted was what Automating Trading Desk came to call the 'ATV' or 'adjusted theoretical value' of the stock in question, a prediction of its price 30 seconds in the future. (ATD experimented with different time horizons, but found the exact choice not to be critical.) What, though, were the predictor variables that could be used? They too had to be public, and knowable in Charleston without undue delay. (In its later years, ATD – like other HFT firms – placed computer servers in the same buildings as exchanges' systems, but originally all its computing was done in first its Charleston and then its Mount Pleasant offices.)

The trading venue from which it was easiest to obtain data was Instinet. Its subscribers were provided with a terminal linked by a modem and telephone lines to Instinet's central computer systems. On the terminal's green screen, 'supply' and 'demand' were displayed in the form of an anonymous list of the bids to buy and offers to sell each stock. By the late 1980s, disappointing volumes of trading had led Instinet to allow not just institutional investors but also other professional market participants to use its system, although it gave institutional investors the capacity to mark their orders 'I-ONLY', allowing only other institutional investors to see them (Instinet, 1988: 12). ATD subscribed to Instinet, and worked out how to 'screen scrape' (interviewee C): to connect the modem directly to an ATD computer that decoded the incoming stream of binary digits that drove the terminal's screen.

The bids and offers on Instinet were, however, a poor guide to the overall balance of supply and demand for a stock, because Instinet remained marginal: in 1989, only around 6 million shares a day were traded on it (about

3 percent of total share trading). The dominant venue was still the New York Stock Exchange, whose trading rooms handled around 150 million shares every day (McCartney, 1990). Since the 1870s, the NYSE had disseminated, originally via the stock tickers described by Preda (2006), the prices at which shares had been bought and sold and the sizes of transactions. By the 1980s, the prices and sizes of the best (i.e. highest priced) NYSE bid and best (lowest priced) offer were also available electronically.

By the late 1980s, a number of companies had rented capacity on communications satellites to transmit these and other financial data to places, such as Charleston, far from major centres. (For all its 'high-tech' nature, satellite transmission is now rarely used in HFT, because of the longer distances signals must travel and the processing delays in the satellite's transponder.) ATD subscribed to one such service, Standard & Poor's ComStock (a manual for which is still in interviewee B's files: S&P ComStock, 1990). ATD received the signals first via the satellite dish on top of Hawkes's garage – in which the ATD programmers worked in a cubicle – and then via a dish on the roof of the old motel on Wappoo Road.

ComStock, however, could not report the full balance of supply and demand on the NYSE: the full 'book' of unexecuted bids and offers was still private to a stock's NYSE 'specialist'. Initially, therefore, ATD's regression model simply used 'the size of the [best] bid relative to the size of the [best] offer', along with 'a short-term trend variable in the transaction prices of the stock' (Whitcomb interview 3). Later, the firm constructed another proxy for the still incompletely known balance of supply and demand on the NYSE. ATD's system calculated two variables, 'down volume' and 'up volume', which

indicated whether transactions were on average taking place at the best price at which there were bids to buy or at the higher price at which there were offers to sell. If the latter, for example (if, in other words, ‘up volume’ exceeded ‘down volume’), ‘that’s indicating well, gosh, everybody seems to be paying up’ (interviewee B), and thus a price rise was likely.

More important for the predictive capacity of ATD’s algorithms, however, than any of these variables was a variable that – paradoxically – had its origins in the most intensely embodied human trading arenas in the US: Chicago’s crowded open-outcry futures ‘pits’. One such pit, in the Chicago Mercantile Exchange’s twin-towered skyscraper, traded futures contracts based on the Standard & Poor’s 500 index, which tracks the changing prices of the US’s leading stocks. The very existence of these contracts was the result of a prolonged battle by the Chicago exchanges. Although originally they traded futures on grain and other agricultural or physical commodities, they had long wished to expand into the trading of futures on stock indices.

The equally long-standing barrier to this trading was that the law of Illinois – and that of many other jurisdictions in the US and overseas – drew the distinction between a legitimate futures contract and a bet based on whether ‘physical delivery’ could occur. (Gambling was against the law in Illinois and most of the US.) It had to be possible for a futures contract to be settled at its maturity by the seller of the future delivering the underlying asset to the buyer. In practice, physical delivery was rare, but if it was not possible then the courts would most likely interpret the contract as an illegal wager (MacKenzie, 2006; Millo, 2007).

Unlike grain, a stock index is a mathematical construct, and settling a future on an index by the delivery of huge numbers of share certificates would be at best clumsy. However, Leo Melamed, who led the Chicago Mercantile Exchange, had initially trained as a lawyer, and knew that the bans on gambling in the US were state bans, which a Federal body could pre-empt. As he put it in his memoirs, 'buried deep in my head was a plan that only a federal agency could carry out: *remove the requirement of physical delivery in the mechanics of our markets*' (Melamed and Tamarkin, 1996: 216, emphasis in original).

