What is an algorithm? Financial regulation in the era of high-frequency trading

What is an algorithm? Financial regulation in the era of high-frequency trading

Nathan Coombs

Abstract

In response to the flash crashes and market manipulations blamed on high-frequency trading, algorithms have been brought inside the regulatory perimeter. This article focuses the most ambitious regulation directed at the practice: the algorithm tagging rule in the German High-Frequency Trading Act. Fifteen interviews with stakeholders in the Act's implementation serve to reconstruct how regulators defined an algorithm and help pose the question of to what extent regulatory definitions and data need to accurately represent financial practices to be useful. Although tentative in its findings, the research suggests that the algorithm tagging rule may be providing valuable signals in the noise to trade surveillance officers and having virtuous effects on the cultures of trading firms. The conclusion argues that sociologists of finance should adopt a more balanced approach when evaluating regulatory technologies and heed MacKenzie's call to open up their black boxes.

Keywords: financial regulation, high-frequency trading, algorithms, compliance, enforcement

Introduction

Since the Global Financial Crisis a drive to control systemic risk in the financial system has led regulators to augment their knowledge (IMF & FSB, 2009). While most pronounced in banking regulation, with steps towards monitoring the credit cycle and mitigating its excesses, greater transparency is also being brought to securities markets (Dorn, 2012). For instance, both the Dodd-Frank Act (US Government, 2010) and the European Market Infrastructure Regulation (European Union, 2012) are helping to build up a detailed database of trades by requiring centralised clearing of derivatives. In a similar spirit, due to difficulties encountered in determining who owns what after the collapse of Lehman Brothers regulators are encouraging financial institutions to adopt the global legal entity identifier (Lee, 2013).

Even trading technologies, traditionally far outside the regulatory perimeter, are becoming subject to transparency requirements. A string of destabilising events associated with algorithmic and high-frequency trading have caused concern: from the 'Quant Meltdown' in 2007 to the errant algorithm that almost bankrupted Knight Capital in 2012 (Kirlenko & Lo, 2013). But the focal point remains the 'Flash Crash'. On 6 May 2010 the Dow Jones dropped 9% in a matter of minutes following a clumsy trade by a mutual fund which resulted in high-

frequency trading algorithms withdrawing their liquidity *en masse* (CFTC & SEC, 2010). With regulators sensitive to emerging risks after the financial crisis such volatilities were certain to set rule-making in motion. The review of Europe's Markets in Financial Instruments Directive responded with stringent algorithm testing requirements and the suggestion that firms should 'notify their competent authority of the computer algorithm(s) they employ, including an explanation of its design, purpose and functioning' (European Commission, 2010, p. 16). More ambitious was Germany's High-Frequency Trading Act (Bundesrat, 2013). The Act's algorithm tagging rule requires firms to identify with a number which algorithm is used to generate a trading decision (Hessisches Ministerium, 2014). The intention was to allow trade surveillance officers in financial exchanges to see individual algorithms at work in their order-books and to help detect attempts at market manipulation (such as the 'spoofing' strategy for which the British trader Navinder Sarao stands accused by the U.S. Department for Justice).[1]

Based on fifteen in-depth interviews with stakeholders in the HFT Act, this article concerns the vexed process through which its algorithm tagging rule was implemented. At its heart are the difficulties regulators encountered in defining an algorithm. As the first attempt to institute the surveillance *of* algorithms in any walk of economic life – inverting conventional understandings of algorithmic governance (Amoore & Piotukh, 2015) – regulators would have to move beyond a theoretical understanding of an algorithm to a definition capturing the salient characteristics of the algorithms employed by algorithmic and high-frequency trading firms. In so doing practical solutions were needed to seemingly philosophical questions: what individuates an algorithm?, where does one algorithm stop and another begin?, should an algorithm be defined with respect to code or is an understanding of its strategy sufficient? The outcome was not only a regulatory definition different to what trading firms see as their algorithms, but also ambiguities as to what constitutes a 'material change' in an algorithm. With compliance officers left to exercise their own judgement on the matter, the result has been substantial, and potentially worrying, discrepancies the number of tags trading firms are generating.

Do these findings support Dorn's (2012, p. 318) contention that regulators' attempts to know markets at a granular level are creating a 'wall of incommensurate, uninterpretable, overwhelming information'? To what extent do regulatory data need to accurately represent financial practices to be useful? In the social studies of finance, where epistemological questions have received the most attention, there is a degree of scepticism about regulatory knowledge practices. Guided by the pragmatist concepts of 'performativity' and 'evaluation cultures' scholars have affirmed that the models and devices made use of by market participants need not be accurate to be useful (Callon, 1998, 2007; MacKenzie, 2006a, 2006b, 2008, 2011; MacKenzie & Millo, 2003; Callon *et al.*, 2007; MacKenzie *et al.*, 2007; Millo & MacKenzie 2009; MacKenzie & Spears, 2014). Such devices simply provide tools, conventions,

[1] Spoofing involves using a trading algorithm to rapidly place orders above or below the median market price with no intention of fulfilling them. In so doing, the 'spoofer' is attempting to profit from inducing a price movement in the market.

and cognitive facilitators for pricing financial instruments. In contrast, studies of regulation within the field have been quite critical of the use of technology to help get to the 'truth' of the market (Williams, 2009; Lenglet, 2011). Stressing the distance between regulators and the practices they attempt to know, these studies provide a sociotechnical twist on the 'Hayekian critique of regulation' (Riles, 2011, chap. 4), arguing that the complexity of automated financial markets is encasing the regulatory enterprise in an impenetrable epistemic fog.

Although tentative in its findings, this study presents a more positive view on the affordances of regulatory technologies. Given the difficulties of evaluating processes of regulatory compliance and enforcement (Parker & Nielsen, 2009), as well as the secrecy of trade surveillance departments, it does not claim that the German HFT Act's algorithm tagging rule has been an unmitigated success. (Further research would be necessary to reach a firm conclusion.) Instead, the analysis seeks to demonstrate two ways in which the rule might be proving useful even given the incommensurate compliance responses of trading firms. It does so, first, by giving a generous hearing to regulators' claim that the algorithm tags are yielding valuable *relational* information about the interactions of trading firms' algorithms. Secondly, by showing that compliance with the rule is reformatting knowledge hierarchies in trading firms. In line a trend that has seen the compliance profession rise in stature since the financial crisis, the algorithm tagging rule, we will see, is empowering compliance officers to render explicit the tacit knowledge of traders and algorithm designers – an unintended yet virtuous 'performation' (Callon, 1998) of market practices addressing some of regulators' concerns.

The discussion begins by asking why regulatory technologies continue to be marginalised in the social studies of finance. The second part locates the motivations for the algorithm tagging rule in problems identified by trade surveillance officers. The third part then teases out the indeterminacies in regulators' definition of an algorithm. The fourth part addresses compliance with, and enforcement of, the tagging rule. The discussion and conclusion reflect on these findings and ask scholars to heed MacKenzie's (2005) call to open up the technological black boxes of global financial regulation.

Knowing regulatory knowledge

The heightened importance of regulatory technologies since the Global Financial Crisis has received remarkably little scholarly attention. It is unsurprising, however, given that regulation remains a peripheral subject in the sociology of finance. Take, for example, the *Oxford Handbook of the Sociology of Finance* (Knorr-Cetina & Preda, 2012) in which regulation plays a role in the stories presented by contributors, but in which no chapter addresses regulation per se. The same applies to key edited volumes in the social studies of finance (Callon *et al.*, 2007; MacKenzie *et al.*, 2007). The neglect may have been understandable until the 2007-8 banking crisis. From the revisions to the Basel Accords in the 1990s onwards regulatory bodies have tended to outsource modelling to the private sector, meaning that that regulation was best comprehended at the legal, political and ideological levels. But the neglect is less understandable in the post-crisis world in which regulators have been introducing their

own sophisticated models and devices. Central bank supervised stress tests of the financial system are one important example (Langley, 2013; Geithner, 2014; Bank of England, 2015). Another are attempts by the International Monetary Fund to map the 'shadow banking' system (Errico *et al.*, 2014). In securities markets regulation, an especially grandiose project in the United States aims to institute a supercomputer-powered Consolidated Audit Trail, unifying the data of the country's fragmented trade reporting organizations (Brigagliano *et al.*, 2012). All these developments would seem to be suitable, indeed pressing, subjects for research at the intersection of the technological and the micro-sociological. And yet, regulation remains stubbornly at the margins of the social studies of finance.

