Economy and Society

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/reso20

Machine learning and social action in markets: From first- to second-generation automated trading

Christian Borch & Bo Hee Min

To cite this article: Christian Borch & Bo Hee Min (2023) Machine learning and social action in markets: From first- to second-generation automated trading, Economy and Society, 52:1, 37-61, DOI: 10.1080/03085147.2022.2050088

To link to this article: https://doi.org/10.1080/03085147.2022.2050088

Published online: 04 May 2022.

Economy and Society Volume 52 Number 1 2023: 37–61
https://doi.org/10.1080/03085147.2022.2050088

Machine learning and social action in markets: From first- to second-generation automated trading

Christian Borch and Bo Hee Min

Abstract

Machine learning (ML) models are gaining traction in securities trading because of their ability to recognize and predict patterns. This study examines how ML is transforming automated trading. Drawing on 213 interviews with market participants (including 94 with people working at ML-employing firms) as well as ethnographic observations of a trading firm specializing in ML-based automated trading, we argue that ML-based ('second-generation') automated trading systems are different to previous ('first-generation') automated trading systems. Where first-generation systems are based on human-defined rules, second-generation systems develop their trading rules independently. We further argue that the use of such second-generation systems prompts a rethinking of established concepts in economic sociology. In particular, a Weberian notion of social action in markets is incompatible with such systems, but we

Christian Borch (corresponding author), Copenhagen Business School, Frederiksberg, Denmark. E-mail: cbo.mpp@cbs.dk
Bo Hee Min, Copenhagen Business School, Frederiksberg, Denmark. E-mail: bhm.mpp@cbs.dk

also argue that second-generation automated trading calls for a reconsideration of the notion of the performativity of financial models.

Keywords: automated trading; economic sociology; machine learning; performativity; social action.

Introduction

Like many other societal domains, securities trading has been extensively automated in recent decades. Where decisions to buy or sell securities previously entailed direct human involvement – in the form of, for example, inter-human pit trading (Baker, 1984) or electronic trading with human traders engaging the market on-screen (Knorr Cetina & Bruegger, 2002) – in several contemporary markets, most orders are placed by fully automated algorithms (MacKenzie, 2018a, 2021). As sociological studies of automated trading have aptly demonstrated, these algorithms are typically designed by (teams of) human traders whose instructions are then transformed into codes by software developers (Lange, 2016; Seyfert, 2016), after which the algorithms operate in electronic markets through complex technological infrastructures that permit them to respond to market developments almost at the speed of light (MacKenzie, 2021).

In this paper, we argue that the automation of securities trading is currently entering a new phase and that this has substantial implications. Since the mid-2010s, new machine learning (ML) techniques are increasingly being deployed for trading purposes. ML encompasses a variety of algorithmic techniques tasked with learning to extract patterns in data and making predictions on that basis. We demonstrate that one of the principal motivations for deploying ML for the purpose of securities trading is to delegate the identification of trading strategies to machines, instead of human traders developing the strategies and then software developers implementing them in code. This reconfiguration of automation does not make humans superfluous. A lot of human labour is needed to clean the data on which the ML systems are trained, just as humans obviously design the overall algorithmic architectures. Still, we argue, it is important to appreciate that, once designed, ML-based trading systems are tasked with producing trading strategies independently.

Our aims in this paper are threefold. First, we wish to provide an empirically grounded understanding of how ML-based trading systems might not only inform human action but take independent action in markets. We suggest that the specificities of ML systems can be captured through a comparison with technological systems that do not possess the same autonomous decision-making capacity. To this end, our analysis is based on comprehensive empirical fieldwork in the automated trading industry, including interviews with informants who specialize in human-defined automated trading (what we refer to as 'first-generation automated trading') and/or in various forms of ML-based automated trading (what we call 'second-generation automated trading'). Second, given that ML encompasses a wide array of data-driven techniques, we aim to provide some nuance to sociological ML debates by teasing out differences and similarities between two types of ML approaches we encountered in our fieldwork: genetic programming and deep neural networks. Finally, our third aim is to think through some of the theoretical implications ML-based trading has for economic sociology. We will do so by considering the notions of social action and performativity.

Note that there are several dimensions of second-generation automated trading that we cannot address in this paper. These include the training of ML-based systems; the types of patterns they can (and cannot) extract from the data; the human-machine interaction involved as well as broader questions about what this form of trading might mean for accountability, regulation, the politics of markets, and so on. These are all interesting and important themes, but we have chosen to leave them aside here to focus our attention on the key transformation ML entails in terms of ML systems themselves coming up with trading rules.

This study contributes to two strands of sociological research. The first concerns sociological discussions of ML. These discussions include a growing body of literature that details how predictive ML algorithms are increasingly deployed in various fields such as credit scoring (Fourcade & Healy, 2017; Kiviat, 2019), criminal justice (Brayne & Christin, 2021), insurance (Cevolini & Esposito, 2020), self-driving cars (Stilgoe, 2018), social media (Fourcade & Johns, 2020), and warfare (Scharre, 2018). It also includes studies that discuss the broader consequences of ML systems, including their ethical and political implications and the ways they reconfigure human-machine relationships (e.g. Amoore, 2019; 2020; Burrell, 2016; Christin, 2020; Coeckelbergh, 2020; Svetlova, 2022). We extend these studies by focusing systematically on the inner workings of ML-based trading systems and how these systems challenge established sociological notions. Our discussion of different types of ML architectures further serves to clarify that ML is not a monolithic phenomenon but rather is characterized by an internal variety that sociologists should take seriously.

Second, by focusing on ML-based trading systems, this study advances economic sociology literature on financial markets, specifically focusing on how these systems are different from the kind of automation that has been examined in previous economic sociology. Of primary importance here are MacKenzie's studies of high-frequency trading and the material dimensions of automated markets (e.g. 2018a, 2018b, 2021, 2022). However, there is now more extensive related literature on the consequences of automation for securities exchanges (Pardo-Guerra, 2019) and regulation (Coombs, 2016; Lenglet, 2011) as well as on the emotional, epistemic and organizational dimensions of automated trading (Borch & Lange, 2017; Lange, 2016; Lange *et al.*, 2019; Min & Borch, 2022; Seyfert, 2016; Souleles, 2019). While none of this research has systematically studied the uptake of ML-based trading systems, recent analyses have taken steps in that direction. This includes Hansen's (2020, 2021) work on the ways in which many brokers and traders who use ML algorithms often prefer simple models since these are more intuitive (see also Hansen & Borch, 2021). Our analysis differs from Hansen's in that we focus on the deliberate deployment of highly complex ML systems, just as we discuss the implications of ML for economic sociology (see similarly Borch, 2021, 2022).

The paper is structured as follows. We begin by presenting our methods and data. We then compare first- and second-generation automated trading systems to single out what is specific to the ML domain. The following section discusses the conceptional implications of the shift to second-generation automated trading by focusing on the notions of Weberian social action in markets and, more briefly, the performativity of financial models. The concluding discussion summarizes and makes suggestions for future research.

