

# Metarules, judgment and the algorithmic future of financial regulation in the UK

## 1. Introduction

In response to calls for a more data-driven approach to regulation,<sup>1</sup> the Financial Conduct Authority (FCA) and the Bank of England are experimenting with the conversion of rulebook content into machine-readable and executable code.<sup>2</sup> These efforts focus primarily on information-gathering, but their implications reach far and are expected to affect all areas and aspects of regulation.<sup>3</sup> The view of human interpretation as an impediment to the effectiveness of regulation is a key premise of the growing appetite for the re-writing rules into code. This is evident in the statement of objectives of the regulators' flagship pilot programme of Digital Regulatory Reporting (DRR).<sup>4</sup> Chief amongst them is 'to make reporting rules and instructions less reliant on human interpretation and implementation, and so improve the quality

---

\* Author's name and affiliation; acknowledgements [omitted for reviewing purposes]

<sup>1</sup> *FCA and Bank of England announce proposals for data reforms across the UK financial sector* (Bank of England News Release 7 January 2020) <<https://www.bankofengland.co.uk/news/2020/january/fca-and-boe-announce-proposals-for-data-reforms-across-the-uk-financial-sector>> accessed 20 August 2021.

<sup>2</sup> FCA, *Data Strategy* (first published: 07/01/2020; last updated: 06/08/2020) <<https://www.fca.org.uk/publications/corporate-documents/data-strategy>> accessed 20 August 2021; Bank of England, FCA et al, *Digital Regulatory Reporting, Phase 2 Viability Assessment* 4, 8 and 17 <<https://www.fca.org.uk/publication/discussion/digital-regulatory-reporting-pilot-phase-2-viability-assessment.pdf>> accessed 20 August 21; Bank of England, *Transforming data collection from the UK financial sector* (Discussion Paper January 2020) ['Bank of England Discussion Paper'] paras 7.16 to 7.24 <<https://www.bankofengland.co.uk/-/media/boe/files/paper/2020/transforming-data-collection-from-the-uk-financialsector.pdf?la=en&hash=6E6132B4F7AF681CCB425B0171B4CF43D82E7779>>.

<sup>3</sup> FCA, *Digital Regulatory Reporting* (first published 1/11/2017; last updated 14/10/2020) <<https://www.fca.org.uk/innovation/regtech/digital-regulatory-reporting>> accessed 20 August 2021; Bank of England, *Transforming data collection from the UK financial sector: A plan for 2021 and beyond* (23 February 2021) ['Bank of England transforming data collection plan'] <https://www.bankofengland.co.uk/paper/2021/transforming-data-collection-from-the-uk-financial-sector-a-plan-for-2021-and-beyond> accessed 20 August 2021.

<sup>4</sup> FCA (n 3).

of regulatory data.<sup>5</sup> The view of human interpretation as an hurdle is also highlighted in a recent Bank of England Discussion Paper which notes how vague regulatory instructions ‘can lead to “pain points” for firms in interpreting instructions’, causing delays and quality issues for the Bank.<sup>6</sup>

While experimentation is still on-going, two schools of thought have emerged as regards the future place of human interpretation in regulation: On the one hand, those in favour of full automation and of ‘taking humans out from large part of the solution development and interpretation phase’ and, on the other hand, those who take a more cautious approach.<sup>7</sup> In this Article, I do not aim to argue against digitalisation, but to explain why, at least from the legal point of view, we have good reasons to ensure that human interpretation remains an indispensable component of data-driven governance.<sup>8</sup> Accordingly, the real challenge is not to find ways to eradicate the process of interpretation, but to design rulebooks which will help their human users take advantage of their own general intelligence and of the specialist intelligence of machines as they go about ascribing meaning to rulebook content against an evolving ecology of commercial practice.<sup>9</sup> I offer three arguments in support of this claim. The first concerns the limited translatability of regulatory content into algorithms. The second draws attention to the finite capabilities of machines in making determinations about the kind of action that is required (eg, with regards to what sort of data needs to

---

<sup>5</sup> FCA, *Pilot Phase 1 Report* (June to December 2018) [‘Phase 1 Report’] 7 <<https://www.fca.org.uk/publication/discussion/digital-regulatory-reporting-pilot-phase-1-report.pdf>> accessed 20 August 2021; See further Bank of England, FCA et al (n 2) 4 and 6 (identifying interpretive problems as major drivers of the high cost of industry compliance in the UK).

<sup>6</sup> Bank of England Discussion Paper (n 2) 2.

<sup>7</sup> Bank of England transforming data collection plan (n 3) para 4.2.

<sup>8</sup> Bank of England Discussion Paper (n 2) para 7.16 (conceding albeit tentatively that interpretation will still be required).

<sup>9</sup> Here I am following Margaret Boden who describes the various types of Artificial Intelligence (AI) as ‘specialist systems’ of intelligence to distinguish them from the ‘general intelligence’ of humans in order to explain why it is so difficult to engineer human level artificial general intelligence. MA Boden, *Artificial Intelligence, A Very Short Introduction* (OUP 2018) 18.

be reported) given the existing and foreseeable development of the relevant technology. The third brings attention to the participatory, deliberative and constructive character of human interpretation as a process, and contends that we would have reasons to preserve it even if it were possible to overcome the limitations discussed under the first two arguments.<sup>10</sup>

I develop my thesis on the following assumptions. Machine-readable and executable regulation consists of *metarules*, namely authoritative micro-directives which are expressed in algorithmic language and specify concrete courses of action (or omission, but I will leave this aside) for rulebook users, while enabling the execution of at least some aspects of that action by machines.<sup>11</sup> I use the prefix ‘meta’ to convey the supervenience of code on regulatory content. I use the term ‘rule’ to mark the normative character of that content. I further assume that metarules are the outputs of algorithmic decision-making which bears the following two features: (1) functional autonomy in the performance of certain tasks (eg, retrieval of specific data); and (2) quasi-decisional autonomy, ie, reliance on machine-learning for the processing of inputs and the determination of outputs in a manner partially independent of human designers and operators so that it remains compatible with the requirements of the Data Protection Act 2018 (DPA 2018) where relevant.<sup>12</sup> Such outputs may be neither entirely predictable nor susceptible to reasoned explanation

---

<sup>10</sup> [Omitted for reviewing purposes].

<sup>11</sup> AJ Casey and A Niblett, ‘Focus Feature: Artificial Intelligence, Big Data, and the Future of Law - Self-Driving Laws’ (2016) 66(4) *University of Toronto Law Journal* 429 (using the term ‘micro-directive’).

<sup>12</sup> Whilst the Data Protection Act 2018 (DPA 2018) applies to personal data, the legislator was clearly concerned with automated decisions without human control. Of particular relevance here are sections 49(1) and 50. While the former stipulates that ‘a controller [of personal data] may not take a significant decision based solely on automated processing unless that decision is required or authorised by law’, the latter adds that a data subject can request an automated decision to be reconsidered or a new decision to be made ‘not based solely on automated processing’. See further K. Yeung, ‘Why worry about decision-making by machines?’ in M Lodge and K Yeung (eds) *Algorithmic Regulation* (OUP 2019) 21, 22-23.

and justification, though this is consistent with humans retaining the capacity to intervene, eg, to validate machine outputs, and the formal discretion to accept or reject them.