Those studies which have engaged regulatory technologies tend to view the practices they support as problematic at best, illusory at worst. Williams's (2009; 2013) ethnographic account of trade surveillance company, for example, focuses on the overwhelming quantity and superficial quality of the data surveillance officers are charged with making sense of. Williams even goes so far as to call for a shift from an 'ocular' paradigm – associating regulatory representations of financial markets with illumination and vision – to a 'penumbral' paradigm – concerned more with what they obscure than illuminate (2009, pp. 485–486). Lenglet (2011) is sceptical too, seeing algorithms as undermining notions of accountability in regulatory enforcement that remain wedded to a human-centric model of agency (an argument taken further by Lenglet and Mol (2015) who argue that with algorithmic trading financial markets have become intrinsically ungovernable). Seyfert (this issue) mines the same critical vein when he identifies the practices of market authorities with an 'epistemic regime of authorisation' bent on reducing emergent phenomena to attributions of individual blame. These scholars thus provide a sociotechnical twist on what Riles (2011, chap. 4) calls the 'Hayekian critique of regulation'. Where Hayek (1948) theorised the inefficiency of state planning based on neo-Kantian postulates about the superiority of individuals' situated knowledge, these studies explore empirically how regulatory knowledge fails in capturing the complexities of electronic and automated markets.

This article proposes a different approach. Recognizing epistemological challenges as an intrinsic feature of the regulatory enterprise, it suggests that we instead ask *to what extent* regulatory definitions and data really need to be accurate. In other words, like the market devices theorised in depth by the social studies of finance, can regulatory data represent only inaccurately and incompletely financial practices but still have their uses? The shift in emphasis has methodological consequences. Rather than focusing on well-established regulatory practices and searching out their lacunae, this approach calls for reconstructing the regulatory *process* in order to link up the problems motivating the introduction of rules to the ways in which the resulting data is employed. With regulators cognizant of the limits to their knowledge and increasingly enrolling private sector actors in rule-making (Black, 2003), the sociological contribution of this approach therefore concerns exploring processes of 'joint regulation' (Huault *et al.*, 2012), or what socio-legal scholars term 'meta-regulation' (Gilad, 2010). It involves investigating how financial practices are defined by regulators in

collaboration with market participants, and the consequences of indeterminacies in regulatory rules, classifications and definitions.

For three reasons, research on the implementation of the German High-Frequency Trading Act is well-suited to the approach. The first relates to the nature of the financial practices. Over the last decade and a half open outcry pits and manual click trading have been crowded out by the rise of algorithmic and high-frequency trading (Arnuk, 2012; Pardo-Guerra, 2012; Patterson, 2012; Lewis, 2014; MacKenzie *et al.*, 2012; MacKenzie & Pardo-Guerra, 2014; Borch *et al.*, 2015; MacKenzie, 2015; Arnoldi, 2016; Thompson, 2016). With computer trading now accounting for the majority of liquidity on U.S. and European markets for equities and standardised derivatives, the use of high-speed algorithms – short pieces of code based on iterative decision-making procedures – have accelerated trading into sub-perceptual speeds. Although evolving out of more rudimentary algorithmic trading practices and often considered a 'sub-set' of them (Aldridge, 2013, pp. 13–14), high-frequency trading is not a means for executing a value investment strategy (Kunz & Martin, 2013). Instead, it focuses on three strategies: market making, index and venue arbitrage, and 'fast alpha' – high-speed price prediction based on quantitative technical analysis (Narang, 2013, pp. 271–273). The upshot is that official reports and scholarship have noted how high-frequency trading strategies have resulted in an enormous increase in the quantity of market data, posing problems for regulatory attempts to know the market (CFTC & SEC, 2010; IOSCO, 2011; Pardo-Guerra, 2012; Foresight, 2012; Joint Staff, 2015). It is therefore not yet clear if traditional forms of regulation will suffice to meet these challenges, or whether, as Kirilenko and Lo (2013) put it, a more radical form of 'regulation 2.0' is needed. An analysis of the German HFT Act contributes to this ongoing debate.

The second reason concerns the obscure epistemic status of algorithms. An algorithm is not a thing that can be reached out and touched, but rather a decision-making *process*. To take the example of an extremely simple limit order execution algorithm: IF financial asset X is available at specified price THEN buy. In more sophisticated high-frequency trading algorithms human mediation is removed altogether so that the IF/THEN sequence of instructions can be executed thousands of times a second in response to changes in the market. Being a process rather than a financial instrument or trading entity it is not immediately obvious how an algorithm can be represented. As Lenglet (2011, pp. 51–52) points out, 'algorithms are not easy to represent… should we assume it [an algorithm] resides in the coding describing a trading pattern and the related actions to take depending on the market's context? Or should we say the algorithm is displayed on the trader's screen, in the window allowing him to set the parameters, which shaped the algorithm?' Perhaps more importantly, algorithms live a double-life: identifiable with their strategies as well as with the computer code in which strategies are operationalised. These representational challenges are not just abstract philosophical concerns. Existing regulatory texts define an algorithm in a variety of different ways, evidencing the difficulties of bringing algorithms into legal codes (Lenglet & Mol, 2015). And yet, this is precisely what the HFT Act's algorithm tagging rule attempts to do: it is the

first attempt to rigorously define an algorithm and to institute a surveillance mechanism rendering algorithms visible for regulatory enforcement. Given its ambition, the vexed story of the rule's implementation speaks directly to the question of to what extent regulatory definitions and data need to accurately represent financial practices to be useful.

The third reason relates to the rule's innovative nature. At the time of writing there are still compliance officers grappling with its metaphysical subtleties. And even if trade surveillance departments in German exchanges are now used to working with the data, the fact that a version of the rule is set to be introduced Europe-wide with the second Markets in Financial Instruments Directive (MiFID II) means that there are still regulatory bodies uncertain as to what to do with it (Interview 01/12/2014; Interview 09/01/2015). Likewise, the rule's novelty means that its implementation remains fresh in the minds of its stakeholders – a felicitous condition for research. More broadly, this study answers the call by regulatory scholars for more research on how regulatory 'rules and models are reshaped during their implementation' (Schneiberg & Bartley 2008, p. 49). Processes of 'joint regulation' (Huault *et al.*, 2012) and 'meta-regulation' (Gilad, 2010), where interactions between regulators and regulated determine the meaning of regulations, are something more often theorised than investigated empirically (an exception is Thiemann & Lepourte, forthcoming). Such fine-grained work is especially important here however since trading algorithms are not the result of regulatory arbitrage (Funk & Hirschman, 2014), or a 'regulatory dialectic' (Kane, 1988) between regulators and financiers, but rather a hitherto unregulated technology. Understanding how the regulatory definition of an algorithm was negotiated across a network of public and private actors therefore provides a window into the challenges of 'regulatory innovation' (Black, 2005) in the 21st century.

For all these reasons, in 2014 and 2015 I conducted fifteen individual and group interviews with stakeholders in the HFT Act's implementation: a pool of twenty interviewees in total. I spoke to people who could take me beyond the regulatory documents and say something about the motivation for the rule, how it was being implemented, and how the data is being used. On the public side, this included a German regulator, a senior trade surveillance officer, a market surveillance and forensics team, and European regulators who face implementing a similar rule with the arrival of MiFID II. On the private side, I spoke to financial technologists, compliance officers at algorithmic and high-frequency trading firms, and industry representatives of the proprietary trading industry. The aim in speaking to both regulators and regulatees was to see how they collectively responded to the definitional problem posed by the algorithm tagging rule. An ancillary goal was to investigate how compliance officers are responding to the rule's indeterminacies. These interviews were supplemented by a review of reports, investigations, and responses to public consultations by American and European regulatory agencies and transnational bodies. The sections which follow present the findings of the research in a more or less chronological reconstruction of the regulatory process. The question of to what extent regulatory definitions and data need be to be accurate to be useful is addressed in final section.