Methods and data

To examine and compare first- and second-generation automated trading systems, we drew on two sets of data: (1) a qualitative study of financial market participants across institutions and markets and (2) observational fieldwork at Tyler Capital Limited, a London firm specializing in deep neural network-based automated trading. The first consisted of 213 interviews with informants at 146 institutions in the United States and Europe (including interviews conducted at Tyler Capital). This pool of interviews included informants working at trading firms, banks, investment management firms such as hedge funds and asset managers, and brokers who deploy automated trading or trade execution systems. Further, it included informants from institutional investors, exchanges, regulators and technology providers. The informants' occupations included traders, quantitative analysts, software developers, researchers and executives.

Interviews focused on informants' backgrounds, the type of work they did, their algorithmic systems, the challenges they faced, and overall developments in automated financial markets. Interviews typically lasted one hour, although some were significantly longer. All informants were offered and granted anonymity, and to avoid identification, we refer to firm names using pseudonyms – except for Tyler Capital where management decided to waive anonymity.

The interviews in the larger qualitative study were conducted in two phases. Between 2014 and 2016, the first author conducted 23 interviews primarily with informants in the Unites States (11 of these interviews were conducted jointly with Ann-Christina Lange). From 2017 to 2020, a research team comprising the authors, Kristian Bondo Hansen, Pankaj Kumar, Nicholas Skar-Gislinge and Daniel Souleles conducted an additional 190 interviews – 78 of which the authors conducted either independently or collaboratively. As Table 1 shows, the interviews particularly focused on trading firms, most of which deployed fully automated trading systems (47 firms and 84 interviews).

Table 1 Summary of qualitative study by algorithmic practice

Institutional type	No. of firms (no. of interviews)					Total
	First-generation automated trading	Exploratory ML	Complementary ML	Second-generation automated trading (ML) only	Non-algorithmic	
Trading firms	20 (27)	5 (9)	11 (13)	5 (29)	6 (6)	47 (84)
Banks	6 (9)		9 (14)			15 (23)
Brokerage firms	2 (4)		2 (5)			4 (9)
Investment management firms	12 (12)	1 (1)	6 (7)		9 (9)	28 (29)
Exchanges and trading venues	9 (21)	1 (1)				10 (22)
Data, technology and analytics providers	5 (7)		9 (9)	5 (5)	4 (4)	23 (25)
Regulators	1 (1)		1 (1)		5 (7)	7 (9)
Other	3 (3)		1 (1)		8 (8)	12 (12)
Total	58 (84)	7 (11)	39 (50)	10 (34)	32 (34)	146 (213)

Informants from automated trading firms were mostly located in Chicago, New York, London and Amsterdam, all of which are hubs for this form of trading. This pool of interviews allowed us to examine the practices and characteristics of automated trading and compare these as they materialized across first- and second-generation automated trading.

That said, interview data of this sort entail important caveats. The first caveat concerns the proprietary nature of the trading systems, meaning that informants were reluctant to reveal much about the specifics of their automated strategies (an issue other researchers studying automated trading are also facing, see for example MacKenzie, 2021). Relatedly, it was not always clear from the interview data whether, in individual firms, ML was deeply implemented for the purpose of automated trading or if it merely served as a glossy term with which to present the firm externally as being up-to-speed with computer science developments.

Considering this, we coded our interview data to best reflect the informants' use of algorithms in trading-related activities. As shown in Table 1, we identified four types of institution-level algorithmic practices in the qualitative study: (1) firms that only deployed first-generation automated trading systems; (2) institutions that explored the potential benefits and applications of ML while using first-generation algorithms for trading and investment-related decisions; (3) institutions that adopted ML to complement and improve the performance of their existing first-generation algorithms; and (4) firms which used sophisticated ML-based algorithms exclusively, either for trading purposes or to identify tradeable information that was then resold to investors and trading firms.[1] As Table 1 shows, our interview pool contained a total of 95 interviews with informants at ML-employing institutions.

Given that we were particularly interested in second-generation automated trading systems in this study, we focused our attention on firms that only did ML-based trading. Specifically, we mainly discuss three firms (Tyler Capital, Vermont Trading and Ragin Capital) that all specialize exclusively in second-generation automated trading. Our interview data made clear that ML was not used as a mere glossy term by these firms, but rather constituted the core of their automated trading systems as well as a source of numerous practical challenges. The authors conducted interviews in all three firms, and our analysis is based on repeat interviews in each. Two of these firms (US-based hedge funds Vermont Trading and Ragin Capital, with approximately 100 and 14 staff, respectively) deploy the same overall type of ML architecture, genetic programming. In contrast, the third firm, Tyler Capital, is a proprietary trading firm (which means that it trades on its own account rather than on behalf of clients) that uses a deep neural network architecture. The final two ML-only firms in the larger interview pool similarly use deep neural networks.

We combined the larger pool of interview data with a second set of data: ethnographic fieldwork at Tyler Capital. This firm focuses on high-frequency trading in futures contracts on the Chicago Mercantile Exchange, although it also trades on other markets worldwide. It has a staff of around 50, making it a medium-sized firm in an industry characterized by huge variety. Tyler Capital's trading is done by a fully automated ML system that the firm has dubbed OPUS. Briefly described, OPUS is a sophisticated deep neural network that, in contrast to first-generation automated trading systems with their human-defined rules, *learns and develops trading rules itself*, based on the large amounts of market data it is fed. Once developed and tested, the trading rules are implemented through now customary high-speed infrastructures that connect the firm to the exchanges on which it trades (MacKenzie, 2018b), including long-distance data transmission infrastructures and 'co-location' services where the firm has computer servers placed physically in exchange data centres to receive exchange data feeds and respond to market developments quickly.

We collected data on Tyler Capital in three phases: (1) interviews by the first author of this paper with a ML engineer and the Chief Executive Officer (CEO) and Chief Technology Officer (CTO) of Tyler Capital in 2017; (2) ethnographic fieldwork by both authors in 2018–2019; and (3) remote interviews in 2020 (due to COVID-19 lockdowns and restrictions). We collected three types of data, as shown in Table 2: ethnographic observations, interviews and documents.

Our ethnographic field visits each had a duration of a couple of days, where the firm provided us with access to the entire site for participant and non-participant observation. During our visits, we followed the work of organizational members with different ranks, roles and team affiliations, including sitting with staff at their desks and watching demonstrations of their technology. Our ethnographic fieldwork was centred on activities that directly concerned OPUS. We also attended several events in the firm, including management meetings, information-sharing meetings, and informal gatherings.