Implicit in the construction of human interpretation as a series of “pain points” for rulebook users, is a deeper concern with the endemic problem of legal uncertainty in regulatory law and finding ways to address it. In view of this, in Section Two, I provide a brief account of the engineering of metarules as a response to the problem of legal uncertainty leaving the discussion of other potential benefits of digitalisation and automation aside for another occasion. I present my three arguments in Sections Three to Five. In Section Six, I conclude with a set of principles for the future design of the financial rulebooks which are animated by the idea that *no stakeholder should become worse off* as a result of the use of metarules and automation.

In the interests of clarity and scope, I will focus on metarules derived from the binding content of the rulebooks of the FCA and the PRA namely regulatory provisions earmarked here (as in the rulebooks) as ‘rules’ so that they are differentiated from non-binding guidance. My examples draw primarily on relatively detailed rules because they explicate the content of the high-level principles of the rulebooks and, as a result, they provide the natural starting point for exploring the conversion of rulebook content into metarules.<sup>13</sup> The statutory objectives of the two regulators alongside the regulatory principles and the threshold conditions, uncodified common law principles, commercial practices and customs are also of relevance to the interpretation of rulebook content and, add an extra layer of complexity, but their

---

<sup>13</sup> This is further discussed below in section 3 below.

examination falls beyond the scope of my inquiry, and of most accounts of algorithmic regulation.

The impact of technology on the use of rules as instruments of social organisation and control has received growing attention in recent legal scholarship. By way of example, Aaron Wright and Primavera De Filippi have coined the term ‘lex cryptographia’ to describe the eventual rise of “rules administered through self-executing smart contracts”.<sup>14</sup> In their turn, Anthony Casey and Anthony Niblett have declared the future ‘death of rules and standards’ thanks to machines translating complex legislative goals into ‘a vast catalog of simple commands for all possible scenarios’.<sup>15</sup> The intersection of artificial intelligence, technology and the law in financial markets has also been researched extensively.<sup>16</sup> The implications of re-writing rules into code for the future of human interpretation in a data-driven governance has attracted less attention, despite the salience of the process of interpretation in regulatory law. My thesis seeks to address this gap in the literature.

---

<sup>14</sup> A Wright and PP De Filippi, ‘Decentralized Blockchain Technology and the Rise of *Lex Cryptographia*’ (March 10, 2015) <<https://ssrn.com/abstract=2580664>> accessed 20 August 2021.

<sup>15</sup> AJ Casey and A Niblett, ‘The Death of Rules and Standards’ (2017) 92(4) *Indiana Law Journal* 1401.

<sup>16</sup> The scholarship focuses primarily on the development and legal implications of Financial Technology (FinTech), technology governance, and competition law issues associated with sandboxes for FinTech experimentation. See E Micheler and A Whaley, ‘Regulatory Technology: Replacing law with computer code’ (2020) 21(2) *European Business Organisation Law Review* 349; S Omarova, ‘Technology v. Technocracy: Fintech as a Regulatory Challenge’ (2020) 6 *Journal of Financial Regulation* 75; E Avgouleas and A Kiayias, ‘The Promise of Blockchain Technology for Global Securities and Derivatives Markets: The New Financial Ecosystem and the ‘Holy Grail’ of Systemic Risk Containment’ (2019) 20(1) *European Business Organisation Law Review* 81; R Van Loo, ‘Making Innovation More Competitive: The Case of Fintech’ (2018) 65 *UCLA Law Review* 232; W-G Ringe and C Ruof, ‘A regulatory sandbox for robo advice’ EBI Working Paper Series 2018 -no. 26 [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3188828](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3188828) accessed 17 June 2022; DW Arner, J Barberis & RP Buckley, ‘FinTech, RegTech, and the Reconceptualization of Financial Regulation’, (2017) 37(3) *North-western Journal of International Law and Business* 371. In parallel to this literature, a more theoretical discourse examines the advent of algorithmic regulation, and the impact of technology on legal concepts and doctrines. See R Brownsword, E Scotford and K Yeung (eds), *The Oxford Handbook of Law, Regulation and Technology* (OUP 2017); M Lodge and K Yeung (eds), *Algorithmic Regulation* (OUP 2019); M Hildebrandt, ‘Law as information in the era of data-driven agency’ (2016) 79(1) *MLR* 1.

## 2. The engineering of metarules as an answer to the problem of legal uncertainty

Financial regulators deploy a variety of rules to communicate their commands, expectations and guidance.<sup>17</sup> In terms of their linguistic structure, which also corresponds with what Frederick Schauer describes as ‘the standard picture of the distinction between rules and standards’, these different types of rules can be seen as making up a spectrum of options, with highly specific rules (rules) standing at one end, and vaguely or broadly phrased rules (standards) standing at the other.<sup>18</sup> Broadly stated rules tend to be durable, and flexible and they allow greater discretion in their interpretation. Detailed rules tend to provide greater certainty, clarity and predictability. Furthermore, their use seems to be more appropriate where the relationship between regulators and regulatees is not one of mutuality and trust. Neither economic analysis nor behavioural studies can provide a definite answer on the choice of legal form.<sup>19</sup> However, they both corroborate to the view that rules are

---

<sup>17</sup> Financial Services and Markets Act 2000 (FSMA 2000), part 9A (providing the legal basis of the rulemaking powers of the FCA and the PRA. Classifications of different types of rules abound in the literature. See notably, F Schauer, *Playing by the Rules: A Philosophical Examination of Rule-Based Decision Making in Law and in Life* (Clarendon Press 1991; reprint 2002) 1-16; R Baldwin, M Cave and M Lodge, *Understanding Regulation, Theory, Strategy and Practice* (OUP 2012) 230 and 296-311; and A Ogus, *Regulation, Legal Form and Economic Theory* (Hart Publishing 1994) 150-179 and J Black, *Rules and Regulators* (Clarendon Press 1997) 7-10.

<sup>18</sup> F Schauer, ‘The convergence of rules and standards’ (2003) 3 *New Zealand Law Review* 303, 307 with further reference to literature demonstrating the diversity of approaches taken to draw the distinction between rules and standards.