The political and the epistemic

In one crucial respect financial and regulatory innovations are asymmetrical. Whereas the former are often presented as the work of economic theorists seeking to rationalise market life, the latter are intrinsically political even at their most technocratic (cf. Porter, 2003). With the German HFT Act, the tension between the political and the epistemic is evident in the split motivations for the Act as a whole and its algorithm tagging rule. Both public and private sector interviewees agreed that the HFT Act was a response to the failure Angela Merkel's coalition to introduce a Financial Transaction Tax after the Global Financial Crisis (Handke & Zimmerman, 2012). As a German regulator put it, the HFT Act served to satisfy the public's appetite for retribution against the financial sector by tapping into the fact that the 'average newspaper reader feels uneasy with the idea of high-frequency trading and things happening within microseconds or nanoseconds.' (Interview 21/04/2015)

The algorithm tagging rule, in contrast, derived from epistemic problems identified by trade surveillance officers in German exchanges. Trade surveillance departments are generally located in financial exchanges and it is their job to watch the aggregate order book of bids and offers for different financial assets with the aim of ensuring the smooth functioning of the market and attempting to detect manipulative activities. Due to the enormous quantity of data, surveillance officers do not watch the order-book directly but instead make use of analytic software based on pre-programmed alerts for suspicious behaviour (Williams, 2009, 2013). After receiving an alert surveillance officers then apply visualisation tools on the suspicious trades in order to decide whether they warrant further investigation.

The way it was described to me is that although the trade surveillance department at the Deutsche Börse Group was aware of the HFT Act's opportunistic origins, they nevertheless sought to take advantage of the new legislation to introduce a rule that would assist their practices.[2] In contrast to scholarship which identifies temporality as the main challenge involved regulating high-frequency trading (Angel, 2014; Thompson, 2016), the senior trade surveillance officer I spoke to stressed that the problems they were facing had little to do with speed. The software platforms trade surveillance make use of already allow them to make sense of market events at any timescale: 'backwards, forwards, for five years, in graphical background, in slow motion, we can do videos on it, everything fine no problem.' If the speed of trading algorithms is a problem it is only insofar as it creates an expectation for faster enforcement response times. 'In the old world… you have two weeks' time and so on, but when you are acting in an algorithm environment it doesn't make sense to wait two, three, four weeks, whatever, until you get the right information.' (Interview 21/04/2015)

Both regulators and trade surveillance identified the main problem with the opaque

[2] The trade surveillance department at the Deutsche Börse Group were able to exercise this level of influence due to the fact that the Frankfurt Stock Exchange and the derivatives exchange Eurex fall under German administrative law. This provided a direct channel of communication for trade surveillance to suggest the inclusion of the algorithm tagging rule into the drafting of the federal Act (Interview 21/04/2015).

cause and effect of market events when algorithms are involved. Reconstructing cause and effect involves determining how a trading strategy influences the subsequent trades of other market actors. So, in what is called a 'momentum ignition' strategy, where a trading firm X attempts to induce a rapid price movement in the market, reconstructing cause and effect involves establishing that it was indeed the trades of firm X that were responsible. That is what has become more difficult in the era of algorithmic and high-frequency trading. As a German regulator put it, when faced with suspicious market behaviour it is now not enough to know not just which trader is making a particular trade, what company they work for, or what type of trading they are engaging in. To be effective, trade surveillance now need to know which algorithm a trader is making use of.

> Is it this algo? Is it a different algo? Or what happens if an algo went wrong? Can we identify that he's wrong? … And that's the reason why we said we need more transparency about the source – I would say data owner, perhaps that's a good definition. We know it's a company XYZ, but we don't know is it the algorithm ABC, what is the difference, can we identify it, and so on. (Interview 21/04/2015)

Trade surveillance were also confronting problems related to inferring trader intentionality. Even if it looks like trades bear the hallmarks of manipulative activity, with the interactions in the order book of a financial exchange being so complex it is difficult to know whether or not they were the result of a deliberate strategy. Since a trader can employ algorithms to pursue legitimate trading strategies and yet also pursue a parallel manipulative strategy, with only individual trader IDs to go by surveillance officers cannot see what is going on in fine enough detail to make an informed judgement. A financial technologist explains:

> You can't clearly see any patterns anymore because just one single trader can create a thousand orders and actually operate six different strategies at the same time, which obviously then makes the job for your surveillance management more difficult … because a trader usually doesn't do 100% market abuse, usually if there is any abuse usually they have 99% legitimate business and then they throw in a couple of fishy trades where they hope to get a benefit or a profit. (Interview 5/11/2014)

The final problem German regulators identified that the use of algorithms has resulted in traders' own understanding of their practices becoming foggy (a matter of concern after the Knight Capital incident). As later discovered in the regulators' consultation with market participants, only some of the largest European and US high-frequency trading firms, which the regulators estimated as 20% of the market, could supply documentation explaining their algorithms' operation (Interview 21/04/2015). For the non-specialist trading desks of banks, who only embraced the use of algorithms due to the demands of clients, the response to regulators' requests for documentation was 'oh no this is a third-party algorithm, we don't

know exactly how it works, oh we don't use it, only in parts.' This a regulator described as a 'bad surprise; we did not expect so many players in the market giving the impression that they do not really know what they are doing.' Regulators were shocked to discover that 'nobody in the firm knows sometimes how the algorithm works. What is a decision path? I don't know. I push this button. Full stop.' (Interview 21/04/2015)

To sum up, although the German HFT Act capitalised on the public's fears about the finance-technology nexus, the idea of algorithm tagging was motivated by the lack of knowledge by trade surveillance and market participants about what algorithms are doing. These motivations will prove important later in this article when considering how trade surveillance are making use of the data supplied by the algorithm tagging rule and when examining its effects on the cultures of trading firms. But first, it is necessary to address how regulators defined an algorithm. In the same way that political and epistemic considerations motivated the HFT Act, these considerations also intermingled in shaping regulators' definition of an algorithm.

Constructing the regulatory definition

The modern understanding of an algorithm is generally traced back the Turing Machine thought experiment of the 1930s, with 'algorithm' denoting the finite set of step-by-step decision procedures involved in solving a calculative problem (Blass & Gurevich, 2003; Goffey 2008, pp. 15–16). In financial markets, the calculative problem facing investors and brokers concerns how to execute an order to buy or sell a financial asset in such a way that minimises its market impact (Kim, 2007, chap. 5).[3] In the case of high-frequency trading, which aims to extract alpha from the trading process itself (Narang, 2013), algorithms are used more widely for everything from sentiment analysis of Twitter (Karppi & Crawford, 2015) through to optimising order routing to exchanges. Indeed, ethnographic research within high frequency trading firms shows that strategists and coders develop their own bespoke algorithms, which are only later combined in a firm's matching engine (Lange, this issue).

Implementing the HFT Act's algorithm tagging rule therefore required regulators to move beyond an abstract, theoretical understanding of an algorithm to a definition which would capture their diverse functions in financial trading. Most challenging would be to settle at a sufficient level of abstraction: generic enough to apply to the procedures of every firm but without being too loosely defined. This section elaborates how regulators approached the definitional process, showing how it required navigating between code-based requirements and the potential for interpretative chaos. It then demonstrates how this these considerations shaped the indeterminacies of the regulatory definition.