In addition to the observational data, we conducted 23 open-ended semi-structured interviews across the firm, interviewing 18 people from all teams and hierarchical levels: the founder, executives, team leaders and team members (see Table 3). This allowed us to cross-check observations from our participant and non-participant work. Since we interviewed members of Tyler Capital according to their roles – such as leadership, ML engineering, trading, operations and infrastructure – we often conducted group interviews with people engaging in similar jobs. In pursuit of including all roles at the firm, we interviewed every executive and team leader. In addition, we conducted repeat interviews with key informants, such as the CEO and CTO. The CTO leads all teams that are at the centre of ML practices in the firm – the Research Engineering Team, Trading Team and Production Team – and he and the CEO are responsible for instituting practices that involve those teams and the ML system. To comprehensively explore these topics, we spent most of a day off-site with the CTO. Finally, management gave us access to internal documents (see Table 2) that describe a wide array of the firm's technological systems and procedures. These documents further complemented our ethnographic data, as they outlined the design and behaviours

Table 2 Data sources for Tyler Capital

Data source	Description	
Observations	2 Teams	• Trading Team • Production Team (in charge of trading operations)
	5 Meetings	• Technology and Research group meeting • Leadership Group meetings (two meetings) • Production Control Group meeting Employer-organized end-of-week social event
In-depth interviews	23 interviews (429 pages of transcripts)	18 members of the firm
Internal documents	32 documents (324 pages)	• Trading Platform Framework • Applied Risk Management in an ML Era • OPUS Principles for an AI Era • Screenshots of Internal Systems • Operational Updates • Governance and Operating Framework • Technology & Research Charter • An Industry in Transition • Others

of the firm's technology that we could not directly observe and served as a further cross-checking of insights from fieldwork and interviews.

Table 3 suggests that our interviews concentrated on top management. We, therefore, cross-checked management's accounts in interviews with the observations of staff members to ensure that findings from our data had wide resonance in the firm. We only report findings that did indeed have wide resonance. Also note that it is not uncommon for executives to be actively involved in the design of trading systems. In the firms discussed in this paper, the executives cited all played a central role in designing their firm's algorithmic systems.

We do not claim that our fieldwork at Tyler Capital is necessarily generalizable in the sense that other firms specializing in second-generation automated trading (more specifically, in deep neural networks) are similarly organized or have parallel emphases. However, for the purposes of this study, we mainly focused on what second-generation automated trading implies in terms of

Table 3 Summary of interviews at Tyler Capital[a]

Functional Unit	Role	No. of informants	No. of interviews
Board of Directors	Founder	1	1
Leadership[b]	Chief Executive Officer	1	10
	Chief Technology Officer	1	11
	Chief Financial Officer	1	2
	Chief Risk Officer	2[c]	3
Technology and Infrastructure	Team Leader	1	4
	Machine Learning Engineer	1	1
	Production Engineer	3	2
	Delivery Engineer	2	2
	Compliance Engineer	1	1
	Infrastructure Engineer	1	1
Trading	Team Leader	1	1
	Senior Trader	1	1
Corporate Support	Human Resources Manager	1	1
Total		18	23[d]

[a]This table only includes formal, recorded interviews. Unrecorded interviews and informal conversations are excluded.
[b]Includes all members of the corporate leadership.
[c]There was a CFO transition during the data collection period, and we interviewed both the new and outgoing CFO.
[d]Because approximately half of the interviews were group or team interviews, the total number of informants and the sum of the numbers of interviews with each role differ from the total number of interviews.

producing automated trading rules, and here we saw full consistency across informants working in competing firms, that is, regardless of the specific ML architecture chosen.

First- and second-generation automated trading

To demonstrate what is new and specific to second-generation automated trading systems and what implications they hold for a sociological understanding of financial markets, it is helpful to compare them to their first-generation forerunners. First-generation automated trading dates to the 1990s but really accelerated in the early 2000s. It is often associated with high-frequency trading – which in some US markets, especially equities and futures markets, now accounts for approximately half the overall trading volume (MacKenzie, 2018b, p. 1639) – even though it is only a subset of automated trading. For example, brokers executing orders on behalf of clients increasingly use automation when they slice large orders into smaller ones and execute them automatically according to particular temporal patterns or by considering the overall volume in markets.

First-generation automated trading consists of human-defined, rules-based strategies that are transformed into computer code and then executed through sophisticated technological infrastructures that connect firms to markets (Lange, 2016; Seyfert, 2016). More precisely, these algorithms have two types of input, *rules* and *data*, which together produce an output in the form of particular orders to buy or sell securities (for example, 'if *C* conditions are in place, then place an order to buy *A* amount of *S* shares at *P* price at trading venue *V*'). Data typically consist of market data in the form of cumulative market events (historical data) and real-time information on current events (continuous flow). The most important source of market data is derived from 'electronic order books' (MacKenzie, 2018b; Pardo-Guerra, 2019). An electronic order book is an exchange's comprehensive record of any outstanding orders, and it provides a snapshot of the state of the market at any given time. While other forms of data are available, for many automated trading firms, order book data remain the primary, if not the sole, source of information about the market.

Moving on to second-generation automated trading, like other domains, there has been a surge of ML applications within automated trading since the mid-2010s. The overall idea of ML-based approaches is to give the algorithmic system an objective function (some version of 'maximize profits under given risk parameters') and then train it on extensive data to learn it to make actionable predictions. Specifically, as mentioned earlier, the chief objective of using ML for trading purposes is to create automated trading rules, that is, a fully computer-generated trading strategy that, for example, would 'decide at time *t* which stocks to buy, hold or sell' (Alonso *et al.*, 2019, p. 266). This ambition encapsulates the central difference between first- and second-generation automated trading (see Figure 1). In contrast to the former's end-to-end human-defined rules, the task for the ML algorithm is to come up with a set of trading rules itself, based on the comprehensive set of market data (input) on which it is trained. So, it is not merely a matter of a human-defined strategy being automatically implemented in the market by an algorithm. Rather, it is a matter of *the strategy itself* being invented by the ML system and then automatically implemented in the market in the form of specific orders to buy or sell securities (output).[2]

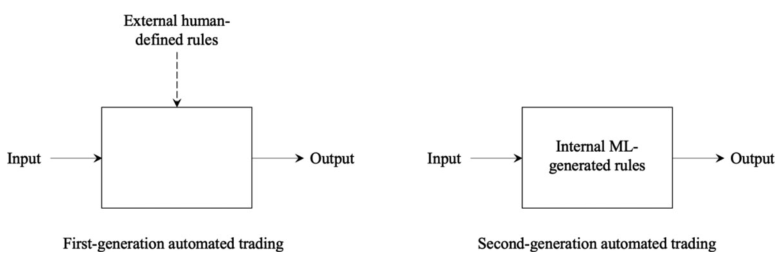

Figure 1 First- and second-generation automated trading systems

Without over-stretching the analogy, first-generation automated trading systems may be understood by paraphrasing Durkheim (1982): These systems face external rules that impose themselves upon the algorithms. By contrast, second-generation automated trading systems are characterized by what we might call an inversion of Durkheim where trading rules are entirely internally generated and, though humans may seek to influence these rules (and wrap them in all sorts of control mechanisms so they do not act disorderly), the ML systems are in effect endowed with proper independence and decision-making autonomy. In these forms of ML-based trading, it is the onus of the ML system to find out what types of trading behaviour might be profitable, and under what conditions, and then to make the actual decisions to buy or sell securities.

This inverse Durkheimianism might be illustrated through a discussion of some of the different kinds of ML-based trading we encountered in our fieldwork. Here, we discuss two different approaches that share the idea that it is the automated system rather than humans that develop the trading rules and hence is the originator of the trading strategy being deployed.