In the early 1970s, Melamed responded to one of the futures markets' periodic crises by supporting the establishment of a new Federal futures regulator. Active lobbying by him and others made possible the 1974 amendments to the Commodity Exchange Act that established the new regulator, the Commodity Futures Trading Commission (CFTC). The new regulatory body did not immediately welcome the abandonment of the traditional way of distinguishing a futures contract from a bet. However, the election of President Reagan changed the regulatory climate, and the subsequent appointment of one of Melamed's allies, Philip McBride Johnson (former Counsel to the Chicago Board of Trade), as CFTC chair facilitated the change. In 1982, the Chicago Mercantile Exchange began to trade S&P 500 futures.

The new contract offered an attractive way of expressing a view on whether overall US stock prices were more likely to rise or to fall, or of hedging against the latter possibility. A verbal exchange or eye-contact and hand signals between two traders in the S&P pit in Chicago was quicker and

simpler than trying to buy or sell hundreds of stocks on the NYSE or Nasdaq (especially given the regulatory constraints on short sales: sales of shares one does not yet own). Far less capital needed to be put up in advance to back up a futures trade than a share purchase of equivalent size, giving futures greater ‘leverage’, as traders put it. It quickly became clear that, in consequence, new information relevant to the overall value of US stocks usually first influenced Chicago futures prices and only slightly later the prices of the underlying stocks (for econometric evidence of this, see Hasbrouck, 2003, and the literature cited there).

‘Pit reporters’ employed by the Chicago exchanges turned traders’ shouted or hand-signalled deals into an electronic stream of futures prices. That stream flowed to data services such as ComStock, and from there into ATD’s computers. Futures prices were ‘the prime market indicator that we were using’, as Whitcomb puts it; ‘definitely the key variable’ (interviewee B). So important were they, that ATD’s system had a ‘futures only’ mode, in which all other predictors were switched off.

Whitcomb’s model, incorporating futures prices and the other variables discussed above, displayed predictive capacity in its first test, in September 1989. Because Instinet was an electronic system, and ATD could trade directly on it, it was the venue chosen for this test. As Whitcomb reported to ATD’s shareholders, the model ‘was not perfect’, and ATD ‘did some fine tuning while we were trading’, but ATD’s experiment still earned gross trading revenues of around 2.5 cents per share traded (Whitcomb, 1989b). To use Instinet, however, ATD had to pay commissions of around 4.5 cents per share.

A portion of those commissions was then returned to ATD as a ‘soft dollar’ reward for using Instinet. (‘Soft dollars’ were, and are, one of the reasons big broker-dealers,⁹ often divisions of major investment banks, kept central roles in share trading. In return for an investment-management firm paying trading commissions to a broker-dealer, it received – and receives – ‘free’ research reports, along with, in the past, perks such as subsidized travel and sometimes outright cash rewards. The economic rationale is that trading commissions are charged not to investment-management firms but to the pension funds, mutual funds and other savings that the firms manage; such commissions are therefore not ‘hard dollars’ – direct expenses – from the viewpoint of investment-management firms. However, cash rewards and other benefits – getting ‘free’ research saves hard dollars – went directly to those firms.) Even Instinet, although designed to allow investment-management firms to trade cheaply with each other, had to offer those firms soft dollar incentives to do so. However, the hedge fund sponsoring the trading experiment required ATD to pass Instinet’s soft dollar payments to it (Whitcomb, 1989b).

ATD had, of course, been aware of Instinet’s commission rates. A problem it had underestimated, however, was that in opening up its system, Instinet had allowed broker-dealers access. Their central roles in all forms of US share trading meant that they possessed information (about their clients’ intended trading, for example) that could not be deduced from an Instinet screen. As Whitcomb told ATD’s shareholders, ‘broker-dealers are always

⁹ A broker executes orders on behalf of a client; a broker-dealer also trades on its own behalf.

active in the market and possess the most current “fundamental” information and orderflow information for the stocks they trade’. Instinet’s own traders (who acted for firms not prepared to use its system directly) told ATD that it was getting ‘bagged’: broker-dealers would buy from ATD when they had information that meant prices were likely to rise, and sell to it when they expected prices to fall (Whitcomb, 1989b). In effect, broker-dealers’ occupancy of a central role in a fixed-role market gave their human traders predictive capacities that, at crucial moments, were greater than those of ATD’s algorithms.