Between code disclosure and interpretative chaos

[3] A common algorithm employed is the volume weight average price (VWAP). An order is normally split up into smaller 'child orders' which are then executed across single or multiple days in reference to the benchmark price (price times the number of assets traded divided by the total volume of the asset traded on a given day).

Industry fears about their intellectual property rights played a large part in shaping the algorithm tagging rule.[4] Although no proposals have suggested that regulators should have direct access to the code of trading firms' algorithms, rules mooted by the review of Europe's Markets in Financial Instruments Directive (European Commission, 2010) were seen as opening the door to the possibility. For instance, the American Chamber of Commerce warned that 'market participants would be strongly against proposals that gave regulators access to the functional specification or the code of the algorithms.' (AmCham EU, 2011, p. 3) The British Banking Association, similarly, insisted that description should be 'limited to high level demonstration of the algorithmic strategy, since a requirement to provide a detailed description of each algorithm… would risk opening the intellectual property concerning the algorithm to the market.' (BBA, 2011, p. 19) Other respondents questioned whether regulators would have the expertise to make use of the information. The Association for Financial Markets in Europe asked whether national regulators will 'be required to review and assess each algorithm (some of which are highly complex) and will they have resources to do this?' (AFME, 2011, p. 31) The Futures and Options Association agreed, writing that such a requirement was 'likely to be meaningless.' (FOA, 2011, p. 15)[5]

While sensitive to these concerns, German regulators saw more immediate problems with code disclosure. My interviewees pointed out that if traders were to attempt manipulative strategies they were not likely to disclose those elements of their algorithms' code to the authorities (Interview 21/04/2015). Moreover, they noted that since algorithms are being modified so frequently – on average every few weeks – then even with a large number of analysts devoted to the task it would be impossible to monitor all the code. Hence, trade surveillance proposed algorithm tagging as an alternative way to gain information about what algorithms are doing in the market. Yet regulators were aware that they were entering uncharted territory in putting the idea into practice. Just as the rule's motivation derived from a lack of knowledge about what algorithms are doing, market participants' conflicting interpretations about what an algorithm is proved no surer guide for rule-making. A German regulator told me that after consultation with the industry over a twelve month period they reached limits to the process beyond which no additional clarity was being gained:

> We came to the question: so, what is an algorithm really? When you discuss with market members you get a multitude of different options. So one says its each IF/THEN clause is an algorithm. And the next one says my strategy is an algorithm. And so we have everything in between. (Interview 21/04/2015)

[4] There have been similar responses to Regulation Automated Trading (CFTC, 2015) in the United States. The proposal that firm lodge their code at an external data repository elicited howls of protest, with industry groups objecting that 'it's inconceivable that any firm should be expected to leave its intellectual property on the doorstep of the government' (Meyer, 2015).

[5] Interestingly, this question mirrors methodological discussions unfolding in the digital humanities with respect to the value of knowing algorithms' code (see Rieder & Röhle, 2012)

The sense that they were addressing a profoundly philosophical question was shared by market participants. A compliance officer at a high-frequency trading firm spoke of consulting with his father, a computer scientist in the 1960s, when seeking to get his head around the problem (Interview 10/04/2015). Other firms advocated understandings of what an algorithm is based on their own proprietary procedures. The divergent responses, then, sometimes hinging on subtle distinctions about what distinguishes a strategy 'as a whole' from an individual algorithm, posed the danger of interpretative chaos in the absence of a more stringent definition.

What regulators took away from the process was that 'we have to define it more closely, what an algorithm is and how we understand it, not leaving it to the members because every member says something different.' (Interview 21/04/2015) Regulators would need a definition of an algorithm at a level of abstraction sufficient to avoid the need for code disclosure, while also not being so abstract – for instance, at the level of the strategy 'as a whole' – where firms would be free to tag whatever they saw fit.

The definition and its indeterminacies

To see how regulators sought to reconcile these demands it is necessary to attend to the Act's implementing guidelines. Reflecting the innovative nature of the rule – the first attempt to institute surveillance *of* algorithms in any walk of economic life – these guidelines have an unfamiliar grammar and epistemic subtlety making them difficult to interpret (even for the generously remunerated compliance teams of trading firms). One of the most challenging issues which faced the regulators drafting the guidelines was how to individuate an algorithm. In response to a discussion paper released by the European Securities and Markets Authority (ESMA, 2014), the Futures Industry Association capture its classically metaphysical difficulties.

> Some trading algorithms are made up of a sum of software sub-parts, which are built to receive and relay data from other correlated markets, or perform other separate calculations to handle complex events… This makes it extremely difficult to identify whether the overarching algorithm is a unique 'whole' or has changed by means of its sub-parts. (FIA, 2014, p. 46)

Put differently, with firms making use of a distributed assemblage of algorithms in their trading processes, how can the part be separated from the whole? The way in which regulators attempted to resolve this problem is the ideational innovation at the heart of this story. The rule's technical guidelines define an algorithm in the following way:

> A trading algorithm is an EDP-operated algorithm containing a well-defined sequence of instructions of finite length… [it] has to be identified [with] the entire sequence of calculation steps (decision path)… the identification obligation is referred to as a sequence of instructions and not to its individual elements, even if the latter could be considered separately as independent algorithms. (Hessisches

Ministerium, 2014, p. 1)

Rather than accepting firms' own understandings of what differentiates their algorithms, the requirement to tag the *entire* sequence of automated decision-steps as *the* algorithm includes all the elements which firms themselves might consider as 'independent algorithms'. So if, for example, a firm employs two algorithms – one for exploiting an arbitrage opportunity, another for determining volume – then the IF/THEN pathways of these algorithms should be combined and tagged with just a single numerical code. In so doing, regulators resolved the question of the part and the whole: they would not be interested in what firms consider their individual algorithms and would only require the identification of the complete automated set of decisions. Consequently, what firms consider an individual algorithm might be a huge number of different 'algorithms' from the perspective of the tagging requirements. A senior market surveillance officer explains that, 'you may have a situation where you [a trading firm] have one arbitrage algorithm based on your *internal* definition, but you have to mark [assign separate algorithm tags to] 5000 different paths.' (Interview 21/04/2015)

According to the rules, it is not enough however just to tag an algorithm once and for all. As previously noted, trading firms modify their algorithms on a regular basis. The rules therefore require a new tag to be assigned with every 'material change' in an algorithm. What does this mean? To determine when a material change has taken place, the rules focus on (1) changes in the decision-making sequence and (2) in the algorithm's parameters. To take the sequencing requirement first, if trading orders are initially determined by passing through what trading firms would consider as algorithms 1,2 and the sequence is changed to 2,1 then a new tag needs to be generated (see figure 1). So, for example, if the volume of an order is initially determined before the trading venue and then this sequence is reversed then a material change in the algorithm is considered to have taken place. Secondly, the rules also dictate that a new tag should be assigned with every material change in an algorithm's parameters. 'Parameter' is the name given to a variable in an algorithm. Parameters can be both qualitative – for example, the financial instrument – and quantitative – for example, order volume.

INSERT FIGURE 1 HERE

Three ambiguities and indeterminacies stand out in these rules. The first is that the

'algorithm' firms are required to tag is not the same what they might see as their algorithms. The regulatory definition concerns only the combined set of IF/THEN decisions responsible for generating a trading order. The distinction is an intelligible one but it provides scope for confusion. The second source of ambiguity concerns the identification of parameters (the rules identify only seven consequential parameters). In reality, these parameters may only partially coincide with the parameters of trading firms' algorithms; and so, some elements of the decision-path taken by firms' automated trading processes may be excluded. Third, the rules displace the question of what constitutes a material change in an algorithm to the open question of what constitutes a material change in a parameter. The aim in doing so was to have trading firms to come up with their own solutions to the problem. A regulator explains:

> of course its always a little undefined what is material… When you have the same parameters but it is only 1.0 or 1.2 for example who cares, it's still the same way… And normally even those guys [in large trading firms] know if its material because they have to hedge something else. (Interview 21/04/2015)

The assumption motivating this loose end is that only firms possess the detailed knowledge required for implementation of the rule. While at first glance then the HFT Act seems a highly prescriptive regulation, then, on close inspection something closer to meta-regulation is in evidence, where regulators allow firms to 'tailor regulation to their individual circumstances, while holding them accountable for the adequacy and efficacy of their internal control systems.' (Gilad, 2010, p. 485) In a practical acknowledgement of the 'finitist' (MacKenzie, 2008, pp. 26–30) perspective on classificatory rules to simply beget more rules *ad infinitum* if trying to account for every contingency, regulators considered it better to leave it to firms to decide what constitutes a material change in an algorithm. Regulators were aware that if specifying their rules at too granular a level then they would be dragged inexorably back towards a code-based definition of an algorithm, the avoidance of which motivated the tagging rule in the first place.