The first approach is known as 'genetic programming'. Taking inspiration from evolutionary selection, the more specific idea of genetic programming is to produce a group of candidate solutions to a problem through evolutionary iteration. The process has the following main steps: (1) define a population of individual candidate 'agents' that seek to fulfil a particular purpose in some environment (for example, agents may consist of trading rules with certain parameters that seek to make money); (2) assign to each a fitness function describing how well the purpose is fulfilled; (3) run simulations of how well the agents perform and then select the best candidates (say, those making the most money under certain risk conditions); (4) generate a new population on this basis by (a) taking some of the selected candidates unmodified into the new population ('reproduction'); (b) creating new candidates in the new population by randomly recombining elements from some of the selected individuals ('crossover'); and (c) creating new candidates in the new population by randomly substituting elements from some of the selected candidates with new randomly generated ones ('mutation'); and (5) repeat steps (3) to (4) until a set of optimized candidates has evolved (Koza & Poli, 2014). It follows that, in this approach, learning is an evolutionary accomplishment cultivated through generations of refinement.

Some of the firms we interviewed deploy genetic programming as their main ML architecture. According to the CEO of one of these firms, Vermont Trading, they use genetic programming 'to evolve rule sets that do the trading'. In other words, their ML architecture is designed such that it produces a set of traders, or trading agents, whose strategies evolve evolutionarily through a process similar to the one described above: 'These traders are evolved over hundreds of thousands of compute nodes over months, and we continuously harvest from these runs [...] right now it's every three hours or so' (CEO, Vermont Trading). The ML architecture is designed so that the trading agents (or rules) which are dynamically evolved in this fashion are able to engage in complex behaviours in markets:

> Each rule triggers an action, and again that action is evolved. The actions that we have, for example, are not just 'go long' or 'go short but cover your short' and so forth, but it's also the order models; for example, 'go short, but go short with a limit order at this particular price normalized by volatility with respect to the prior close'. (CEO, Vermont Trading)

Importantly, given the thousands of iterations of reproduction, crossover and mutation, the rules and actions that are evolved cannot be anticipated. So, though humans have designed the overall ML architecture, the result is not a human-conceived one. In fact, the entire aim of this approach is the exact opposite of a human-designed strategy: to automate the generation of a set of trading agents that will display strategies superior to what humans could conceive. To be sure, at Vermont Trading, they have a group of quantitative analysts ('quants') who define a set of indicators, including 'technical indicators that are simply aggregations of time series of price, volume, volatility, and so forth' (CEO, Vermont Trading). These indicators are a way for human experts to guide the ML system in particular ways, say, by seeking to endow some of the agents with particular volatility preferences. However, this kind of external human influence is at best partial, indirect and non-deterministic. 'The system being evolutionary, if the indicator has value, it will be picked up. If it doesn't, it won't' (CEO, Vermont Trading).

The CTO of another firm, Ragin Capital, whose work also revolves around genetic programming, similarly expressed that they have quants who come up with ideas for trading opportunities but, in the end, it would be the ML system itself that 'is constantly generating new strategies and new signals [i.e. trading opportunities] that feed into the overall strategy'. When asked if it is indeed the ML system that detects these signals rather than, say, human employees, the CTO responded, 'That's right'. The point is that human input will only matter for the ML system if the evolutionary iterations and their unanticipated upshots can confirm its importance. Except for these types of indirect, non-deterministic influence, the ML system is completely automated.

Another key ML architecture deployed by several trading firms is that of artificial neural networks (ANNs), which include deep neural networks (DNNs). The basic idea of ANNs is to create a network of information processing units (neurons) in which learning is achieved iteratively by the network itself. This idea was originally (but much less so today) inspired by the notion that the brain operates as a neurological network. The neurons are connected via so-called synapses, which transfer information through the network. Synapses also connect neurons to input and output, with each synapse being given a variable weight. An example of a simple, single-layer ANN is depicted in Figure 2.

In contrast to such 'shallow', single-layer ANNs, 'deep' neural networks contain two or more 'hidden' layers. An example of a DNN is presented in Figure 3. This DNN has five sources of input data, and it makes one output prediction, with two hidden layers placed between the input and output

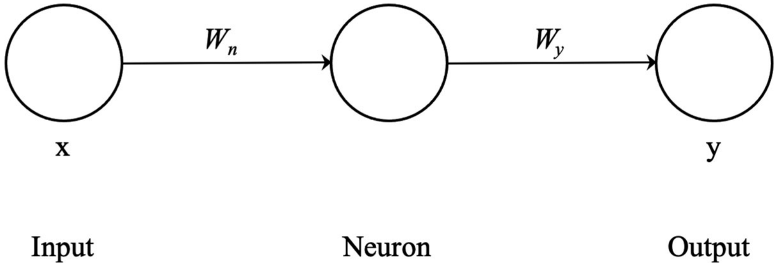

Figure 2 Simple neural network

layers. Note that for simplicity, Figures 2 and 3 only have one neuron in the output layer (they make just one prediction), but depending on the task at hand, there could be multiple outputs.

As mentioned earlier, Tyler Capital's automated trading system is centred on a DNN architecture called OPUS. Feeding the neural network huge amounts of historical market data, the idea is that OPUS will develop a set of automated trading rules, that is, come up with and execute independent trading strategies on its own. Like Vermont Trading and Ragin Capital, the firm has staff employed who may suggest possible trades and feed such suggestions into the system. However, in the end, the ML system has full discretionary power to disregard these suggestions. As the CEO of Tyler Capital put it, 'we don't write rules for [OPUS]'. The firm's CTO added that:

> OPUS has the entire responsibility to take information in aggregate over years and come up with an action policy. So basically, all the humans can ever do is

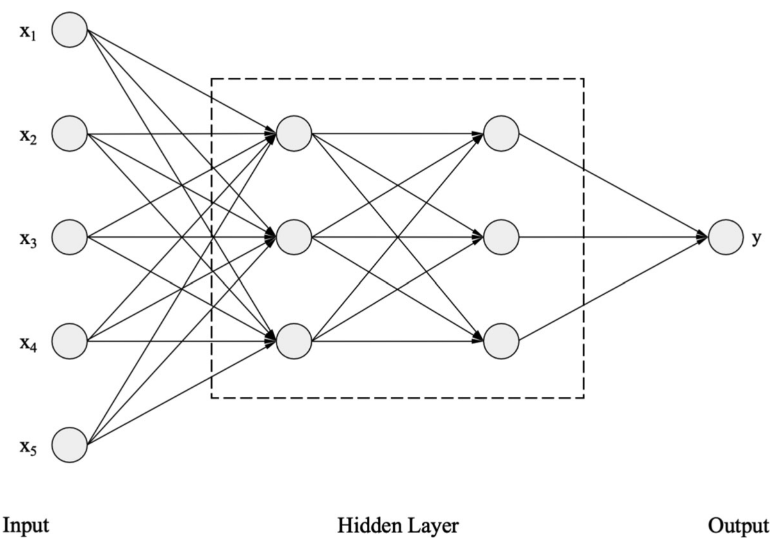

Figure 3 Deep neural network

> give OPUS information, incentives and penalties, and it's up to OPUS entirely to come up with the trading policy. We can't do it. We can inform, we can enforce risk controls, like 'you can't put on more than 100k of margin', right?[3] Of course, we have controls that are of that form, but we don't have rules where a human comes up with a rule. We try to capture their intuition as information representation. A wiggly line that OPUS gets to see, and then OPUS decides if that wiggly line is meaningful or not. So OPUS, on some level, gets to choose from the options that the people give it. That's part of it. Also, we can give it sometimes quite broad options so that it can choose on its own, and it can just ignore what it doesn't want. (CTO, Tyler Capital)