For its next trading experiment, ATD shifted from the convenient but peripheral Instinet system to the New York Stock Exchange. ATD could trade on the NYSE only via a member firm, and these firms’ usual commissions – which in the late 1980s averaged nearly 7 cents per share (Berkowitz, Logue and Noser, 1988: 104) – would have rendered it impossible for ATD to trade profitably. However, Whitcomb found an NYSE member, a major investment bank, that had an internal trading group that also used predictive models in share trading (albeit with a much longer time horizon than ATD’s), and therefore understood immediately that because ATD’s trading was automated, the bank would not have to provide it with many of the services that most clients needed. In the light of this difference, ATD was able to negotiate specially reduced commissions of only 3 cents per share. The investment bank also provided high speed modems, set up two dedicated telephone lines between ATD’s Charleston office and the bank’s Manhattan headquarters, and allowed ATD’s electronic orders to flow from there through its high-speed connection to the specialists’ booths in the NYSE’s trading rooms. During the

test, conducted in April and May 1990, ATD traded 3.2 million shares, and its model again showed predictive power: the firm's gross trading profit averaged 1.9 cents per share (Whitcomb, 1990c). That, however, was less than the commissions it was paying the bank. Again, ATD was trading at a net loss to the hedge fund sponsoring the experiment.

The investment bank, however, knew that it itself could trade at a cost much lower than even the reduced commissions it charged ATD. It therefore proposed to ATD an arrangement via which ATD would pay only an estimate of the bank's actual costs of trading (around 1.4 cents per share), while the bank would keep the bulk of ATD's profits, on a sliding scale that equated to a roughly 75:25 split between it and ATD (Whitcomb, 1990c). As noted above, it was this arrangement that allowed ATD to begin trading profitably, and thus secured – at least temporarily – the firm's survival.

Viewed through the lens of what in 1990 was still an almost intact fixed-role market, the arrangement could be viewed as equitable. It 'was a very fair and honourable deal', says Whitcomb: 'I've only praise for [the investment bank]' (interview 2). By 1990, however, a tiny, initially scarcely noticed breach in the fixed-role system had been created. Within less than a decade, it was to start the system unravelling, and ATD was to be at the heart of the process.

'Among the despised'

The breach was in Nasdaq, the National Association of Securities Dealers (NASD's) Automated Quotation System. Set up in 1939, the NASD was an

SEC-encouraged response to widespread fraud in ‘over-the-counter’ – i.e., not exchange-based – share trading. (The NYSE and other exchanges imposed requirements on companies seeking to list their shares, and shares that traded ‘over-the-counter’ were originally usually those whose issuers could not meet those requirements. Later, technology companies such as Apple, Microsoft, Intel and Cisco, which could have listed on the NYSE, often chose Nasdaq instead.) Launched in February 1971, again in part in response to pressure from the SEC for greater price transparency, Nasdaq was an electronic system for the on-screen dissemination of price quotations (bids to buy and offers to sell) from the NASD’s authorized marketmaking firms.

In order directly to trade on Nasdaq, a securities firm had to become a member of the National Association of Securities Dealers, which involved ‘enormous bureaucratic hurdles’, Whitcomb told ATD’s shareholders (Whitcomb, 1995). Once accepted into membership, a sufficiently well-capitalized firm could then register with the NASD as a marketmaker for one or more stocks. Only then did a firm get the ‘Level III’ access, via a Nasdaq terminal, needed to post bids and offers on Nasdaq’s screens. To trade without that access, even an NASD member firm had generally either to strike a deal by telephone with a marketmaker or use Nasdaq’s Small Order Execution System (SOES, set up in 1982), via which a member could send orders – typically from retail customers – for a thousand shares or fewer to be executed automatically against a marketmaker’s quote. (Unlike an exchange such as the NYSE, the NASD never had a face-to-face trading floor.)

In a period in which the minimum unit of price (and thus the minimum ‘spread’ between the prices at which a marketmaker would buy a stock and

sell it) was still an eighth of a dollar, being a Nasdaq marketmaker was a profitable business. The role brought with it both formal obligations and an informal norm – vigorously policed by the harassment by their fellow marketmakers of those who violated it – not to display ‘odd-eighths’ price quotes such as $\$20\frac{1}{8}$, $\$20\frac{3}{8}$, etc. In the racist terminology then still common, to post a bid or offer with an odd-eighth price was to ‘make a Chinese market’, and (unless the bid or offer was very fleeting) could trigger an abusive telephone call from another marketmaker. The effect of the norm was to widen the ‘spread’ between the highest bid and lowest offer (which is the main source of marketmakers’ revenues) from 12.5 cents to a typical 25 cents (Christie and Schultz, 1994).