<sup>19</sup> RB Korobkin, ‘Behavioral analysis and legal form: Rules v standards revised’ (2000) 79 *Oregon Law Review* 23, 58; and K M Clemont, ‘Rules, standards and such’ (2020) 68 *Buffalo Law Review* 751, 760. For a classic economic analysis of legal rule-making, see I Ehrlich and RA Posner ‘An economic analysis of legal rule-making’ (1974) *Journal of Legal Studies* 257, 286 (identifying four types of costs: the cost of rule-making, the cost of enforcement, the cost of compliance, which is defined as the cost imposed on the industry, and social costs understood as the costs imposed by regulatory offences). On the economic analysis of law-making and the ‘rules versus standards’ debate, see further L Kaplow, ‘Rules versus standards: An economic analysis’ (1992) 42 *Duke Law Journal* 557; F Schauer, ‘The tyranny of choice and the rulification of standards’ (2005) 14 *Contemporary Legal Issues* 803; CS Diver, ‘The optimal precision of administrative rules’ (1983) 93 *Yale Law Journal* 65, 73-74. On the contribution of behavioural economics into the traditional economic analysis of law, see notably C Jolls, C Sunstein and R Thaler, ‘A behavioural approach to law and economics’ (1998) 50 *Stanford Law Review* 1471, 1474 and 1545; and E Zamir and D Teichman, *Behavioral Law and Economics* (OUP 2018) 1. For a classic criticism, see RA Posner, *Rational Choice, Behavioral Economics, and the Law* (1998) 50 *Stanford Law Review* 1551.

more expensive than standards for regulators to make, while standards are usually more expensive for regulatees to apply compared to rules.<sup>20</sup>

Where the law is not sufficiently clear on a particular matter, regulatees need to spend time and resources to figure out whether their behaviour complies with the law.<sup>21</sup> Financial regulators try to reduce these costs as much as possible pursuant to the regulatory principles of the Financial Services and Markets Act 2000 (FSMA 2000) and specifically that of proportionality, hence, internalising part of the cost of legal uncertainty.<sup>22</sup> When the regulated behaviour is frequent and homogeneous, (nearly) full information is assumed to be available ex ante and, as a result, detailed rules are not prohibitively costly to make as means to ameliorate legal uncertainty.<sup>23</sup> When the regulated behaviour is infrequent and heterogeneous, full information is not available at the point of rulemaking and vague rules (standards) are promulgated instead as a more affordable option when seeing from the perspective of rulemakers.<sup>24</sup>

Recent technological developments have the potential of cutting the costs of rulemaking significantly for the benefit of regulators and regulatees alike.<sup>25</sup> Specifically, predictive, data storage, and communication technology promises to improve the ability of financial regulators to collect, and store information, make projections, and design finely calibrated rules in algorithmic language, update the

---

<sup>20</sup> Korobkin (n 19) 56.

<sup>21</sup> Casey and Niblett (n 15) 1407.

<sup>22</sup> FSMA 2000, section 3B(1)(b).

<sup>23</sup> Casey and Niblett (n 15). H-B Schaefer, *Legal Rules and Standards* in CK Rowley and F Schneider (eds) 2 *The Encyclopedia of Public Choice* 347, 347-348; and Korobkin (n 19) 46 and 56 (arguing that the cost of using standards is higher because more cases will be litigated, and that self-serving bias increases the likelihood of litigation).

<sup>24</sup> Casey and Niblett (n 15) 1408.

<sup>25</sup> Micheler and Whaley (n 16) 364-365 (noting that the cost of switching to the relevant technology is enormous but once done there will be massive savings collectively).

content of those rules and communicate them in (almost) real time.<sup>26</sup> In short, it pledges to unlock the mutation of the existing financial rulebooks into a vast catalogue-like rulebook of metarules the purpose of which would be to provide constantly updated context-specific granular instructions to rulebook users.<sup>27</sup>

To see how this new type of financial rulebooks might work, consider the use of traffic lights for the regulation of the flow of traffic:<sup>28</sup> In a world without traffic lights, drivers would have to consult times, tables that matched intersections and directives with prescribed intervals of stopping. With traffic lights, all these complexities are translated into a simple instruction: a red light or a green light depending on whether drivers are required to stop or to continue driving. This is the simplest function of traffic lights, but more sophisticated functions are also possible. For example, traffic lights adjusting the duration of intervals or giving priority to an ambulance in an emergency thanks to sensors which detect and predict the flow of the traffic in real-time.

In a similar way that traffic lights produce red and green signals to regulate traffic in public roads, it is at least conceivable that machines could be trained to produce metarules for the regulation of financial markets. Powered by advanced data

---

<sup>26</sup> James Proudman, ‘Supervisor-centred automation – the role of human-centred automation in judgment-centred prudential supervision’ (27 April 2020) speech given at the Bank of England, CEPR and Imperial College Conference on the ‘Impact of AI and Machine Learning on the UK Economy’ (discussing how machine learning and other artificial intelligence is transforming the operational capabilities of the PRA in relation to financial supervision) <<https://www.bankofengland.co.uk/speech/2020/james-proudman-supervisor-centred-automation-speech>> accessed 20 August 2021; Mark Carney, ‘New economy, new finance, new Bank’ (21 June 2018), speech on the on-going work of the Bank to explore how new technologies could improve the compliance of the regulated industry and streamline regulatory processes while enhancing the ability of the Bank of England to analyse data. <<https://www.bankofengland.co.uk/speech/2018/mark-carney-speech-at-the-lord-mayors-bankers-and-merchants-dinner-mansion-house>> accessed 20 August 2021.

<sup>27</sup> Casey and Niblett (n 15).

<sup>28</sup> *ibid* 1416-1417.



and predictive analytics, machines would receive data input, identify the relevant rulebook provision, create a metarule that is a simple tailor-made instruction (eg, in relation to a reporting requirement), and then communicate that instruction in real-time or even execute the action that is required automatically (eg, retrieval and submission of specific bit of data in compliance with a reporting requirement) on behalf of human rulebook users.<sup>29</sup> Subject to further improvements, the same technology could be also deployed for the automatic detection of violations of regulatory law and even the automation of enforcement:<sup>30</sup> Administrative fines could become immediately payable, or human operators may be automatically restrained from taking any further action.<sup>31</sup> A machine could also be programmed to produce a compliance score as a metarule to warn, for instance, a mortgage advisor that a particular recommendation would be in breach of the suitability requirements of the FCA rulebook and, if required, even proceed to log them out from their desktop office computer automatically, so that they are unable to proceed to the completion of the transaction.<sup>32</sup>

From the point of view of computer engineering, we have technology in place with the potential of supporting the massive production and execution of metarules.<sup>33</sup>

Machine learning is a type of artificial intelligence which enables machines not just to

---

<sup>29</sup> Casey and Niblett (n 15) 1411-1412.

<sup>30</sup> *ibid* 1404.

<sup>31</sup> In a similar way that a 'smart' red-light camera can impose a fine by deducting the relevant amount of money from the bank account of a car driver. On the constitutional implications of the automation of enforcement through the use of red-light cameras in the US, see JO Christensen, 'Note, Wrong on Red: The constitutional case against red-light cameras (2010) 32 Washington University Journal of Law and Policy 443, 446 (arguing among other things that the use of these cameras violates the defendant's right of due process).

<sup>32</sup> Suitability requirements mainly apply to the provision of financial advice, investment management and the management of assets of certain pension schemes. FCA Handbook of Rules and Guidance, Conduct of Business Sourcebook (COBS), chapter 9 <<https://www.handbook.fca.org.uk/handbook/COBS/9/?view=chapter>> accessed 12 August 2021.