It is common to first- and second-generation automated trading systems that once the algorithms are active in markets, they operate without direct human involvement. One of the things the people designing first-generation automated trading systems need to learn is to abstain from the temptation to suddenly interfere with the system's operations on the basis of an impulse idea (Borch & Lange, 2017). The same applies to second-generation systems. According to the CEO of Vermont Trading:

> People come in [in the morning] and they 'set up' the desks. [...] There are certain set up steps that they take [...]. They're monitoring; they have CNBC on and a Bloomberg terminal and all that, but [in terms of the ML system] it's hands-off. I like to joke, 'look but don't touch'. (CEO, Vermont Trading)

It is also important to note that, while firms would typically have various types of control wrappers around their systems, including the ability to pull all orders if necessary, the trading rules generated by the ML systems would usually be directly implemented in automated market action. In other words, the ML trading rules do not merely inform human action in markets, they replace it entirely. We now discuss some of the conceptual implications of this.

Social action and performativity in an age of second-generation automated trading

Weber (1978, pp. 635–636) famously argues that markets are arenas for social action in a twofold sense, as buyers and sellers not only take account of each other (exchange) but also of rival buyers and sellers (competition), and that such interaction constitutes 'the archetype of all rational social action' (Swedberg, 2000, p. 379). Of course, Weber's analyses of financial markets revolve around the face-to-face interaction of human traders that characterized the markets during his time. Though present-day automated markets look markedly different, certain features identified by Weber continue to be of relevance. For example, Pardo-Guerra (2019, pp. 304–308) notes that Weber's (2000, pp. 310, 323) emphasis on how stock and commodity exchanges offer a physical location for trading to take place remains important, although nowadays the situatedness of trading revolves around particular material infrastructures that tie market participants physically to exchange data centres such as through co-location services.

But what about social action in algorithmically driven markets? Is this a contradiction in terms given that fully automated algorithms are non-human entities? According to Weber's famous definition:

> We shall speak of 'action' insofar as the acting individual attaches a subjective meaning to his [*sic*] behavior – be it overt or covert, omission or acquiescence. Action is 'social' insofar as its subjective meaning takes account of the behavior of others and is thereby oriented in its course. (1978, p. 4)

Using this definition as our starting point, we discuss a series of key elements of Weber's theory of social action – alter-orientation, subjective meaning, means-end relation, and 'explanatory understanding' (focusing on the 'motive' or 'intended meaning' behind an action; 1978, pp. 8–9) – considering both first- and second-generation automated trading systems. Beginning with the 'social' or alter-orientation part of Weber's definition, this is an element that resonates with how many first-generation automated trading systems operate. As MacKenzie (2018a) has argued – and this is supported by our interview data – due to their order book emphasis, first-generation automated trading systems are often designed such that they consider how other market participants are acting in the market. When other market participants place orders in the market, this is reflected in changes in the order book, and even seemingly minor changes might reflect larger movements that could have a considerable impact on markets. Therefore, several types of strategies have been developed that aim at detecting and predicting whether, for example, a series of smaller orders may reflect a larger underlying order.

Where the alter-orientation of first-generation automated trading systems matches Weber's definition of social action, the same cannot be said about his emphasis on subjective meaning. After all, automated trading systems are not human beings who are able to attach subjective meaning to their behaviour. On closer inspection, however, first-generation systems can be seen as devices that enact human instrumentally rational action. This is the interpretation proposed by Gane (2012) in his analysis of first-generation automated trading. Gane cites Weber's comment that 'every artifact, such as for example a machine, can be understood only in terms of the meaning which its production and use have had or were intended to have' (Weber, 1978, p. 7). Given that their rules are human-defined, first-generation automated trading systems may be seen as constituting tools humans use as 'an instrumentally rational design that seeks to secure an advantage in an ongoing struggle over price' in markets (Gane, 2012, p. 70). Or, as put by Weber himself:

> That which is intelligible or understandable about [the machine] is thus its relation to human action in the role either of means or of end; a relation of

which the [human] actor or actors can be said to have been aware and to which their action has been oriented. (1978, p. 7)

Combining this machine-as-human-means-or-end with the fact that first-generation automated trading systems are designed to take account of the behaviour of other market participants, makes it reasonable to conceive of this type of automated trading systems as engaging in social action at a distance, that is, human social action as mediated and enacted by machines.

However, what complicates this Weberian interpretation of automated surrogate instrumentally rational action is that these first-generation automated trading systems are typically designed in an evolving fashion that tends to render them difficult to comprehend and potentially makes it difficult to trace their action back to human subjective meaning (see also Coombs, 2016). As one of our informants put it, first-generation automated trading systems have a 'bolt on' nature, in which new rules are gradually added on top of existing rules. This budding process is intensified by the fact that due to fierce competition, the period in which first-generation automated strategies are profitable (before exploitable opportunities are seized by others or otherwise disappear) is constantly diminishing. According to our informants, this means that trading strategies may currently only be viable for a couple of months. To deal with this 'performance degradation', as some of our informants phrased it, traders spend a lot of their time developing new strategies or tweaking old ones to postpone their decay (Borch & Lange, 2017). As a result, the overall trading system – the so-called 'black box', which collates all the individual strategies – may grow increasingly complex, and its operations and inner connections may consequently grow more difficult to comprehend. Indeed, as Lange (2016) observed based on her ethnographic fieldwork on first-generation automated trading, high-frequency traders are often unable to account for and explicate how their overall trading system works, concluding that 'the traders are unlikely to know much about the internal operations of the black box' (2016, p. 237). This was echoed by the CTO of Tyler Capital. Comparing the firm's ML-based system to first-generation automated trading systems, with which several Tyler Capital staff members have extensive prior experience, the CTO argued that these systems, which may consist of 'millions of lines of code', tend to 'become quite brittle because basically what happens is you keep inventing new rules that get added on new rules that get added on new rules and eventually you go down some path that is intentional but not intentional really'.

So, *in principle*, first-generation automated trading systems can be seen as algorithms that enact human instrumentally rational action and they do so in a setting where they are designed to take account of the behaviour of other market participants. An explanation of the ensuing surrogate social action of a first-generation automated trading system therefore appears attainable through an inquiry into 'the meaning which its production and use have had or were intended to have' for the people who designed it (Weber, 1978,

p. 7). However, *in practice*, budding black-box systems – the continuous addition of new layers of rules, some of which may not be compatible with existing rules – means that despite their seemingly simple rules-based nature ('if this happens, do that'), first-generation automated trading systems can be exceedingly difficult to explain, and their eventual action need not directly reflect human subjective meaning.