The breach in this fixed-role system was created quite inadvertently by the SEC. During the 1987 stock market crash, many Nasdaq marketmakers – fearing continuing precipitous price falls – stopped processing SOES orders. After the crash, the SEC successfully pressured the NASD to make it obligatory for marketmakers to fill SOES orders at the prices they were displaying on Nasdaq’s screens. That ruling opened the breach. If a marketmaker’s staff were not monitoring Nasdaq’s screens sufficiently attentively, they might not alter their price quotations fast enough as market prices changed — and traders who were paying closer attention could then use SOES to send orders (which a marketmaker now *had* to fill) that ‘picked off’ these stale quotes, for example buying shares at a price that the

marketmaker had not increased as prices rose.¹⁰ Increasingly large numbers of semi-professional traders (pejoratively dubbed ‘SOES bandits’ by the official marketmakers) seized these opportunities, using Nasdaq screens and SOES access provided – in crowded, sometimes makeshift trading rooms, often in rundown buildings in lower Manhattan – by ‘day trading’ firms that had succeeded in becoming NASD members. By the mid-1990s, there were more than 2,000 such ‘bandit’ traders (Harris and Schultz, 1998: 41).

SOES gave these outsiders a direct route into the heart of a major fixed-role market. Nasdaq’s marketmakers tried everything they could to seal the breach – trying to bar access to SOES by ‘professional’ traders, trying to get the SEC’s permission to replace SOES with a new system without compulsory execution of trades, even death threats to individual bandits – but nothing worked. (Because Nasdaq was designed primarily to facilitate trading by telephone, it could not be anonymous. The firms displaying orders on-screen and those sending them via SOES were all identifiable, their telephone numbers were easily to hand, and in practice the identities of individual leading ‘bandits’ were well known to Nasdaq’s marketmakers.) Tempers flared on both sides. Whitcomb, for example, watched, horrified, as a leading bandit

¹⁰ An NASD rule barred marketmakers from installing an automated system to update their quotes in the light of changes in other marketmakers’ quotes. According to interviewee EZ, who worked for the NASD in the 1990s, the rationale was to stop marketmakers evading their responsibilities by automatically ‘fading’ their bids and/or offers: altering their quotes so that they would never be executed.

‘in effect threatened ... physically’ a top NASD official: ‘He made moves towards him while using exceedingly profane language’ (interview 2).

With the profits of ATD’s early-1990s’ NYSE trading shrinking, Whitcomb had realized ‘we needed a completely different act’ (interview 2). An acquaintance of Whitcomb’s, a professor of accounting, had a former student who had become a ‘SOES bandit’. ATD formed a joint venture with the professor and the trader, seeking to develop what Whitcomb calls ‘an automated SOES bandit system’ (Whitcomb interview 2; in his letters to ATD’s shareholders, Whitcomb used the more neutral term ‘SOES activist’). ATD’s traders ‘just sat down’ with bandits, watching what they did and asking them why they did it (interviewee C). It quickly became clear that ‘bandits’ predicted price changes by carefully monitoring changes in the array of marketmakers’ (non-anonymous) bids and offers on Nasdaq’s screens. If, for example, even a small number of marketmakers lowered their bids (and especially if they were marketmakers whose actions were regarded by ‘bandits’ as likely to presage price movements in the stock in question), then ‘bandits’ would use SOES to sell as quickly as they could. ‘[T]hey’d [the ‘bandits’ would] be like, “two ... people [marketmakers] left [lowered their bids]” or “Goldman leads in this stock, and when Goldman leaves, everybody leaves, so I saw Goldman leave [lower its bid], so I hit it [i.e. sold the stock]”’ (interviewee C).¹¹

ATD automated ‘bandit’ predictive reasoning of this kind. ‘Our computer scans the NASDAQ digital data feed for several hundred stocks’,

¹¹ That ‘bandit’ reasoning was of this general form, and was profitable, was later confirmed by the quantitative analysis by Harris and Schultz (1998).

Whitcomb told ATD's shareholders in January 1995, 'watching for indications that the market makers in a stock are about to "fade their quotes" on one side [i.e. either bids or offers].' Despite the distance between ATD's Mount Pleasant headquarters and Nasdaq's computer centre in Trumbull, Connecticut, ATD's 'automated SOES bandit' was faster than the human beings whose predictive reasoning it mimicked. 'What we were competing against [was] the ... SOES guys', says interviewee C. 'Goldman leaves the offer, there would quickly be a thousand people trying to hit the offer [i.e. buy shares] ... most of them ... "point and click" [manual traders]'. HFT was still only nascent in the mid-1990s, and in my wider dataset I found only one other firm that had operated an automated SOES system akin to ATD's (interviewee AG).

The pocket of predictive 'structure' that human and automated SOES 'bandits' exploited was created by the sociotechnical organization of Nasdaq: as noted, only registered marketmakers could post bids and offers on Nasdaq's screens, and those bids and offers were therefore limited in number and not anonymous. The most likely reason why monitoring them enabled 'bandits' to predict price movements was that marketmaking firms were often also broker-dealers that executed large orders from institutional investors (Smith, Selway and McCormick, 1998: 34). If, for example, a marketmaking broker-dealer was in the process of executing a large sell order (or had learned that such an order was being executed), then it would lower its bid prices to avoid buying shares at a price that was likely to fall. In effect, Nasdaq's non-anonymous arrays of on-screen bids and offers thus broadcast, to anyone prepared to monitor them attentively enough, broker-dealers'

private information – the information that had been the undoing of ATD's Instinet trading.