<sup>33</sup> A notable example is the ISDA Common Domain Model (CDM) -a logical model to express reporting rules for EMIR and MiFIR- which was tested in the DRR Pilot. <<https://www.isda.org/2019/10/14/isda-common-domain-model/>> accessed 12 August 2021.

do specific tasks but also learn (eg, to retrieve, transmit, submit or update a specific bit of information) without being explicitly programmed.<sup>34</sup> A key advantage of this technology is that it allows real-time analysis of vast volumes of information for the identification of unusual correlations, patterns, and emerging risks, and for predictions about the future. Blockchain and other types of Distributed Ledger Technology, quantum computing, the Internet of Things as well as the convergence of Big Data and Big Compute are expected to increase data storage, access and processing all of which are crucial for the further advancement of the analytical capabilities of computers with machine learning software.<sup>35</sup> The combination of Machine Learning with Natural Language Processing or other types of semantic technology could also enable machines to read and process legal content for the execution of reporting and other tasks.<sup>36</sup> An advanced type of hybrid Machine Learning, which is of particular interest here, is that of Deep Learning. One of the intriguing features of this technology is that it uses the so-called ‘thought vectors’ to deconstruct language with almost mathematic precision and to translate and simulate the usage of natural

---

<sup>34</sup> For a classic definition of machine learning, see T Mitchell, *Machine Learning* (New York: McGraw-Hill Education; 1997) 2. JD Lohr, WJ Maxwell and P Watts, ‘Legal practitioners’ approach to regulating AI risks’ in M Lodge and K Yeung (eds) *Algorithmic Regulation* (OUP 2019) 224, 225 (providing examples of the multiple applications of Machine Learning).

<sup>35</sup> The National Cyber Security Centre defines Distributed Ledger Technology as ‘a data storage mechanism in which data is stored at multiple locations on a shared network.’ It can be permissioned or permissionless, depending on the presence or absence of a central authority which may be private or public. National Cyber Security Centre, *Distributed Ledger Technology - The nature and applications of Distributed Ledger Technology* (White Paper 20 April 2021) <[https://www.ncsc.gov.uk/whitepaper/distributed-ledger-technology#section\\_2](https://www.ncsc.gov.uk/whitepaper/distributed-ledger-technology#section_2)> accessed 20 August 2021. Big Data describes an extremely vast set of accessible data (eg, the Internet of Things). Big Compute refers to a wide range of tools and approaches to run large-scale applications for business, science, and engineering performing complex modelling, simulations etc. Cloud computing is an example of Big Compute. The Internet of Things connects machines and enables them to transfer data over the network autonomously that is without human to human or human to machine interaction. See further ‘FinTech jargon buster’ *Practical Law UK Practice Note w-010-983*; and DH Wolpert and WG Macready, ‘No Free Lunch Theorems for Optimization’ (1997) 1(1) *IEEE Transactions on Evolutionary Computation* 67-82 (for a sceptical account on the allegedly positive impact of Big Data on Machine Learning).

<sup>36</sup> XBRL is an example of semantic technology which provides a framework for exchanging business information. <<https://en.wikipedia.org/wiki/XBRL>> and <<https://en.wikipedia.org/wiki/XML>> accessed 12 February 2020. For a critical examination of current initiatives to expand the use of AI and other technologies in the legal profession in the US, see F Pasquale, ‘A rule of persons, not machines: The limits of legal automation’ (2019) 87 *George Washington Law Review* 1, 17-43.

language.<sup>37</sup> While these types of technology evolve, the growth of machine readable and executable financial regulation is incremental and largely a work in progress which nevertheless already calls into question the case for retaining human interpretation as an aspect of the emerging algorithmic governance.<sup>38</sup> In the next three sections, I offer three grounds for its preservation.

### 3. The limited translatability of regulatory content into algorithms

The financial rulebooks consist of a combination of high-level principles, detailed rules and non-binding guidance.<sup>39</sup> High-level principles take the form of the FCA ‘Principles for Business’ (PRIN) and the PRA ‘Fundamental Rules’ (FR) which replicate eight out of the eleven FCA Principles for Business.<sup>40</sup> The high-level principles are ‘standards’ in terms of their function. They encapsulate benchmarks against which regulatees are to be assessed in terms of their professional conduct and their financial health and soundness.<sup>41</sup> The remaining myriad of rules and guidance are more detailed statements of the high-level principles of the rulebooks and are to

<sup>37</sup> Y LeCun, Y. Bengio and G Hinton, ‘Deep Learning’ (2015) 521 *Nature* 436.

<sup>38</sup> For an overview of notable examples, see Bank of England transformation data collection plan (n 3) para 5.3; ISDA, ‘CDM Factsheet’ (2019) <https://www.isda.org/a/z8AEE/ISDA-CDM-Factsheet.pdf> accessed 20 August 2021.

<sup>39</sup> Traditionally, debates were cast as a choice between rules and standards, however, more recent literature in the field has moved beyond this rigid binary choice. On this point, see C Ford, *Innovation and the State: Finance, Regulation and Justice* (CUP 2017) 12; and Schauer (n 18).

<sup>40</sup> The letter ‘R’ in ‘COBS 2.1.1 R (1)’ earmarks COBS 2.1.1 as a ‘rule’. The word ‘Rules’ in the title ‘Fundamental Rules’ of the relevant part of the PRA rulebook serves the same purpose. The legal status of the PRIN and the FR as ‘rules’ is also evident in the FCA and the PRA Glossaries both of which refer to section 417(1) of the FSMA 2000 as the latter defines the term ‘rule’ broadly to include principles. <<https://www.handbook.fca.org.uk/handbook/glossary/G1036.html?starts-with=R>>; and <<https://www.prarulebook.co.uk/rulebook/Content/Chapter/211144/07-06-2022>> accessed 1 July 2022.

<sup>41</sup> Examples of high-level principles applicable to the regulatees of both the PRA and the FCA include the principle of integrity, the principle of due care skill and diligence, and the principle of financial prudence. See respectively PRIN 2.1.1 R <<https://www.handbook.fca.org.uk/handbook/PRIN/2/?view=chapter>>; and FR 2.1 <<https://www.prarulebook.co.uk/rulebook/Content/Chapter/211141/30-11-2015>> accessed 1 July 2022.

be interpreted in the light of those principles.<sup>42</sup> The high-level principles allow the exercise of discretion, amplify and reinforce the meaning of the more detailed provisions of the rulebooks. While the prominence of high-level principles in the rulebooks testifies to the survival of elements of a principles-based approach to regulation, their earmarking of PRIN and FR as ‘rules’ in the rulebooks suggests two things: The abandonment of the rhetoric of ‘Principles Based Regulation’ and that the choice of regulatory approach and the choice of rule type are not necessarily aligned.<sup>43</sup>

The creation of machine readable and machine executable regulation requires coding.<sup>44</sup> If we want a machine to do something for us, we need to give it an algorithm, ie, a single list of rules presented in the right order for the machine to follow. A set of algorithms make up a code, while a system of codes makes up a computer software, for instance, a computer software that supports machine learning. Broadly speaking, software developers have three options when they code legal text: They can code the text themselves or develop an algorithm that trains machines to do

---

<sup>42</sup> *British Bankers Association, R (on the application of) v Financial Services Authority and another* [2011] EWHC 999 (Admin) para [170]. For a more recent application, see *Berkeley Burke SIPP Administration Ltd v Financial Ombudsman Service Ltd* [2018] EWHC 2878 (Admin).