Turning now to second-generation automated trading systems, and considering them in light of Weber's theory of social action, in some ways, these systems resemble their first-generation forerunners. This includes their alter-orientation in the markets. Like a first-generation automated trading system, a second-generation system might detect that other market participants' orders to sell a security are suddenly amassing in the order book, relative to the orders to buy it, and use this as an exploitable signal about an immanent price drop. The central difference is that for a second-generation automated trading system such strategies would be learned through data and potentially through ongoing market interaction rather than being encoded by hand.

Similarly, where both the means and ends of first-generation automated trading systems are entirely human defined, only the overall ends (the objective function) are fully human defined in second-generation systems. In contrast, the means – the trading rules derived to fulfil the objective function – are algorithmically generated. Of course, to repeat, there are humans designing the specific ML architectures and as demonstrated earlier, the design choices as well as the data fed into the ML systems will affect how they come up with their trading decisions. However, precisely how the ML objective function is fulfilled is not decided upon by humans. Accordingly, the trading system's underlying 'motivation' is entirely different across first- and second-generation systems. In contrast to first-generation systems, it is conceivable that a second-generation system develops means (trading rules) that conflict with the intentions or values of the people who designed it. For example, potentially, an ML trading system might learn to fulfil its objective function by engaging in illegal trading behaviours, such as spoofing (MacKenzie, 2022), for which reason second-generation automated trading systems demand more comprehensive control wrappers to avoid this than do first-generation systems.

When it comes to Weber's emphasis on explanatory understanding, the difference between different kinds of second-generation automated trading systems is important. For example, the CEO of Vermont Trading stated that, since their trading system is based on genetic programming (rather than neural networks), it is in fact 'inherently explainable. […] In evolutionary computation and genetic programming, we can make the substrate explainable quite easily and still have non-linear relationships and create more sophisticated, complex behaviour out of our machine learning system'. In line with this, the firm has designed their software such that when the ML system is actively trading, it is possible to 'pop up the rules that fire, that were responsible for that trade'. In other words, for each trade made, it is possible to single out exactly which rule(s) were triggered. The firm would therefore know precisely why the system acted as it did. According to the CEO of Vermont Trading, the situation is different for DNNs, where explainability is much harder to attain.

This view might be seen as a self-serving comment in the sense that, as a CEO, one would prefer to be able to say that one's ML systems are explainable. However, several of our informants, including people specializing in DNNs, confirmed that ML systems based on genetic programming are explainable, whereas DNN explainability is exceedingly hard to obtain. For example, the CEO of a smaller firm that transitioned from having a combination of first- and second-generation automated trading to implementing a pure DNN-based trading system said that whereas in genetic programming 'you actually have a program where you can go in and analyse what the program does', things are vastly more complex with DNNs: 'a deep learning model can contain [...] maybe hundreds of millions of variables. [...] It's very difficult to understand what's actually going on'. It is simply not clear how input moves between neurons and how a particular output is eventually produced (Christin, 2020). This was indeed a common argument voiced by several of our informants against using DNNs in the first place. While the predictive capabilities of DNN automated trading systems are generally considered highly impressive, many were reluctant to use DNNs because of the opacity characterizing the ways they arrive at their decisions (similarly Hansen, 2020, 2021).

The opaque decision-making logic of complex DNNs is just one form of algorithmic opacity. Burrell (2016) helpfully delineated two additional forms: 'opacity as intentional corporate or state secrecy', which includes the protection of the proprietary nature of such systems (see similarly Lange, 2016), and opacity as 'technical illiteracy' (that coding is a specialist skill not possessed by all citizens). However, particularly the decision-making opacity has potentially huge consequences. If the decision-making logic evades human understanding, the likelihood decreases that adequate safety mechanisms can be set up to prevent these algorithms from either triggering or exacerbating market crashes (World Economic Forum, 2018). In one of Tyler Capital's internal documents, the challenge is described as follows:

> The difficulty of tracing how decisions have been made by ML applications often makes it very difficult to prevent in advance, or to correct afterwards, undesirable model outcomes. For example, the neural net may discover complex, non-linear 'hidden' correlations that are difficult or impossible to anticipate or discover. Further, it is very difficult to predict how a model trained on known historical data but 'making its own decisions' will react when it is live in the market with a much larger dataset and it encounters events that have not been seen before in the data that was used to train it. (Internal document, Tyler Capital, 2020)

On a more conceptual level, the explainability challenge means that, from a Weberian perspective, second-generation automated trading systems have a wobbly sociological status: Like their first-generation forerunners, the action of order-book focused second-generation automated trading systems would surely take account of the behaviour of other market participants and be thereby oriented in their course. While this might look like social action, the fact that they independently generate trading rules means that the action of these systems cannot be traced back to human subjective meaning, rendering a Weberian conception of social action difficult to apply. Weber's work suggests that the lack of subjective meaning and the inability to explain the action of DNNs would tend to relegate these systems instead to what he called 'merely reactive behavior to which no subjective meaning is attached' (Weber, 1978, p. 4).

Does this suggest that economic sociologists should not associate second-generation automated trading systems with social action? Or does it suggest that Weber's categories need updating and that his emphasis on subjective meaning – more generally, the *verstehen* dimension so central to his work – is overly restrictive and not sufficiently attentive to modern ML technologies? We certainly find his notion overly restrictive and unavailing for understanding ML technologies and see its inadequacy in this domain as a further encouragement to explore theories that either bid farewell to classical human-centred action theory (e.g. Luhmann, 1995) or completely rethink how to conceive of social action (e.g. along the lines of Latour, 2005).

That said, are there ways to rescue a Weberian notion of social action when considering second-generation automated trading systems? It might be argued that, to the extent that humans could eventually arrive at some level of explanation of such systems, including DNN-based ones, these could be seen as engaging in Weberian social action, although this action may not be directly derived from human subjective meaning (and may not be a direct enactment of human instrumentally rational action). Computer science research on so-called 'explainable AI', which tries to address DNN opacity by seeking ways to attain some degree of explainability or interpretability (e.g. Samek *et al.*, 2019), might be said to move in this direction. The aim of the tools and devices developed to this end echoes Weber's quest for '[a] correct causal interpretation of a concrete course of action', which 'is arrived at when the overt action and the motives have been both correctly apprehended and at the same time their relation has become meaningfully comprehensible' (1978, p. 12). In other words, by offering tools that render DNN decision-making more transparent, a semi-Weberian explanation of the motivation underlying specific predictions might eventually be obtained. However, not only are the explainability tools that have been developed to date widely seen as being unsuccessful at addressing the fundamental opacity problem (e.g. Gilpin *et al.*, 2019; Kindermans *et al.*, 2019), but the quest for full transparency has also been debated. For example, it has been argued that, just as humans are only able to offer partial accounts of themselves, only partial accounts of algorithmic systems such as DNNs should be expected (Amoore, 2020).