The result of these indeterminacies has not only been to make the rule difficult to implement, but to also leave to free interpretation critical aspects that would ensure consistent compliance across firms. As the next section shows, this has led to significant divergences in what firms are tagging and troubling discrepancies in the number of tags being generated. After presenting these findings, the question of to what extent regulatory data need to accurately represent financial practices is addressed, with reasons presented as to why the tagging rule might still be proving useful.

Compliance and enforcement

The vocation of compliance is paradoxical. Compliance officers are on the payroll of firms and yet are tasked with a para-regulatory function which involves translating rules into standardised procedures and monitoring their observance. As Lenglet (2012, p. 59) notes, compliance

officers have to invest considerable interpretative labour in the process, since 'even for financial actors, describing financial objects can prove challenging'. While regulatory indeterminacies are not of course unique to the HFT Act, the accuracy of algorithm tags would seem especially important given that the data supplied is used in real-time trade surveillance (indeed, this is why regulators put so much stress on defining an algorithm in their rule-making). The sheer diversity of compliance responses uncovered by this study was therefore a surprising and possibly concerning finding. Such discrepancies will take us to the question animating this article: to what extent do regulatory definitions and data need to accurately represent financial practices to be useful?

The smaller trading firms I spoke to were somewhere between nominal compliance and outright non-compliance with the rule. Acknowledging that they lacked the technical and budgetary resources to comply with the letter of the law, those charged with managing compliance in these firms were defensive about the situation. They claimed not to see the point of the rule and questioned trade surveillance's competency to make use of the data they already have access to (a sentiment I encountered repeatedly in my engagement with the industry). As a compliance manager put it: 'I think they're [trade surveillance are] just a little but clueless on where they need to look.' (Interview 29/03/2015) Elsewhere, in somewhat larger firms, lax interpretation of the rules allowed compliance officers to decide that they needed no systematic methodology for implementing the rule. This was justified on the grounds that since material changes in their algorithms were so rare ad hoc responses would be sufficient if and when the situation arises (Interview 10/04/2015).

Discrepancies in how the rule is being implemented are not however solely the result of differences in firms' compliance budgets. In other cases, in equally sophisticated high-frequency trading firms, there are substantial differences in how they determine if a 'material change' in an algorithm has taken place. For instance, one large high-frequency trading firm has devised their own automated system for determining the tipping points at which a change in an algorithm could be classed as material. They assigned differential weights to algorithms' parameters, allowing a formula to determine whether changes were material or not (Interview 09/01/2015). At another firm, prior to the introduction of the tagging requirement they already had procedures in place for tracking changes in their algorithms. Compliance issued example sheets to coders of the sort of changes which require authorisation, with the intention of building up traders' ability to make their own judgement calls about when a material change has taken place. The firm in question then just transposed this methodology to compliance with the HFT Act (Interview 29/03/2015).

The different ways in which compliance officers interpreted the rule sometimes led them to lock horns with one another. After convening in Canary Wharf, London,

> we got in a big argument what a different strategy is… And to be honest, someone said to me 'I don't think the rules say this'… Well, that's the point. If it's not clear and someone else takes an interpretation, or even if it is clear and it doesn't make sense,

then I change it. (Interview 10/04/2015)

Not only did compliance officers adopt their own idiosyncratic interpretations of the rule; at some firms, creative interpretation elided into rule revision.

Hunting signals in the noise

How problematic are these interpretative discrepancies? A worrying metric is inconsistency in the number of tags being generated across firms. For example, one firm claimed to have only assigned four tags since becoming 'compliant'. Another reported that they were issuing thousands per year. (Trade surveillance at the Deutsche Börse Group confirmed this diversity.) If one wanted to infer from the tagging figures the number of algorithms firms are 'really' using it would thus be an impossible task. In this sense, the tags fail in accurately representing firms' algorithmic trading procedures. On the other hand, to evaluate the data in this way would be to ignore the rule's motivations. The tagging device was never intended to represent trading firms' algorithms as they 'really are'. The point was instead to generate information that would simply render visible *relational* interactions between algorithmic strategies. A senior trade surveillance officer stressed that the intention was to let surveillance see

> what algo is reacting with another algo… And so, everything is not done manually because for example we can identify now if there is a connection between algorithms; how they interact with each other; we can identify the numbers of trades; we can identify how is the relation between an algorithm of firm A and an algorithm of firm B. Is it suspicious? Is it not suspicious? (Interview 21/04/2015)

The matter at hand concerns *to what extent* a regulatory representation of a financial practice needs to be accurate in order to be useful. Clearly, in a perfect world the algorithm tags would accurately 'mirror' or 'reflect' the algorithms used by firms and do so through the application of consistent and commensurate compliance procedures across firms. And yet, given the impossibility of such a scenario with the limited budgets and resources of regulatory bodies, there has to be a degree of tolerance for inconsistency and representational uncertainty. The critical question therefore concerns not if the tags accurately represent trading firms' algorithms – they don't – but rather to what degree the information can be inaccurate and yet still prove useful for regulatory surveillance and enforcement. According to trade surveillance, the tagging rule succeeds in hitting this threshold for usefulness, providing a weak but valuable signal in the noise which can help them to detect attempts at market manipulations such as 'layering' and 'spoofing'. A financial technologist explains:

> It [the algorithm tag] gives you a better footprint of the machine who's dealing with you... Because you can drill it down to a certain strategy… and ask specific questions [of traders]… like we've seen with algo number 5 this is the trading activity can you

> explain how that interacts with the systems you have in place like algo number 4, algo number 3. (Interview 05/11/2014)

Explaining abductive forms of reasoning associated with detective work presents a number of explanatory challenges, and it should also be acknowledged that this study did not gain access to front-line trade surveillance officers who would have been best placed to articulate how they are using the data on a day-to-day basis. But from what I could gather, the value of the tags lies in providing a more granular view of the market by disaggregating algorithmic strategies.

For instance, take a manipulative strategy known as 'layering'. When layering a firm places sell orders above the market bid price and then increases bids for the asset. Once the initial requested bid price is reached and the assets are sold, then the bids are cancelled allowing an illegitimate profit to be extracted by the algorithm in the millisecond time frame. Prior to the introduction of the tagging rule, what might have seemed to trade surveillance like a firm pursuing a layering strategy may in fact have just been different algorithms pursing parallel strategies which leave behind a suspicious-looking footprint. Vice versa, seemingly innocent algorithmic strategies hidden behind the anonymity of trader IDs might be cast in a very different light once the trading patterns are disaggregated by the tags. It is these *relational* dynamics that the tagging rule was introduced to render visible. As a regulator describes, 'when you have a layering algorithm… it doesn't matter how it is [the algorithm is] programmed because at the end of the day the result is still there.' (Interview 21/04/2015) The value of the tags lies in disaggregating market data in such a way that enables trade surveillance to more efficiently pursue a process of enquiry. On their own the tags do not determine whether or not a market manipulation is being intentionally pursued. But they do provide a better informational proxy from which to start a process of investigation.