<sup>43</sup> J Black, ‘Paradoxes and Failures: ‘New Governance Techniques and the Financial Crisis’ (2012) 75(6) *Modern Law Review* 1037, 1042-1044 (drawing a distinction between ‘rulebook’ and ‘operational’ principles-based regulation); and J Black, ‘The Rise, Fall and Fate of Principles Based Regulation’ (LSE Law, Society and Economy Working Papers 17/2010) <<http://ssrn.com/abstract=1712862>> accessed 25 July 2022. See further [omitted for reviewing purposes] (on the reincarnation of principles-based regulation as ‘judgment-led’).

<sup>44</sup> Software development involves coding and computer programming. Coding is the process of translating and writing codes from one language to another, while programming is the process of building an executable programme that can be used to carry out machine level outputs. <[https://en.wikipedia.org/wiki/Computer\\_programming](https://en.wikipedia.org/wiki/Computer_programming)> accessed 12 February 2020. See further JA Kroll, J Huey, S Barocas, EW Felten, JR Reidenberg, DG Robinson & H Yu, ‘Accountable Algorithms’ (2017) 165(3) *University of Pennsylvania Law Review* 633, 646 (describing ‘software code’ as ‘a rigid and exact description of itself: the code both describes and causes the computer’s behavior when it runs.’).

the coding and produce outputs accordingly or opt for a combination of both.<sup>45</sup> A system of algorithmic financial regulation could make use of these options and produce algorithmic micro-directives of at least two kinds: First, metarules for the execution of tasks like the automatic submission of data for reporting purposes.<sup>46</sup> Second, metarules communicating compliance scores on the basis of a statistical model to enable humans pre-test a step in order to see if it complies or not while retaining the formal discretion to accept or reject them.<sup>47</sup>

The crafting of metarules presupposes the translatability of rulebook content into its algorithmic equivalent but the conversion of legal rules into code is an extremely challenging task.<sup>48</sup> The root cause of these difficulties can be traced back to the kind of intelligence that machines are equipped with. Compared to humans, machines come with specialised intelligence. They exceed human capacity at specific tasks, but their focus is narrow and domain-specific and, therefore, of limited transferability across domains. Machines are pre-programmed to deliver a specific goal, namely the one that is encoded in their software. Accordingly, their ‘smartness’

---

<sup>45</sup> M Zalnieriute, L Bennett Moses and G Williams, ‘The Rule of Law and Automation of Government Decision-Making’ (2019) 82(3) MLR 425, 432-433.

<sup>46</sup> Currently, regulators make rules to instruct regulatees to submit data for reporting purposes and expect them to comply voluntarily or under the threat of enforcement. Regulatory technology bears the potential of replacing the existing reporting arrangement, which is known as the ‘push model’, with a ‘pull model’ of reporting in which regulators pull data themselves. For present purposes, I am assuming that the ‘push model’ of reporting remains in place namely one in which regulatees have the legal obligation to submit information in compliance with regulatory law. On the distinction between different models of reporting, see Bank of England Discussion Paper (n 2) 42-45 <https://www.bankofengland.co.uk/-/media/boe/files/paper/2020/transforming-data-collection-from-the-uk-financial-sector.pdf?la=en&hash=6E6132B4F7AF681CCB425B0171B4CF43D82E7779> accessed 1 July 2022.

<sup>47</sup> For a discussion on the use of technology in the private sector for the prediction legal outcomes, see notably DM Katz, ‘Quantitative legal prediction -or how I learnt to stop worrying and start preparing for the data-driven future of the legal services industry’ (2013) Emory Law Journal 909, 914-915; and A Porat and LJ Strahilevitz, ‘Personalising default rules and disclosure with Big Data’ (2014) 112 Michigan Law Review 1417, 1436. The use of this technology by the industry may create problems where the incentives of the industry do not align with the regulators’ incentives. On this point, see Micheler and Whaley (n 16) 363 and 366.

<sup>48</sup> FCA Pilot Phase 1 Report (n 5) 10 and 14 (for a discussion of relevant difficulties in the context of digital reporting).

is a function of how effective they are in producing outputs that attain the goal in question (eg, submission of specific data).

The more access they have to data, the more capable the machines become in performing pre-programmed tasks.<sup>49</sup> Nevertheless, data accessibility is not enough. Machines do not process equally well all types of data.<sup>50</sup> To perform well, they need to be fed with highly structured data, ie, standardised bits of information with an exhaustively defined meaning. This is not to say that machines cannot cope at all with more open-ended and less-clearly pre-defined data. They do, but the less structured the data, the more their capabilities diminish. The machines' reliance on data of the highly structured type sets an important obstacle to the conversion of rulebook content into code. It requires software developers to, first, break rulebook content down into granular instructions, and then convert those instructions into algorithmic language so that machines can read them and perform certain tasks. The trouble is that, even when it is feasible to reach the requisite degree of granularity, there is always the risk of loss of meaning.

Consider, for instance, SUP 16.11.5.R of the FCA Handbook which like most of the rulebook provisions lacks exhaustive precision. SUP 16.11.5.R specifies that 'A sales data report must contain sales data in respect of the following products: (1) retail investments; (2) pure protection contracts; (3) regulated mortgage contracts (but not further advances); (4) home purchase plans; (5) home reversion plans; (6) regulated

---

<sup>49</sup> Lohr et al (n 34) 231 (noting that '[d]ata is the raw material of machine learning'); and Wolpert and Macready (n 35) 67-82 (taking a more sceptical stance and noting amongst other things that mining a larger pool of historical data will not make machine outputs correct or meaningful).

<sup>50</sup> Bank of England and Financial Conduct Authority, 'Machine Learning and the UK financial services' (October 2019) 21-23 <<https://www.bankofengland.co.uk/report/2019/machine-learning-in-uk-financial-services>> accessed 21 August 2021 (drawing a distinction between three different types of data: 'structured', 'semi-structured' and 'unstructured-data').

sale and rent back agreements; (7) high-cost-term credit and; (8) home credit loan agreements.’ Even though it is possible to rephrase the content of this legal rule into more granular attributes such as ‘sale’, ‘data’, ‘products’, ‘retail’, ‘investments’, each one of those attributes also needs to be translated into more precise metadata to enable algorithmic conversion.