We have discussed the notion of social action in light of the rise of second-generation automated trading systems, but we believe that these similarly prompt a reconsideration of other key sociological notions. To demonstrate this, we end with a brief discussion of the idea of performativity. Notably, MacKenzie has argued that, in some situations and under particular circumstances, economic formulas and models can have performative effects on markets when used in practice (MacKenzie, 2006; MacKenzie & Millo, 2003). Using the Black–Scholes options pricing model as his main example, MacKenzie's key point is that the practical use of an economic model may make economic processes look more like the way the model depicts them.

The performativity thesis has inspired a lot of critical discussion and further elaboration (e.g. Bamford & MacKenzie, 2018; Esposito, 2013; Fligstein & Dauter, 2007; Svetlova, 2012; Tellmann, 2020). In the present context, our interest is focused on its applicability to the domain of automated trading. MacKenzie himself has taken some steps in that direction. For example, he and Millo (2003, p. 128) describe how, in the early 2000s, the increasing use of handheld computers in the trading pits of the Chicago Board Options Exchange did not fundamentally affect the performativity of financial models: 'It can reasonably be said of this technosystem that it performs theory', they argued, meaning that financial theories were now increasingly technologically performed. It could similarly be argued that, in the case of first-generation automated trading systems, these might performatively enact any theories underlying their human-defined rules.

However, performativity looks markedly different for second-generation automated trading systems. To be sure, human expertise (including financial models and theories) may play an important role for such systems, and quants may define various indicators or seek to influence the ML system through various rewards and penalties. However, given the rule-generating independence of the ML system, any performative effects of human-defined models and theories would, at best, be highly convoluted, as they would be filtered through the ML system's internal *modus operandi*. It would, in other words, be difficult to establish a direct connection between a human-defined financial theory and the output and impact of the ML-based system's trading behaviour. It seems more reasonable to argue that rather than models potentially having performative effects through their incorporation into ML-based algorithmic trading systems, these systems may be generating their own financial models. With a nod to the title of MacKenzie's (2006) book on this topic, they become true model 'engines'. This is precisely how Tyler Capital conceives of it. The CTO stated, 'we're automating the process of creating financial models […] In the most extreme case, you could say [OPUS] develops a new financial model every day'. This shift is yet another instantiation of ML's inverse Durkheimianism we referred to earlier, that rules – or here, models – are not predefined and external, but developed from below, from inside the system itself. This process also means that to the extent that sociologists are interested in examining the performativity of ML-based systems, an important first step would be to understand how these systems develop their models and what their underlying rationale is. However, this is precisely the difficulty, given the opacity of DNNs.

Conclusion

As elsewhere in society, ML technologies are making inroads in financial markets, including automated securities trading. In this study, we have demonstrated some central ways in which the uptake of ML is transforming automated trading, focusing on how second-generation automated trading systems differ from their first-generation forerunners. In particular, we showed that, where first-generation systems are based on human-defined rules, second-generation systems develop their trading rules independently – in what amounts to an inverse Durkheimianism. We further suggested that where first-generation systems can be understood according to Weber's notion of social action – namely, as enacting human instrumentally rational action by taking account of the behaviour of other market participants – the situation is markedly different for second-generation automated trading systems. Here, decisions cannot be traced back to human subjective meaning, nor may they even be explainable in human terms, for which reason we consider Weber's notion inadequate to account for these types of ML systems. Finally, we argued that second-generation automated systems place the notion of performativity in a new light. Rather than these systems performing human-defined models, they produce new models independently of humans.

Although we have discussed some key aspects and implications of ML-based securities trading, there are several additional dimensions that call for further analysis. In line with those briefly listed in the introduction, further research is needed on the minute ways in which humans design and curate second-generation automated trading systems (addressing questions pertaining to data preparation, training methods, model testing, overfitting, and so on). Additionally, it would be important to examine how human-machine interactions in trading firms change with the adoption of ML systems, as well as what the limitations of using such systems for trading purposes are, including whether the deployment of self-learning trading systems makes markets more vulnerable to crashes and instability, particularly when these systems are characterized by opacity. DNN opacity itself begs further analysis, including what its combination with independent automated decision-making might mean for ethics and liability. Important directions for such work include those outlined by Amoore (2020), who proposes to make opacity the starting point for ethical considerations about ML rather than to call for a form of transparency which might never be attained; and Svetlova (2022), who argues for going beyond focusing on the opacity of individual algorithmic systems and consider the ethical questions their collective behaviours raise on a more systemic level.

Finally, we hope that empirical work on second-generation automated trading will inspire theoretical reflections on the viability of notions anchored in conceptions of human decision-making and subjectivity. Despite all the human efforts that go into designing and setting up ML-based systems, their independent rule-generating ability suggests that economic sociologists need to take seriously that – within securities trading but elsewhere as well – non-human actors now seem to be granted a level of agency previously reserved for humans. This not only takes us beyond Weber but possibly calls for a more fundamental rethinking of the social.

Notes

1 As exchanges do not engage in trading, we coded these interviews in accordance with the principal way of using of algorithms for matching orders ('order matching engines').
2 As mentioned earlier, this shift does not entail an absence of humans or human labour. For example, firms specializing in second-generation automated trading all report that cleaning and preparing data for ML training is extremely labour intensive (Borch, 2022; see also Lange *et al.*, 2019, pp. 609–610).
3 Margin trading means trading securities from borrowed funds from a broker, in effect, allowing for taking leveraged positions.

Acknowledgements

We thank Zachary David, Kristian Bondo Hansen, John Levi Martin, Alex Preda, Daniel Souleles, the *E&S* editors, and three anonymous reviewers for their valuable comments. We are particularly grateful to our informants, especially Mike Bushore, Chris Donnan, and the rest of the Tyler Capital community for allowing us to follow their work. We would also like to thank Caitlin Welch for her research assistance.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

Research for this paper received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 725706).

Ethical approval statement

Prior to starting this study, we obtained ethical approval for all protocols from the ERC and a local institutional review board at the Copenhagen Business School. Informants quoted in this paper gave their written informed consent prior to submission.

ORCID

Christian Borch http://orcid.org/0000-0001-8217-5880
Bo Hee Min http://orcid.org/0000-0003-2800-7370

References

Alonso, M. N., Batres-Estrada, G. & Aymeric, M. (2019). Deep learning in finance: Prediction of stock returns with long short-term memory networks. In T. Guida (Ed.), *Big data and machine learning in quantitative investment* (pp. 250–277). Wiley.

Amoore, L. (2019). Doubt and the algorithm: On the partial accounts of machine learning. *Theory, Culture & Society*, *36*(6), 147–169.

Amoore, L. (2020). *Cloud ethics: Algorithms and the attributes or ourselves and others*. Duke University Press.

Baker, W. E. (1984). The social structure of a national securities market. *American Journal of Sociology*, *89*(4), 775–811.

Bamford, A. & MacKenzie, D. (2018). Counterperformativity. *New Left Review*, *113*(September/October), 97–121.

Borch, C. (2021). Machine learning and social theory: Collective machine behaviour in algorithmic trading. *European Journal of Social Theory*. Retrieved from https://doi.org/10.1177/13684310211056010

Borch, C. (2022). Machine learning, knowledge risk, and principal-agent problems in automated trading. *Technology in Society*, *68*. Retrieved from https://doi.org/10.1016/j.techsoc.2021.101852

Borch, C. & Lange, A.-C. (2017). High-frequency trader subjectivity: Emotional attachment and discipline in an era of algorithms. *Socio-Economic Review*, *15*(2), 283–306.