Certainly, some threshold of diligent compliance with the rule needs be met for the tags to be useful for enforcement activities (the nominal or non-compliant responses of smaller trading firms are obviously a problem). But according to the regulatory staff I spoke to, the intention was to work towards enforcing compliance among smaller firms, starting with the most blatant offenders. They recognised the impossibility of enforcing absolutely commensurate forms of compliance across firms or realising a situation in which the number of tags being generated by firms would be situated within a narrow band. And yet, they claim even imperfect data on algorithms is still better than no information.

Could it be the case that these accounts are somewhat rose tinted? It is of course necessary to take account of possible biases in interviewees' responses. Regulatory staff, for instance, are not likely to admit to the rule providing data of limited use after investing so much effort in its development. The same problem would likely have held even if this study had gained access to frontline surveillance officers at German financial exchanges, whose interactions with external persons are tightly proscribed. Without data verifying the number of cases brought against firms after the rule was introduced, it is hard to know how successful the initiative has proven. Nonetheless, it is at least credible that the tags may be proving useful. The

robust compliance procedures of particular firms will deliver the most useful data for the investigative actions of trade surveillance, but given their intention to simply disaggregate market data even the weaker responses of some firms might be affording trade surveillance a better view of the market.

Cultural shifts and the performation of the market

The second way in which the tagging rule may be proving useful is in its effects on the cultures of trading firms. Sometimes conveyed with a note of melancholy, a theme that recurred repeatedly in my interviews was how the introduction of the algorithm tagging rule was affecting a move towards the 'professionalization' of the industry (Interview 29/03/2015; Interview 22/04/2015). The early days of 'cowboy' proprietary trading, where quants and programmers could trial algorithms live on the market, look numbered. In the larger firms, the cultural shift in has been more subtle but no less significant. Compliance officers recalled that prior to the introduction of the tagging rule they were not respected and were frequently brushed off by traders in their attempts to ensure that regulations are followed. Traders saw no need to interrupt their trading activity to engage in compliance exercises and saw no need to explain, or even share, the details of their algorithms with compliance officers. A result of this antagonism was a siloing of knowledge about algorithms within firms, exacerbating some of the knowledge problems discussed earlier in this article. As Lange's (this issue) account of 'organisational ignorance' shows, the black boxing of algorithms is not necessarily due to their technical complexity but is rather actively maintained by the division of labour within firms.

With the German HFT Act coming into force, however, compliance officers spoke of feeling empowered. In line with a general increase in the power of compliance officers in recent years in response to the raft of regulations introduced since the financial crisis (Effinger, 2015), the fact that new technical systems needed to be implemented in order to comply with the rule has given compliance officers increased leverage over traders. It has allowed them to force traders to translate their tacit knowledge into an explicit conceptual form. Rules issued directly by regulatory authorities were instrumental for this work. As a compliance officer comments, 'when I just started out I thought there was value in convincing the traders of the needs. After working 15 years in compliance, I stopped. Because they're not interested, they don't care.' (Interview 10/04/2015) The need to satisfy legal requirements provides a strong mandate for compliance officers to open up the black boxes within their firms. A related development has been a drive towards technical upskilling. It is no longer adequate for compliance officers in trading firms to approach their vocation through a legalistic lens; after the introduction of the algorithm tagging rule they now also need to be proficient in the technical aspects of trading. In one trading firm, this has resulted in compliance officers sitting in the firm's proprietary trading room and watching when traders modify the code of their algorithms, making traders justify to compliance why they do not consider the changes they are making to be 'material' (Interview 09/01/2015).

Compliance with the tagging rule is therefore reformatting the power relations and

knowledge hierarchies of trading firms in such a way that is addressing some of the problems which motivated its introduction: namely, a lack of knowledge in trading firms about the algorithms they are making use of. To be clear, this is not to invoke the 'strong Barnesian' notion of performativity deployed in MacKenzie and Millo's (2003) analysis of how the Black-Scholes-Merton options pricing formula brought market prices in line with its predictions. Nor is it to indulge in a tautology where the mere fact of compliance indicates that regulations have 'performed' their function (what Callon (2007, p. 324) dismisses as a 'prescriptive' model of performativity). The relevant notion here is rather Callon's concept of 'performation'. Prior to the strong notions of performativity which have since become the norm in the social studies of finance, Callon (1998, p. 23) observed with respect to market devices that the

> most interesting element is to be found in the relationship between what is to be measured and the tools used to measure it. The latter do not merely record a reality independent of themselves; they contribute powerfully to shaping, simply by measuring it, the reality that they measure.

That the algorithm tagging rule was devised principally as a way to allow trade surveillance to know the market but has also had the unintended consequence of opening up the algorithmic black boxes within firms fits well with Callon's notion.

The cultural effects of the rule therefore support the central contention of this paper: that the inaccuracies of regulatory definitions and data do not necessarily undermine their usefulness. In combination with the way in which surveillance officers are using the algorithm tags to better understand interactions between algorithms, these research findings suggest that just as the social studies of finance has revealed how the models and devices made use of by market participants do not need to be accurate to be useful neither it seems do regulatory rules and technologies.

Discussion and conclusion

This article argues that regulatory definitions and data need not accurately represent financial practices in order to be useful. The algorithm tagging rule of the German HFT Act was an especially appropriate test for this counter-intuitive claim. Insofar as the rise of algorithmic and high-frequency trading has led to scholarly introspection about the ability of regulators to know markets, regulatory technologies would seems to meet their ultimate challenge in this domain. Yet rather than interrogating the knowledgeability of automated markets through an abstract epistemological lens, the foregoing analysis places its emphasis firmly on social practices, looking at the practices which motivated the algorithm tagging rule, shaped regulators' definition of an algorithm, and are being mobilised in compliance and enforcement. Based on these findings it was then argued that despite the incommensurate compliance procedures of different firms, for two reason the data may be proving useful. First, a generous hearing was given to trade surveillance's claims that even if imperfect the tags are still granting

an improved *relational* vision of the market to them. In contrast to an all or nothing criteria for usefulness, this perspective reveals the potential for degrees of usefulness. Second, we saw that compliance with the rule is reformatting traditional knowledge hierarchies within trading firms. In allowing compliance officers to open up their firms' algorithmic black boxes, an unintended consequence of the rule has been to address some of the knowledge problems which motivated its introduction.

Beyond the focus on algorithmic and high-frequency trading, what this paper hopefully shows is that it is possible for the social studies of finance to more productively engage regulation. In particular, hitherto lacking has been research which elucidates the lacunae of regulatory rules and technologies without automatically defaulting to narratives of regulatory hubris, inadequacy, and failure. The shift in emphasis therefore has an implicitly normative, even political, dimension. It does not aim to absolve regulators of their failures, but rather to be aware that our conceptual frameworks for judging what counts as success can be subtly complicit with anti-regulatory rhetoric. By its very nature regulation is a fraught epistemological enterprise and if one goes looking for gaps in regulators' knowledge, vision, or ability to make sense of market data, then one will not have to look far. Morgan Ricks (2016, p. 179), former advisor in the US Treasury Department, is candid on the limits of regulatory tools: 'There are no magic solutions: a measure of imperfection and arbitrariness has to be tolerated.'

What of Dorn's (2012) argument that regulators' attempts to know markets only distracts them from more important work like, for example, devising better means of bank resolution? The problem here is that Dorn assumes that all the issues facing financial regulators can be reduced to a singular fix; that a whole range of regulatory foci can be losslessly dispensed with. Such a position does not seem credible. In seeming to call into question the very need for securities market regulation, the argument leans implicitly on a faith in the self-regulation capacities of markets not borne out by events of recent decades. Furthermore, Dorn's stress on simplicity is in acute tension with his recognition of the complexity of financial markets. Why should financial complexity not be met with regulatory complexity? Indeed, the theme of complexity needs to be handled with care. As with the notion's deployment following the Global Financial Crisis, complexity can quickly become an alibi for regulatory inaction and fatalism (Datz, 2013). In the social studies of finance, too, what seems to have taken root is a conflation of what Davies and McGoey (2012) call 'strategic ignorance' and Hayekian 'social ignorance'. With a focus on the work of market participants, the sense that regulators do not understand the minutiae of their practices has been translated into the normative proposition that regulators *cannot* know enough to regulate effectively. This, I have argued, is simply a permutation on the Hayekian critique of regulation – an idea dating back much further than the Global Financial Crisis but which has since morphed into a complexity-oriented argument.