Being 'among the despised' SOES bandits, as Whitcomb puts it (interview 2), was sometimes uncomfortable for ATD ('we knew – or we felt – that NASD would like nothing better than to shut down a SOES bandit firm for a violation'), but this form of trading was crucial to ATD's survival in the mid-1990s: 'it saved us', says Whitcomb (interview 3). It was perfectly possible sometimes to make a profit of 25 cents per share traded: 'you just made, whatever, \$250 off of them [on the maximum SOES order, for 1,000 shares]. So they [Nasdaq's marketmakers] hated us' (interviewee C).

More important, however, in the long run was ATD's involvement in a set of new trading venues established in the mid-1990s, often initially to serve the needs of 'bandits' and other 'day-traders' of shares listed on Nasdaq. While 'bandits' used SOES to create their trading positions, they could not, unless they were very lucky, use it profitably to close those positions (Harris and Schultz, 1998), so needed other trading venues on which to do so. The most important of these was Island, set up in 1995, based in the lower Manhattan offices of Datek (a leading 'bandit' firm). Island was led by a talented programmer, Josh Levine, who combined immersion in bandits' trading culture with a deep commitment to 'information libertarianism' (see MacKenzie and Pardo-Guerra, 2014).

Any Island user could place bids and offers direct into Island's order book, and Levine provided a fast communications protocol, OUCH, for doing this algorithmically. Execution of orders was fully automatic and also very fast

– there was no equivalent of an NYSE specialist having first to press ‘enter’ – and (equally crucially from the viewpoint of automated marketmaking) orders that had become stale could quickly be cancelled. There was no need to ‘screen scrape’: Island provided ‘ITCH’, a fast direct datafeed which a computer could use to construct a mirror of Island’s order book. In short, Island was an environment (at its core, a technical system) materially suited to trading algorithms.

Island’s origins in the world of ‘SOES bandits’ caused it initially to be viewed as disreputable: Island staff were told ‘never, ever [to] mention ... who was a client of Island ... people did not want it to be known that they were trading on Island’, says interviewee C. ATD had no such prejudice, and quickly saw how well suited the new venue was to algorithmic trading. ATD became the first firm, Datek aside, to trade on Island, and soon became a very heavy user, boosting Island’s trading volumes and thus its attractiveness as a trading venue. (Island initially had a daily limit of 999,999 orders per connection. Interviewee B realized one day that ATD was about to overrun that limit – ‘I was like, oh crap’ – but fortunately had another connection available to him.)

ATD and the other nascent HFT firms also became important marketmakers on the other new share-trading venues (unlike on Nasdaq, any firm could make markets on them), continuously posting bids and offers with prices – often informed by Chicago’s futures prices (interviewee AB) – frequently marginally superior to those available on Nasdaq (Biais, Bisière and Spatt, 2003). Year after year, the new ‘algorithm-friendly’ venues increased their share of trading, until first Nasdaq and then the NYSE had to

copy their features to remain competitive, a process that was facilitated by Nasdaq buying Island and the NYSE buying another of the new electronic trading venues, Archipelago.

Conclusion

A correlate of many historic changes is a transformation of 'common sense'. Even as late as the mid-1990s it seemed 'obvious' that a well-organized financial market needed a fixed-role structure, with (for example) officially designated marketmakers, and traders who 'picked off' marketmakers' stale prices could therefore be seen as doing something reprehensible, even despicable. Within less than a decade, however, it was marketmakers' privileges that had begun to seem improper.

Whitcomb and other financial economists played a role in the transformation. Whitcomb helped set up an Electronic Traders Association to make the case to the SEC and Congress that 'SOES bandits were not doing anything evil or dishonest, and might even be performing a service by putting some pressure on [bid-offer] spreads' (Whitcomb interview 3). The Association's efforts were helped by nationwide news coverage of the discovery by financial economists William Christie and Paul Schultz (1994) of the informal norm among Nasdaq marketmakers to keep spreads wide by avoiding odd-eighths price quotations. The odd-eighths scandal helped pave the way for the SEC's 1997 'Order Handling Rules', which forced Nasdaq's official marketmakers to display customers' bids and offers on-screen when their prices were better than their own. Having an order book that was visible

to some participants and not others (as on the NYSE) began to seem unacceptable. Island mounted a pointed advertising campaign around the slogan, ‘We’ll show you our book. Why won’t they?’ (interviewee C).

None of this, however, would have been effective in undermining fixed-role markets without the concrete demonstration – on Island and other new electronic venues – that high volumes of trading and competitive prices could be achieved without official marketmakers. ATD’s algorithms, those of the other new, unofficial ‘electronic marketmakers’, and HFT more generally, were crucial to that. A decisive shift in US share trading towards all-to-all markets was underway.