Brayne, S. & Christin, A. (2021). Technologies of crime prediction: The reception of algorithms in policing and criminal courts. *Social Problems*, *68*(3), 608–624.

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, *3*(1). Retrieved from https://doi.org/10.1177/2053951715622512

Cevolini, A. & Esposito, E. (2020). From pool to profile: Social consequences of algorithmic prediction in insurance. *Big Data & Society*, *7*(2). Retrieved from https://doi.org/10.1177/2053951720939228

Christin, A. (2020). The ethnographer and the algorithm: Beyond the black box. *Theory and Society*, *49*(5–6), 897–918.

Coeckelbergh, M. (2020). *AI ethics*. MIT Press.

Coombs, N. (2016). What is an algorithm? Financial regulation in the era of high frequency trading. *Economy and Society*, *45*(2), 278–302.

Durkheim, É. (1982). *The rules of sociological method*. Macmillan.

Esposito, E. (2013). The structures of uncertainty: Performativity and unpredictability in economic operations. *Economy and Society*, *42*(1), 102–129.

Fligstein, N. & Dauter, L. (2007). The sociology of markets. *Annual Review of Sociology*, *33*(1), 105–128.

Fourcade, M. & Healy, K. (2017). Seeing like a market. *Socio-Economic Review*, *15*(1), 9–29.

Fourcade, M. & Johns, F. (2020). Loops, ladders and links: The recursivity of social and machine learning. *Theory and Society*, *49*(5–6), 803–832.

Gane, N. (2012). *Max Weber and contemporary capitalism*. Palgrave Macmillan.

Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M. & Kagal, L. (2019). *Explaining explanations: An overview of interpretability of machine learning*. Retrieved from https:/rxiv.org/pdf/1806.00069.pdf

Hansen, K. B. (2020). The virtue of simplicity: On machine learning models in algorithmic trading. *Big Data & Society*, *7*(1). Retrieved from https://doi.org/10.1177/2053951720926558

Hansen, K. B. (2021). Model talk: Calculative cultures in quantitative finance. *Science, Technology, & Human Values*, *46*(3), 600–627.

Hansen, K. B. & Borch, C. (2021). The absorption and multiplication of uncertainty in machine-learning-driven finance. *British Journal of Sociology*, *72*(4), 1015–1029.

Kindermans, P.-J., Hooker, S., Adebayo, J., Alber, M., Schütt, K. T., Dähne, S., … Kim, B. (2019). The (un)reliability of saliency methods. In W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen & K.-R. Müller (Eds.), *Explainable AI: Interpreting, explaining and visualizing deep learning* (pp. 267–280). Springer.

Kiviat, B. (2019). The moral limits of predictive practices: The case of credit-based insurance scores. *American Sociological Review*, *84*(6), 1134–1158.

Knorr Cetina, K. & Bruegger, U. (2002). Traders' engagement with markets: A postsocial relationship. *Theory, Culture & Society*, *19*(5/6), 161–185.

Koza, J. R. & Poli, R. (2014). Genetic programming. In E. K. Burke & G. Kendall (Eds.), *Search methodologies: Introductory tutorials in optimization and decision support techniques* (pp. 143–185). Springer.

Lange, A.-C. (2016). Organizational ignorance: An ethnographic study of high-frequency trading. *Economy and Society*, *45*(2), 230–250.

Lange, A.-C., Lenglet, M. & Seyfert, R. (2019). On studying algorithms ethnographically: Making sense of objects of ignorance. *Organization*, *26*(4), 598–617.

Latour, B. (2005). *Reassembling the social: An introduction to actor-network theory*. Oxford University Press.

Lenglet, M. (2011). Conflicting codes and codings: How algorithmic trading is reshaping financial regulation. *Theory, Culture & Society*, *28*(6), 44–66.

Luhmann, N. (1995). *Social systems*. Stanford University Press.

MacKenzie, D. (2006). *An engine, not a camera: How financial models shape markets*. MIT Press.

MacKenzie, D. (2018a). 'Making', 'taking' and the material political economy of algorithmic trading. *Economy and Society*, *47*(4), 501–523.

MacKenzie, D. (2018b). Material signals: A historical sociology of high-frequency trading. *American Journal of Sociology*, *123*(6), 1635–1683.

MacKenzie, D. (2021). *Trading at the speed of light: How ultrafast algorithms are transforming financial markets*. Princeton University Press.

MacKenzie, D. (2022). Spoofing: Law, materiality and boundary work in futures trading. *Economy and Society*, *51*(1), 1–22.

MacKenzie, D. & Millo, Y. (2003). Constructing a market, performing theory: The historical sociology of a financial derivatives exchange. *American Journal of Sociology*, *109*(1), 107–145.

Min, B. H. & Borch, C. (2022). Systemic failures and organizational risk management in algorithmic trading: Normal accidents and high reliability in financial markets. *Social Studies of Science*, *52*(2), 277–302.

Pardo-Guerra, J. P. (2019). *Automating finance: Infrastructures, engineers, and the making of electronic markets*. Cambridge University Press.

Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K. & Müller, K.-R. (Eds.) (2019). *Explainable AI: Interpreting, explaining and visualizing deep learning*. Springer.

Scharre, P. (2018). *Army of none: Autonomous weapons and the future of war*. W. W. Norton & Company.

Seyfert, R. (2016). Bugs, predations or manipulations? Incompatible epistemic regimes of high-frequency trading. *Economy and Society*, *45*(2), 251–277.

Souleles, D. (2019). The distribution of ignorance on financial markets. *Economy and Society*, *48*(4), 510–531.

Stilgoe, J. (2018). Machine learning, social learning and the governance of self-driving cars. *Social Studies of Science*, *48*(1), 25–56.

Svetlova, E. (2012). On the performative power of financial models. *Economy and Society*, *41*(3), 418–434.
Svetlova, E. (2022). AI ethics and systemic risks in finance. *AI and Ethics*. Retrieved from https://doi.org/10.1007/s43681-021-00129-1
Swedberg, R. (2000). Afterword: The role of the market in Max Weber's work. *Theory and Society*, *29*(3), 373–384.
Tellmann, U. (2020). Beyond performativity: Material futures of finance. *Economy and Society*, *49*(3), 345–363.
Weber, M. (1978). *Economy and society: An outline of interpretive sociology*. University of California Press.
Weber, M. (2000). Stock and commodity exchanges ['Die Börse' (1894)]. *Theory and Society*, *29*(3), 305–338.
World Economic Forum. (2018). *The new physics of financial services: Understanding how artificial intelligence is transforming the financial ecosystem*. World Economic Forum.

Christian Borch is a Professor of Economic Sociology and Social Theory at the Copenhagen Business School. His latest book is *Social avalanche: Crowds, cities and financial markets* (Cambridge University Press, 2020).

Bo Hee Min is an Assistant Professor of Economic Sociology at Copenhagen Business School. Her research examines how technology, algorithms and artificial intelligence transform work, organizations and markets.