There is certainly some truth to the Hayekian critique. The complexity of financial markets today should not be underestimated. Still, to draw conclusions about the necessary insufficiency of regulatory knowledge is, from a Latourian perspective, to place the insights of social theory over and above those of regulators involved day in day out with attempting to

know markets and uphold the law. I would therefore recommend further empirical research with regulators to ward against a drift towards sociological hubris. If that message is taken away then this article might succeed in recommencing MacKenzie's (2005) call to open the technological black boxes of global financial regulation, showing how sociologists can engage the reporting requirements, standards, and data repositories integral to the regulatory response to the financial crisis.

What consequences for adjacent research programmes might this article's analysis have? One vein of scholarship which might also benefit from engagement with this study is that on 'algorithmic governance' (Amoore, 2014; Saurwein *et al.*, 2015; Johns, 2016; Amoore & Piotukh, 2015). As noted earlier in this article, this study of the German HFT Act is the first that looks not at governance *by* algorithms, but governance *of* algorithms. As algorithms play an ever greater role in economic life, diffusing from high finance into the prosaic details of the 'real' economy, it seems likely that they will eventually be submitted to regulatory oversight. When addressing these new surveillance mechanisms, scholars of algorithmic governance might do worse than to reflect on the German HFT Act's tagging rule. They may also consider, when thinking about their challenges, to what extent regulatory data needs to be accurate to be useful.

Acknowledgements

This research has benefited from presentations and discussions at the workshop series Investigating High-Frequency Trading, the first of which was held at Copenhagen Business School in November 2014 and the second at the University of Konstanz in June 2015. For feedback on earlier drafts I thank Donald MacKenzie, Stephen Kemp, Theo Bourgeron, Ann-Christina Lange, Marc Lenglet and Robert Seyfert. I also thank the four anonymous reviewers for their helpful comments. I am especially grateful to André Broome for suggesting this research theme early on in the project's development.

References

AFME. (2011). European Commission Public Consultation: Review of the Markets in Financial Instruments Directive (MiFID). A Response by the Association for Financial Markets in Europe (AFME).

Aldridge, I. (2013). *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems.* 2nd ed. Hoboken, NJ: John Wiley & Sons.

AmCham EU. (2011). AmCham EU's Response to Consultation on the Review of the Markets in Financial Instruments Directive (MiFID).

Amoore, L. (2014). Security and the incalculable. *Security Dialogue,* 45(5), 423–39.

Amoore, L. & Volha, P. (2015). Life beyond big data: Governing with little analytics. *Economy and Society,* 44(3), 341-66.

Angel, J. (2014.) When Finance Meets Physics: The Impact of the Speed of Light on Financial

Markets and Their Regulation. *Financial Review,* 49(2), 271–81.

Arnoldi, J. (2016). Computer Algorithms, Market Manipulation and the Institutionalization of High Frequency Trading. *Theory, Culture & Society* 33(1), 29–52.

Arnuk, S. (2012). *Broken Markets: How High Frequency Trading and Predatory Practices on Wall Street Are Destroying Investor Confidence and Your Portfolio.* Upper Saddle River, NJ: FT Press.

Bank of England. (2015.) The Bank of England's Approach to Stress Testing the UK Banking System. Retrieved from http://www.bankofengland.co.uk/financialstability/Documents/stresstesting/2015/approach.pdf.

BBA. (2011). European Commission's Consultation on the Review of MiFID. A Response by: The British Bankers' Association.

Black, J. (2003). Enrolling Actors in Regulatory Processes: Examples from UK Financial Services Regulation. *Public Law,* Spring, 62–90.

Black, J. (2005). What Is Regulatory Innovation? In J. Black, M. Lodge, & M. Thatcher (Eds.), *Regulatory Innovation: A Comparative Analysis,* (1–15). Cheltenham, UK and Northampton, USA: Edward Elgar.

Blass, A., & Gurevich, Y. (2003.) Algorithms: A Quest for Absolute Definitions. *Bulletin of the European Association for Theoretical Computer Science.* Retrieved from http://research.microsoft.com/en-us/um/people/gurevich/Opera/164.pdf.

Brigagliano, J., Campion, K., Katz, D. & Blake, A. (2012). SEC adopts consolidated audit trail. *Journal of Investment Compliance,* 13(4), 18–24.

Bundesrat. (2013). Gesetz zur Vermeidung von Gefahren und Missbräuchen im Hochfrequenzhandel (Hochfrequenzhandelsgesetz).

Callon, M. (1998). Introduction: The Embeddedness of Economic Markets within Economic Theory. In M. Callon (Ed.) *The Laws of the Markets* (1-57). Oxford and Malden, MA: Blackwell Publishers.

Callon, M. (2007). What does it mean to say that economics is performative? In D. MacKenzie, F. Muniesa & L. Sui (Eds.), *Do Economists Make Markets? On the Performativity of Economics,* (311-57). Princeton & Oxford: Princeton University Press.

Callon, M., Yuval, M. & Muniesa F. (Eds.) (2007). *Market Devices.* Oxford: Blackwell Publishing.

CFTC, and SEC. (2010). Findings regarding the market events of May 6, 2010: Report of the staffs of the CFTC AND SEC to the Joint Advisory Committee on Emerging Regulatory Issues. Retrieved from http://www.cftc.gov/ucm/groups/public/@otherif/documents/ifdocs/staff-findings050610.pdf.

CFTC. (2015). Regulation Automated Trading. Notice for Proposed Rule Making. Washington, DC: Commodities and Futures Trading Commission. Retrieved from http://www.cftc.gov/idc/groups/public/@newsroom/documents/file/federalregister112415.pdf.

Datz, G. (2013). The narrative of complexity in the crisis of finance: Epistemological challenge

and macroprudential policy response. *New Political Economy,* 18(4), 459–79.

Davies, W. & McGoey L. (2012). Rationalities of Ignorance: On Financial Crisis and the Ambivalence of Neo-Liberal Epistemology. *Economy and Society,* 41(1), 64–83.

Dorn, N. (2012). Knowing markets: Would less be more? *Economy and Society,* 41(3), 316–34.

Effinger, A. (2015). The rise of the compliance guru—and banker ire. *Bloomberg.com.* June 25. Retrieved from http://www.bloomberg.com/news/features/2015-06-25/compliance-is-now-calling-the-shots-and-bankers-are-bristling.

Errico, L., Harutyunyan, A., Loukoianova, E., Walton, R., Korniyenko, Y., Amidžić, G., AbuShanab, H. & Shin, H. S. (2011). Mapping the Shadow Banking System through a Global Flow of Funds Analysis. IMF Working Paper WP/14/10. International Monetary Fund.

ESMA. (2014). European Securities and Markets Authority Discussion Paper MiFID II/MiFIR. 548.

European Commission. (2010). Review of the Markets in Financial Instruments Directive. Public Consultation.

European Union (2012) *European Market Intrastructure Regulation (EMIR).*

FIA. (2014). FIA Association Reply Form for the ESMA MiFID II/MiFIR Discussion Paper.

FOA. (2011). European Commission Public Consultation: Review of the Markets in Financial Instruments Directive (MiFID). A Response by the Futures and Options Association.

Funk, R. J. & Hirschman, D. (2014). Derivatives and Deregulation Financial Innovation and the Demise of Glass–Steagall. *Administrative Science Quarterly,* 59(4), 669–704.

Saurwein, F., Just, N. & Latzer, M. (2015). Governance of algorithms: Options and limitationsnull. *Info,* 17(6), 35–49.

Foresight. (2012). Foresight: The Future of Computer Trading in Financial Markets: Final Project Report. The Government Office for Science, London.