The consequences for exchanges have been wrenching: in 2013, for example, the NYSE, with its history stretching back more than two centuries, was bought by an electronic newcomer, the Intercontinental Exchange. Previously central roles disappeared or lost privileges and responsibilities; the economic value of occupying them plummeted. At the end of the 1990s, being an NYSE specialist was still very profitable, with operating margins that could exceed 50 percent (Vinzant, 1999). In 2000, Goldman Sachs paid \$6.5 billion to acquire the NYSE specialists Spear, Leeds & Kellogg. In 2014, Goldman sold its NYSE business to the Dutch HFT firm IMC for less than a hundredth of that: \$30 million (Dolgoplov, 2015: 7).

In ATD’s early years, ‘we were begging to get on [exchange] systems’, and meeting resistance, recalls interviewee B. ‘New York [Stock Exchange]: “oh, don’t even come here.”’ By the early 2000s, exchanges discovered that this attitude to automated trading was no longer tenable. They *needed*

marketmaking algorithms and other forms of HFT to sustain trading volumes in the face of competition from new venues such as Island. '[T]here came a point where they were begging to have us bring our volume to their systems. ... New York: "oh, please come here."'

To attract firms that traded algorithmically, traditional exchanges such as the NYSE had to reorganize trading fundamentally, creating technical systems within which algorithms had the data to predict prices and the capacity to act on those predictions with minimal delay. (This can be seen as a strong performative effect: the use of algorithms helped create markets materially better suited to algorithms.) For example, almost all US share-trading venues had to offer trading firms colocation: the capacity to place their servers in the same building as an exchange's systems. Almost the entire range of interviewee B's 'game pieces' – the data needed to identify pockets of predictive 'structure' – were made available. (The chief exception to the availability of data is firms' identities. In fixed-role markets, which firm has placed which order is often knowable, and that is usually seen as acceptable, even necessary as a way of making participants' economic behaviour visible and thus subject to informal sanctions. In contrast, every financial market of which I am aware that has more of an all-to-all organization is anonymous.)

Because so much of our world (not just financial markets), has been reshaped – sometimes by algorithms – to become materially suited to algorithms, it is now easy unconsciously to naturalize algorithmic prediction: not to see extent to which core practices of prediction (by HFT and other algorithms) are possible because of 'structure' resulting from material

arrangements that are the outcome of political-economic struggles.¹²

Consider, for example, the continuing importance to HFT (amply testified to by the interviewees in my wider dataset) of changes in index-futures prices and order books as predictors of share-price movements. That pocket of predictive structure is in no sense ‘natural’: it depends, ultimately, on the fact that futures and shares fall under separate Federal regulators – the CFTC and SEC – with higher leverage permitted in futures. That arrangement has been challenged repeatedly, but unsuccessfully, most recently in 2012, when Congressional Representatives Barney Frank and Michael Capuano put forward a bill to merge the CFTC and SEC. The Frank-Capuano bill failed, in particular because of opposition from the Senate Agriculture Committee, to which the CFTC reports because of that body’s roots in regulating agricultural futures.

As I have emphasized, ATD’s history throws light on fundamental changes in US share trading. I should, however, also note that this was the only market in which ATD was involved in any large scale. Other HFT firms were and are involved in the trading of other kinds of asset (futures, foreign exchange, benchmark US Treasury bonds, options, interest-rate swaps, and some physical commodities), and in trading in Europe, East Asia, Brazil and Australia as well as in the US. Any adequate treatment of the overall political economy of high-frequency trading thus needs to cast its net far more widely than this article has. Furthermore, even the political economy of US share trading has important facets that have not been examined here in any detail. Regulation, for example, has shaped US share trading profoundly, but in this article I have treated it in the way ATD experienced it, as a predominantly

¹² I am extremely grateful to Martha Poon for comments that led me to this formulation.

external sphere on which the firm had only limited direct impact.¹³ Nor have I examined in any detail the relations between the regulation of financial markets and the political system.

Future publications will tackle such issues. However, it is already clear from my interviews that the political economy of algorithmic prediction developed in this article throws light on the other domains in which HFT is active. There too, pockets of predictive structure exist, some of which are the result of political-economic processes. Who or what has access to the data needed to identify potentially profitable pockets is a political-economy matter, as is the capacity actually to realize those profits (in some domains, that capacity is restricted by the blatant exercise of market power). When access to data and the constraints on and opportunities for trading change, the balance of power among institutional actors is threatened. In some cases, the status quo has so far proved resilient; in others, fixed roles are now beginning to crumble there too.