Translating SUP 16.11.5.R into more granular instructions for coding purposes is not as straightforward as it seems. Perhaps, the best way to illustrate this point is by considering HLA Hart’s well-known example of a rule according to which ‘no vehicle may be taken in the park’.<sup>51</sup> Suppose that software developers are training a machine to identify a vehicle and that, for the purposes of that training they define ‘vehicle’ to mean ‘passenger cars’. This still leaves unclear whether the definition of the word ‘vehicle’ should include a truck, a wheelchair, or a pram. No account of the meaning of the word ‘vehicle’ can include everything that is a vehicle and exclude everything that is not a vehicle. The relationship between the various uses of the word ‘vehicle’ is, as Ludwig Wittgenstein famously noted, like the relationship between members of a family.<sup>52</sup> A resemblance exists, but it is not possible to give this resemblance any rigid definition. Accordingly, the algorithmic conversion of legal rules of relative precision is much more complicated than –say, the identification and submission of the reference number of a specific product provider.

Compared to SUP 16.11.5.R, certain more technically detailed reporting requirements of the existing financial rulebooks seem to be more susceptible to

---

<sup>51</sup> Hart uses this example to explore the ‘open texture’ of the law. HLA Hart, *The Concept of Law* (OUP 1994) 128-129.

<sup>52</sup> GEM Anscombe and R Rhees (eds), *Philosophical Investigations* (Blackwell 1953; GEM Anscombe translation) para 65-66 (on family resemblance). M McGinn, *Routledge Philosophy Guidebook to Wittgenstein and the Philosophical Investigations* (Routledge 2003) 33-72.

algorithmic conversion from the point of view of computer engineering. Take, for example, SUP 16, Annex 21R of the FCA Handbook of Rules and Guidance. This is a legal rule on the content of the sales data report in relation to mortgage reporting requirements known as ‘PSD001’.<sup>53</sup> SUP 16, Annex 21R specifies in extreme detail the data reporting fields that must be completed and further guidance.<sup>54</sup> For instance, one data field concerns the reference number of the product provider, which comes with a six-digit code. Another concerns the reference number of the product sold, which also comes with a six-digit code. Finally, a further data field is about the provision of financial advice at the point of sale which comes with code ‘Y= advised’ or ‘N= not advised’. A technical legal rule of the type of SUP 16, Annex 21 R of the FCA Handbook is most probably so exhaustive of meaning that its conversion into code is the least challenging. That being said, even the most technically detailed rules are meant to be read in conjunction with rules of high or medium linguistic vagueness to ensure that they are used properly.<sup>55</sup>

The fact that the meaning of legal rules is in varying degrees context-dependent points to a further difficulty.<sup>56</sup> Even when it is possible to convert the semantic content of legal rules into algorithmic language, it is not possible to capture the perpetually changing context within which these rules are meant to apply.<sup>57</sup>

---

<sup>53</sup> Home finance lenders and administrators submit relevant reports to the FCA, but the FCA shares the relevant data with the PRA. For an overview of recent changes on mortgage reporting requirements, see *FCA and PRA changes to mortgage reporting requirements* (Policy Statement FCA-PS19/23 PRA-PS22/19 September 2019) <<https://www.fca.org.uk/publication/policy/ps19-23.pdf>> accessed 12 February 2020.

<sup>54</sup> According to SUP 16.11.7 R, the submission of PSD001 data reports must further meet the specifications of SUP 16 Annex 21 R <[https://www.handbook.fca.org.uk/form/sup/SUP\\_16\\_ann\\_21\\_20191028.pdf](https://www.handbook.fca.org.uk/form/sup/SUP_16_ann_21_20191028.pdf)> accessed 20 August 2021.

<sup>55</sup> SUP 16.11.7 R.

<sup>56</sup> Schauer (n 17) 142.

<sup>57</sup> Micheler and Whaley (n 16) 354 (noting in relation to Natural Language Processing that the relevant technology is not sophisticated enough to cope with social context and the linguistic nuance of



Modern regulatory law does not shy away from the fact that its constituent rules and principles are anticipatory in nature and that, as a result, they communicate information that will be in need of refinement and clarification after the point of their making. In fact, it addresses this problem of uncertainty in the following manner: It expects rule users to interpret the content of the rule in question in the light of their cultural, social and economic background.<sup>58</sup> The deployment of rigid and inflexible language is by and large counter-productive for the regulation of complex and dynamically evolving issues.<sup>59</sup> To be sure, the ambiguity of open-ended legal language gives rise to a multitude of interpretations and engenders arbitrariness, but this problem is to be kept under control through a series of mechanisms of checks and balances. Being embedded into procedural aspects of interpretation, these mechanisms are deliberative, and they rely on the rulebook users' reflection, reasoning and constructive contestation.<sup>60</sup>

Compared to modern regulatory law, the architecture of machine-executable regulation is conspicuously different. We know from legal practice that what a rule means is not given *ex ante*. It is context-sensitive and requires argument and contestation. Machine-executable regulation assumes the opposite. The mathematical

---

regulatory law). See further, E Zamir and D Teichman, *Behavioural Law and Economics* (OUP 2018) 141, 152 (considering how context affects people's heuristics, biases and ultimately decisions including decisions relevant to rule following).

<sup>58</sup> Bank of England transforming data collection plan (n 3) para 3.3 (describing how regulatees engage with financial regulators, industry associations, peers and other stakeholders in their attempt to interpret reporting requirements, hence, highlighting the participatory and dialectical character of that process). The procedural aspects of human interpretation and their importance are further discussed in section 5 below.

<sup>59</sup> In relation to private law transactions, lawyers are also accustomed to the use of open-ended legal language in their attempt to eliminate the need to draft fully contingent agreements. For a general discussion, see JM Skarloff, 'Smart Contracts and the Cost of Inflexibility' (2017) 166 *University of Pennsylvania Law Review* 263; and further EP Schuster, 'Cloud Crypto Land' (November 21, 2018). LSE Legal Studies Working Paper 17/2019, (2021) 84 *Modern Law Review* [\_\_] (in press) 24-26 (discussing the practical value and cost efficiency of open-ended legal language in the context of smart contracts) <<https://ssrn.com/abstract=3476678>> accessed 20 August 2021.

<sup>60</sup> Hildebrandt (n 16) 25. See further section 5 below, where I consider this issue in further detail.

compression of information into code makes information inaccessible (at least temporarily) while under the control of its *owner* which is here understood in the broad sense. At the same time, the compressed information is presumed to have the following characteristics: It is uniform in the course of time and independent from the person to whom the information is addressed and from other pieces of information.<sup>61</sup> In short, the mathematical expression of information turns the communicated bits of data into interpretations of themselves.

The collapse of space between information and its interpretation is deeply problematic from the legal point of view. This becomes clear once we think that something that is information to you is not information to me, and what is relevant or useful information to you, may not be relevant or useful information to me.<sup>62</sup> Training programmes for software developers and other professionals involved in coding and data validation, and the development of codes of best practice for the regulation of their professional conduct are only some examples of ways in which this concerning implication of the mathematical compression of information may be addressed.<sup>63</sup> These measures ensure that machine outputs are checked and validated by humans with the necessary expertise and they also help human rulebook users understand how machines process data and what assumptions are embedded into their statistical modelling in order to be able to work out the grounds behind a metarule as a machine output. However, it is important to note that they do not restore the lost public space

---

<sup>61</sup> Here I draw on the mathematical concept of information of Claude E Shannon. See CE Shannon, 'A mathematical Theory of Communication' (July and October 1948) 27 Bell System Technical Journal 379-423, 623-656. For an insightful discussion of how maths is used to compress, process and transmit information (eg, information about law) with speed integrity and confidentiality and further reference to literature on the mathematical theory of information and cybernetics, see Hildebrandt (n 16) 15-21.