Geithner, T. F. (2014). *Stress Test: Reflections on Financial Crises.* London: Random House Business Books.

Gilad, S. (2010). It Runs in the Family: Meta-Regulation and Its Siblings. *Regulation & Governance,* 4(4), 485–506.

Goffey, A. (2008). Algorithm. In M. Fuller, *Software Studies: A Lexicon,* (15–21). Cambridge, MA: MIT Press.

Hanson, S. G., Kashyap, A. K. and Stein, J. C. (2011). A macroprudential approach to financial regulation. *The Journal of Economic Perspectives,* 25(1), 3–28.

Hayek, F. A. 1948. *Individualism and Economic Order.* University of Chicago Press.

Hessisches Ministerium. (2014). Guidelines to the Adherence to the Requirement of the Labelling of Trading Algorithms. Wiesbaden.

Huault, I., Lazega, E. & Richard, C. (2012). Introduction: The Discreet Regulator. In I. Huault & C. Richard (Eds.) *Finance: The Discreet Regulator: How Financial Activities Shape and Transform the World,* (1–16), Houndmills: Palgrave Macmillan.

IMF & FSB. (2009). The Financial Crisis and Information Gaps: Report to the G-20 Finance Ministers and Central Bank Governors. Retrieved from https://www.imf.org/external/np/g20/pdf/102909.pdf.

IOSCO. (2011). Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency. FR09/11. Retrieved from http://www.iosco.org/library/pubdocs/pdf/IOSCOPD354.pdf

Johns, F. (2016). Global governance through the pairing of list and algorithm. *Environment and Planning D: Society and Space*, 34(1), 126-149.

Joint Staff. (2015). The U.S. Treasury Market on October 15, 2014.

Kane, E. J. (1988). Interaction of Financial and Regulatory Innovation. *The American Economic Review*, 78(2), 328–34.

Karppi, T. & Crawford, K. (2016). Social Media, Financial Algorithms and the Hack Crash. *Theory, Culture & Society*, 33(1), **73-92.**

Kim, K. (2007). *Electronic and Algorithmic Trading Technology: The Complete Guide*. Burlington, MA: Elsevier.

Kirilenko, A. A. & Lo, A. W. (2013). Moore's law versus Murphy's law: Algorithmic trading and its discontents. *The Journal of Economic Perspectives*, 27(2), 51–72.

Knorr-Cetina, K. & Preda, A. (2012). *The Oxford Handbook of the Sociology of Finance*. Oxford Handbooks. Oxford: Oxford University Press.

Lange, A. (this issue). Organisational Ignorance: An Ethnography of High-Frequency Trading. *Economy and Society*.

Langley, P. (2013). Anticipating uncertainty, reviving risk? On the stress testing of finance in crisis. *Economy and Society*, 42(1), 51–73.

Lee, A. (2013). LEIs Will Hit Hurdles. *International Financial Law Review*, February.

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

Lenglet, M. (2012). Ambivalence and Ambiguity: The Interpretative Role of Compliance Officers. In I. Huault & C. Richard (Eds.) *Finance: The Discreet Regulator: How Financial Activities Shape and Transform the World*, (59–84). Houndmills: Palgrave Macmillan.

Lenglet, M. & Mol, J. (2015). A sorry state? Financial regulation in an algorithmic age. *Unpublished Manuscript*.

Lewis, M. (2014). *Flash Boys: Cracking the Money Code*. London: Allen Lane.

MacKenzie, D. (2005). Opening the black boxes of global finance. *Review of International Political Economy*, 12(4), 555–76.

MacKenzie, D. (2006a). *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.

MacKenzie, D. (2006b). Is economics performative? Option theory and the construction of derivatives markets. *Journal of the History of Economic Thought*, 28(1), 29–55.

MacKenzie, D. (2008). *Material Markets: How Economic Agents Are Constructed*. Oxford:

Oxford University Press.

MacKenzie, D. (2011). The credit crisis as a problem in the sociology of knowledge. *American Journal of Sociology,* 116(6), 1778–1841.

MacKenzie, D. & Millo, Y. (2003). Constructing a Market, Performing Theory: The Historical Sociology of a Financial Derivatives Exchange. *American Journal of Sociology*, 109(1), 107–45.

MacKenzie, D., Muniesa, F. & Siu, L. (Eds.) (2007). *Do Economists Make Markets? On the Performativity of Economics.* Princeton & Oxford: Princeton University Press.

MacKenzie, D. & Spears, T. (2014). 'The formula that killed Wall Street': The Gaussian copula and modelling practices in investment banking. *Social Studies of Science,* 44(3), 393–417.

Meyer, G. (2015). US Regulators Propose Powers to Scrutinise Algo Traders' Source Code. *Financial Times*, December 1.

Millo, Y., & MacKenzie, D. (2009). The Usefulness of Inaccurate Models: Towards an Understanding of the Emergence of Financial Risk Management. *Accounting, Organizations and Society,* 34(5), 638–53.

Narang, Rishi K. (2013). *Inside the Black Box a Simple Guide to Quantitative and High-Frequency Trading.* 2nd edition. Hoboken, New Jersey: John Wiley & Sons, Inc.

Pardo-Guerra, J. P. (2012). Financial automation, past, present, and future. In K. Knorr-Cetina and A. Preda (Eds.), *The Oxford Handbook of The Sociology of Finance* (567–86). Oxford: Oxford University Press.

Parker, C., and Nielsen, V. (2009). The Challenge of Empirical Research on Business Compliance in Regulatory Capitalism. *Annual Review of Law and Social Science,* 45, 45-70.

Patterson, S. (2012). *Dark Pools: The Rise of A. I. Trading Machines and the Looming Threat to Wall Street.* Random House.

Porter, T. (2003). Technical Collaboration and Political Conflict in the Emerging Regime for International Financial Regulation. *Review of International Political Economy,* 10(3), 520–51.

Ricks, M. (2016). *The Money Problem: Rethinking Financial Regulation.* Chicago and London: University of Chicago Press.

Rieder, B. & Röhle, T. (2012). Digital humanities: Five challenges. In D. M. Berry (Ed.), *Understanding Digital Humanities,* (67–84). Houndmills: Palgrave Macmillan.

Riles, A. (2013). Market collaboration: Finance, culture, and ethnography after neoliberalism. *American Anthropologist,* 115(4), 555–69.

Riles, Annelise. 2011. *Collateral Knowledge: Legal Reasoning in the Global Financial Markets.* Chicago and London: The University of Chicago Press.

Schneiberg, M., and Bartley, T. (2008). Organizations, Regulation, and Economic Behavior: Regulatory Dynamics and Forms from the Nineteenth to Twenty-First Century. *Annual Review of Law and Social Science,* 4(1), 31–61.

Seyfert, R. (this issue) Bugs, Predations or Manipulations? Incompatible Epistemic Regimes of High-Frequency Trading. *Economy and Society.*

Thiemann, M. & Lepourte, J. (forthcoming). Stitched on the Edge: Rule Evasion, Regulatory Networks and the Evolution of Markets. *American Journal of Sociology.*

Thompson, G. F. (2016). Time, Trading and Algorithms in Financial Sector Security. *New Political Economy,* 0(0): 1–11. doi:10.1080/13563467.2016.1183116.

US Government. (2010). *Dodd-Frank Wall Street Reform and Consumer Protection Act.*

Williams, J. W. (2009). Envisioning financial disorder: Financial surveillance and the securities industry. *Economy and Society,* 38(3), 460–91.

Williams, J. W. (2013). *Policing the Markets: Inside the Black Box of Securities Enforcement,* Abingdon: Routledge.

Nathan Coombs is a Leverhulme Early Career Research Fellow in the School of Social and Political Science, University of Edinburgh. His book, *History and Event: From Marxism to Contemporary French Theory,* was published by Edinburgh University Press in 2015. His current research focuses on knowledge technologies in financial regulation.