Acknowledgements

I am hugely grateful to David Whitcomb and my other ATD interviewees, without whom this article could not have been written. The research was funded by the European Research Council (FP7, grant 291733). Comments

¹³ As noted, Whitcomb was involved in efforts to increase the perceived legitimacy of ‘SOES bandit’ trading and to encourage decimalization (pricing in cents). Such influence as he had, however, flowed more from his status as a well-respected academic economist than from his role in ATD.

from Whitcomb, Martha Poon and two anonymous referees helpfully influenced the redrafting of this article, but responsibility for all remaining errors is mine.

References

- Abolafia MY (1996) *Making Markets: Opportunism and Restraint on Wall Street*. Cambridge, MA: Harvard University Press.
- Aspers P (2007) Theory, reality, and performativity in markets. *American Journal of Economics and Sociology* 66(2): 379-398.
- Barnes B (1974) *Scientific Knowledge and Sociological Theory*. London: Routledge.
- Berkowitz SA, Logue, DE and Noser EA (1988) The total cost of transactions on the NYSE. *Journal of Finance* 43(1): 97-112.
- Beunza D and Stark D (2004) Tools of the trade: The socio-technology of arbitrage in a Wall Street trading room. *Industrial and Corporate Change* 13(2): 369-400.
- Biais B, Bisière C and Spatt C (2003) *Imperfect Competition in Financial Markets: ISLAND vs NASDAQ*. Available at: <http://papers.ssrn.com/abstract=302398> (accessed April 27, 2011)
- Borch C, Hansen KB and Lange A-C (2015) Markets, bodies, and rhythms: A rhythmanalysis of financial markets from open-outcry trading to high-frequency trading. *Environment and Planning D: Society and Space* 33(6): 1080-1097.
- Callon M (ed.) (1998) *The Laws of the Market*. Oxford: Blackwell.
- Callon M (2007) What does it mean to say that economics is performative? In: MacKenzie D, Muniesa F and Siu L (eds) *Do Economists Make Markets? On the Performativity of Economics*. Princeton, NJ: Princeton University Press, 311-357.
- Callon M and Latour B (1981) Unscrewing the big Leviathan: How actors macro-structure reality and how sociologists help them to do so. In: Knorr Cetina K and Cicourel AV (eds) *Advances in Social Theory and Methodology: Toward an Integration of Micro- and Macro-Sociologies*. Boston: Routledge and Kegan Paul, 277-303.
- Christie WA and Schultz PH (1994) Why do NASDAQ market makers avoid odd-eighth quotes? *Journal of Finance* 49(5): 1813-40.
- Cohen KJ, Maier S, Schwartz RA and Whitcomb DK (1981) Transaction costs, order placement strategy and the existence of the bid-ask spread. *Journal of Political Economy* 89(2): 287-305.
- Collier J (2002) Lowcountry firm among high-tech's best. *The State*, August 26. Available at

- <http://www.thestate.com/news/business/article/14327990> (accessed August 26, 2015)
- Coughenour J and Harris LE (2003) *Specialist Profits and the Minimum Price Increment*. Available at <http://www.usc.edu/schools/business> (accessed December 23, 2012)
- Dolgoplov S (2015) *Regulating Merchants of Liquidity: Market Making from Crowded Floors to High-Frequency Trading*. Available at <http://tabbforum.com> (accessed January 13, 2016).
- Engelen E, Ertürk I, Froud J, Johal S, Leaver A, Moran M, Nilsson A and Williams K (2011) *After the Great Complacency*. Oxford: Oxford University Press.
- Fisher RA (1936) The use of multiple measurements in taxonomic problems. *Annals of Eugenics* 7(2): 179-88.
- Fligstein N (1996) Markets as politics: A political-cultural approach to market institutions. *American Sociological Review* 61(4): 656-73.
- Fligstein N (2001) *The Architecture of Markets*. Princeton, NJ: Princeton University Press.
- Gillespie T (2014) The relevance of algorithms. In: Gillespie T, Boczkowski PJ and Foot KA (eds) *Media Technologies: Essays on Communication, Materiality, and Society*. Cambridge, MA: MIT Press, 167-193.
- Harris JH and Schultz PH (1998) The trading profits of SOES bandits. *Journal of Financial Economics* 50(1): 39-62.
- Hasbrouck, J (2003) Intraday price formation in U.S. equity index markets. *Journal of Finance* 58(6): 2375-99.
- Instinet (1988) Customer Computer to Instinet Application Layer Specification, Version 2.6. May 6. In interviewee B's private papers.
- Kearney B (2012) Trading firm co-founder James Hawkes pursues math 'courseware'. *Post and Courier*, December 10. Available at <http://www.postandcourier.com> (accessed August 26, 2015).
- Knorr Cetina K and Bruegger U (2002) Global microstructures: The virtual societies of financial markets. *American Journal of Sociology* 107(4): 905-951.
- Lévi-Strauss C (1966) *The Savage Mind*. London: Weidenfeld & Nicolson.
- MacKenzie D (2006) *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.
- MacKenzie D (2009) *Material Markets: How Economic Agents are Constructed*. Oxford: Oxford University Press.
- MacKenzie D, Beunza D, Millo Y, Pardo-Guerra JP (2012) Drilling through the Allegheny Mountains: Liquidity, materiality and high-frequency trading. *Journal of Cultural Economy* 5(3): 279-276.
- MacKenzie D and JP Pardo-Guerra (2014) Insurgent capitalism: Island, bricolage and the re-making of finance. *Economy and Society* 43(2): 153-182.
- McCartney RJ (1990) Trading via computer: It's already here. *Washington Post*, July 1: H5.
- Melamed L and Tamarkin B (1996) *Leo Melamed: Escape to the Futures*. New York: Wiley.
- Millo Y (2007) Making things deliverable: The origins of index-based derivatives. In: Callon M, Millo Y and Muniesa F (eds) *Market Devices*. Oxford: Blackwell, 196-214.