<sup>62</sup> Hildebrandt (n 16) 20.

<sup>63</sup> J Burrell, 'How the Machine "Thinks": Understanding Opacity in Machine Learning Algorithms' (2016) 3(1) Big Data and Society 12; Wolpert and Macready (n 35) 67-82; Zalnieriute et al (n 45) 425; Kroll et al (n 44), 633; FCA, Pilot Phase 1 Report (n 5) 12.

for constructive deliberations through which the meaning of rules becomes alive and which regulatory law affords. They are not meant to change the architecture of machine readable and executable regulation; they work within that architecture.

To conclude, the limited translatability of the content of the financial rulebooks makes plain that human decision-making will be indispensable for the interpretation of regulatory content given the present and foreseeable stage of development of the relevant technology. On the one hand, it will be necessary for the application of those parts of the financial rulebooks that are not fully amenable to coding. On the other hand, it will be essential for addressing errors in relation to the fraction of rulebook content that is possible to translate into algorithmic language. For example, it will be required for weighing factors that were not thought of at the time of the making of the computer programme that generated the metarule in question, and for identifying any mistakes in the decisions that were informed by a faulty metarule.<sup>64</sup>

#### **4. Specialist intelligence and its limitations in interpreting regulatory content**

The interpretation of a legal text involves analogical reasoning.<sup>65</sup> In its simplest form, this type of reasoning can be broken down into the following tasks: Once a decision

---

<sup>64</sup> M Oswald, 'Algorithm-assisted decision-making in the public sector: framing the issues using administrative law rules governing discretionary powers' 376(2128) *Philosophical Transactions Royal Society* 1, 7 <<https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2017.0359>> accessed 25 August 2020.

<sup>65</sup> CR Sunstein, 'On Artificial Intelligence and legal reasoning' (2004) 8 *University of Chicago Law School Roundtable* 29 at 31-32; and B Sheppard, 'Incomplete Innovation and the premature disruption of legal services' (2015) *Michigan State Law Review* 1787 at 1870. For a general discussion on the nature of legal reasoning, see N McCormick, *Legal Reasoning and Legal Theory* (OUP 1978; reprinted 1993).

maker has all the available information about the factual background of the issue at hand, they look at the history of the applications of the relevant law to figure out how that law may apply in the present situation.<sup>66</sup> Specifically, they locate similar past situations and their determinations, they compare them with the current case, and they identify similarities and differences to discern the most relevant ones and form a conclusion about what the law requires in the current situation. In performing all these functions, a decision maker does not just try to discover what others thought about a similar situation in the past.<sup>67</sup> To do the job properly, they need to take a reflective stance, namely, to come up with a principle that makes the best sense of the relevant past decisions or determinations and then apply that principle on the case at hand.<sup>68</sup>

In interpreting a legal rule, humans are able to grasp the meaning of a legal rule because they have a shared understanding of what a rule is, what following a rule is, and what words mean.<sup>69</sup> They attribute meaning taking into account the context within which the rule in question applies. Furthermore, they scrutinise and review each other's attributions as active participants in a community of interpreters.<sup>70</sup> Machines do not grasp the meaning of what they read the same way as humans do. A key feature of machine learning is that it is driven by a statistical model, whose design embeds a system of scoring and typically involves impenetrably complex

---

<sup>66</sup> Here, I am only outlining a series of tasks that interpreters typically carry out. I am not concerned with the past decisions as formally binding precedent, which is a different matter. For a general discussion, see G Lamond, 'Precedent and Analogy in Legal Reasoning' *The Stanford Encyclopedia of Philosophy* (Spring 2016) <<https://plato.stanford.edu/entries/legal-reas-prec/> accessed> 20 August 2021.

<sup>67</sup> R Dworkin, *Law's Empire* (Hart Publishing 1998) 65.

<sup>68</sup> *ibid* 67-68.

<sup>69</sup> Here I draw on the work of Ludwig Wittgenstein on 'language games' as aspects of a shared 'form of life'. Anscombe and Rhees (n 52) paras 23, 203 and 241; McGinn (n 52) 44 and 55.

<sup>70</sup> [Footnote omitted for reviewing purposes]

calculations.<sup>71</sup> The statistical model serves a specific goal in relation to which machines learn to mine data from vast datasets, identify correlations and patterns, infer information, make predictions and produce outputs.<sup>72</sup> This goal may address a legitimate concern, for example, that of the consistent submission of specific data; but from that, it does not follow that it fully captures the policy objectives of financial regulators, or that it indeed yields correct outputs when assessed from the perspective of regulatory law.

Think of a computer software which trains machines to help regulatees comply with their legal obligation to share information and cooperate openly with the financial regulators as it is set out in PRIN 2.1.1 R and its identical FR 2.7 of the FCA and the PRA rulebooks respectively. To pre-empt the disclosure of all available information at all times and in all circumstances, suppose also that the computer software of my example trains machines to generate metarules to the effect of recommending non-disclosure each time the statistical model predicts that financial regulators are unlikely to follow up and, as a result, detect the non-release of pertinent information. Being informed by an irrelevant consideration –that of the likelihood of being caught to withhold information, metarules of that kind would almost certainly lead to wrong regulatory outcomes. And, if I am right on this observation, then, it is almost certain that these metarules would neither serve the policy objectives of financial regulation nor the delivery of strategic goals like, for instance, the avoidance of a culture of creative compliance.

---

<sup>71</sup> For a detailed analysis of different methods of machine learning, see G James et al, *An Introduction to Statistical Learning* (Springer 2017); I Goodfellow, Y Bengion and A Gourville, *Deep Learning* (MIT Press 2016). The complex calculations underpinning machine learning is known as algorithmic opacity. J Cobbe, 'Administrative Law and the Machines of Government: Judicial Review of Automated Public Sector Decision Making' (2019) 39 *Legal Studies* 635, 638-639.

<sup>72</sup> Cobbe (n 71) 637-639.

Suppose now that a machine is trained to generate metarules to help regulatees comply with a wide range of FCA rulebook provisions which mandate them to act with honesty when dealing with their clients.<sup>73</sup> To differentiate between truthful and deceptive statements and produce a metarule accordingly, the machine of my example analyses digital records of client communications for linguistic markers of deception and then generates a compliance score to a given set of statements. The software does not capture the essence of dishonesty. Instead, it works with indirect (and not always transparent) clues of dishonest communication to predict whether a particular communication will be ‘earmarked’ as honest or not.<sup>74</sup> Examples of indirect clues might be the frequency of avoiding first person singular pronouns, and the ratio of negation, equivocations and other linguistic patterns.<sup>75</sup> Predictions of that sort are grounded on certain assumptions -eg, the idea that there is a causal interface between one’s language and cognition and that telling a lie is more cognitively taxing than telling the truth, but they ignore a range of other factors which also affect linguistic patterns. Examples include underlying medical conditions and background noise.