- Muniesa F (2005) Contenir le marché: La transition de la crie à la cotation électronique à la Bourse de Paris. *Sociologie du travail* 47(4): 485-501.
- Muniesa F (2011) Is a stock exchange a computer solution? Explicitness, algorithms and the Arizona Stock Exchange. *International Journal of Actor-Network Theory and Technological Innovation* 3(1): 1-15.
- Philips M (2013) How the robots lost: High-frequency trading's rise and fall. *Bloomberg Businessweek*, June 10-16. Available at <http://www.bloomberg.com/bw/articles/2013-06-06/how-the-robots-lost-high-frequency-tradings-rise-and-fall> (accessed 18 January 2016.).
- Poon M (2009) From New Deal institutions to capital markets: Commercial consumer risk scores and the making of subprime mortgage finance. *Accounting, Organizations and Society* 34(5): 654-674.
- Preda A (2006) Socio-technical agency in financial markets: The case of the stock ticker. *Social Studies of Science* 36(5): 753-82.
- Preda A (2013) Tags, transaction types and communication in online anonymous electronic markets. *Socio-Economic Review* 11(1): 31-56.
- S&P ComStock (1990) Interrogation mode specifications, Version 1.0.0, August. In interviewee B's private papers.
- Schmerken I (2011) Steve Swanson: Renaissance man. Available at <http://old.wallstreetandtech.com/top-innovators/2011/Steve-Swanson> (accessed 18 January 2016).
- Seyfert R (in press) Bugs, predations or manipulations? Incompatible epistemic regimes of high-frequency trading. *Economy and Society*.
- Smith JW, Selway JP and McCormick, DT (1998) *The Nasdaq Stock Market: Historical Background and Current Operation*. N.p.: NASD Department of Economic Research, Working Paper 98-01.
- Swanson SD and Whitcomb DK (2004) Letter to ATD shareholders, March 10. David Whitcomb, private papers.
- Vinzant C (1999) Do we need a stock exchange? *Fortune*, November 22. Available at <http://archive.fortune.com> (accessed January 25, 2016).
- Whitcomb DK (1989a) Letter to prospective ATD shareholders, June 12. David Whitcomb, private papers.
- Whitcomb DK (1989b) Letter to ATD shareholders, November 11. David Whitcomb, private papers.
- Whitcomb DK (1990a) Letter to ATD shareholders, January 30. David Whitcomb, private papers.
- Whitcomb DK (1990b) Letter to [hedge fund manager], April 2. Interviewee B, private papers, 'Portfolio Info' file.
- Whitcomb DK (1990c) Letter to ATD shareholders, October 28. David Whitcomb, private papers.
- Whitcomb DK (1995) Letter to ATD shareholders, January 5. David Whitcomb, private papers.
- Whitcomb DK (2001) Letter to ATD Shareholders. April 8. David Whitcomb, private papers.
- Wipperfurth H (1999) Hyper-aggressive day trading firm wants to take on third market giants. *Investment Dealers' Digest* 65 (42) October 25: 11.

High-frequency traders (AA-BX)	51
Dealers, brokers and broker-dealers (DA-DT)	20
Exchange and trading venue members (EA-GB)	54
Investment-management firm traders (IA-ID)	4
Manual traders (MA-MF)	6
Other algorithmic trading (OA-OM)	13
Regulators (RA-RE)	5
Suppliers of technology and telecommunications links (SA-SR)	18
Researchers/market analysts (UA-UN)	14
Total	185

Table 1. Overall set of interviewees.

Senior management:	6	9%
Administration/compliance/marketing:	17	25%
Technical:	29	43%
Quantitative analysis:	2	3%
Trading:	5	7%
Mixed roles including trading:	9	13%
	<hr/> 68	

Table 2. Staff roles at ATD, early 2000s.

Source: staff list in interviewee B's files, exact date unknown. Mixed roles include, e.g., 'Trading/Research' and 'Trader & Modeler'.