---

<sup>73</sup> A notable example is rule COBS 2.1.1 R (1) known as ‘the client’s best interest rule’ according to which ‘[a] firm must act honestly, fairly and professionally in accordance with the best interests of its clients.’ COBS 2.1.1R, FCA Handbook of Rules and Guidance <<https://www.handbook.fca.org.uk/handbook/COBS/2/1.html>> accessed 12 February 2020. For a recent interpretation of the FCA’s client’s best interest rule, see *Quinn v IG Index Ltd* [2018] EWHC 2478 (Ch).

<sup>74</sup> F Tomas, O Dodier and S Demarchi, ‘Computational measures of deceptive language: Prospects and Issues’ (2022) Volume 7 *Frontiers in Communication* 1, 10 (pointing out that ‘the engineering literature relies almost exclusively on powerful algorithms without commonly mentioning the underlying theory explaining the difference between deceitful and sincere narratives.’) <<https://www.frontiersin.org/articles/10.3389/fcomm.2022.792378/full>> accessed 20/5/22.

<sup>75</sup> The Linguistic Inquiry and Word Count (LIWC) is a well-known example of computer-based textual analysis. See ML Newman, JW Pennebaker, DS Berry and JM Richards, ‘Lying Words: Predicting deception from linguistic styles’ (2003) 29(5) *Personality and Social Psychology Bulletin* 665, 666. For a critical overview of computational approaches to the detection of verbal deception, see Tomas et al (n 74).

Computer-based textual analysis is also blind of context. To address this issue at least in part, software developers often introduce an element of randomness to improve the learning capabilities of machines. For instance, they design software which programmes machines to gather additional data, learn from it and adjust their ‘thinking’ accordingly, each time the machines detect linguistic patterns that are not pre-programmed to identify as ‘dishonest’. While randomness improves the accuracy of machine outputs and helps pre-empt ‘gaming’ namely strategic behaviour that aims to abuse the system, at the same time, it increases opacity.<sup>76</sup> Undoubtedly, computer engineers can be called upon to explain how the machine produces a certain output, and therefore their role in restoring a degree of transparency is crucial. That being said, they neither have the legal training nor indeed the *de jure* power to clarify the circumstances under which a certain kind of conduct would amount to ‘dishonesty’ in the eyes of the law.

This is not to say that human recommendations are always fully transparent. Think, for instance, the provision of legal advice.<sup>77</sup> Clients often defer to the recommendations of their solicitor without always requiring a detailed account of all the factors that informed their legal advice. However, from that, it does not follow that their legal advice is unjustifiable or unexplainable. The solicitor operates within established frameworks of professional competence, independence and accountability and they are prepared -when required- to explain and justify their recommendations. Like the solicitor of my example, financial regulators and other public decision makers also operate within established frameworks of competence, independence and

---

<sup>76</sup> Kroll et al (n 44) 653-656. On the inscrutability of various types of machine learning, see MH Jarrahi, ‘In the age of the smart artificial intelligence: AI’s dual capacities for automating and informing work’ (2019) 36(4) *Business Information Review* 178, 182-183 (describing AI decision-making as a ‘black-box’ compared to human decision-making).

<sup>77</sup> Oswald (n 64) 6 (using the provision of medical advice by way of example).

accountability through various mechanisms of appeal and scrutiny. Their decisions are reviewable, explainable and justifiable.

Finally, a further point of difference between human interpretation and machine-outputted determinations of regulatory content is that, in the latter case, the formation of an underlying structure of meaning is absent. The development of semantic web is of course possible, but machine learning technology does not require an initial concept of a pre-programmed rule structure.<sup>78</sup> While this is an important advantage in one sense, it is troubling in another. Precisely because it does not presuppose the modelling of a semantic web, this type of technology exhibits greater scalability. However, the internal logic of machine learning does not follow established rules of inference, as humans do when they engage in interpretation and, consequently, it cannot guarantee interpretive coherence.<sup>79</sup> What it does is to conduct statistical analysis and, accordingly, generate micro-directives with the sole criterion being that of the delivery of the automated system's goal. The generated metarules may be connected with the specific goal that the automated system is designed to serve, but they are not connected to each other.<sup>80</sup>

To conclude, machines can carry out a series of tasks of analogical reasoning. They can retrieve factual information, identify matching past legal facts, enlist their similarities and differences, rank them in terms of relevance and use statistical modelling to output compliance scores in great speed. What they cannot do -at least not yet- is to root interpretive determinations on judgments of principle according to

---

<sup>78</sup> B Sheppard, 'Warming up to inscrutability: How technology could challenge our concept of law' (2018) 68 (Supplement 1) *University of Toronto Law Journal* 36, 46-47.

<sup>79</sup> *ibid* 46.

<sup>80</sup> *ibid*.



public criteria that are open to intelligible scrutiny and contestation. Humans do better in normative reasoning because of their moral imagination, empathy and capacity of critical judgment, but they are slow in navigating through voluminous legislation and case law, for the retrieval of factual information or the identification of matching past legal facts without the aid of machines. To be sure, humans make mistakes, but at least we have a fairly comprehensive understanding of the nature of human error, and we can anticipate it. For example, judicial review offers a robust albeit imperfect pathway for the contestation of mistaken decisions in public governance, while supervisory visits and investigations help expose and scrutinise errors made by the regulatees.<sup>81</sup> By contrast, our understanding of machine error is rudimentary at present and, as a result, our tools to respond to it are lacking in sophistication.

## **5. A proceduralist justification for the preservation of human interpretation**

Were we to overcome the limitations discussed under the first two arguments and, as a result, able to train machines to produce the perfect metarules, would we still have reason to preserve human interpretation? If regulatory law were there only to communicate to regulatees what they may or may not do, we would most probably be better-off just by switching to metarules. However, the picture is more complex. In the UK, as in other modern democracies, it seems to be the case that regulatory law does more than that. It harbours an interpretive practice of constructive deliberation which cuts across decisions on the making and application of rules despite the

---

<sup>81</sup> For a similar observation focusing on regulators, see Lodge and Yeung (n 12) 23.

procedural differences between the two processes.<sup>82</sup> Accordingly, a question to ask is whether -absent that interpretive practice- an algorithmic scheme of financial markets governance would be legitimate.<sup>83</sup>

The question of legitimacy is not confined to the interests of the members of the regulated financial industry. It also concerns consumers and virtually everyone affected by the decisions and actions of financial regulators. That being said, in the remainder of this section, I shall construe financial regulation as a dyadic relationship between regulators and the regulated industry as a matter of priority and on the grounds of simplicity. My thinking is the following: If I can show that my procedural justification applies to the relationship between regulators and regulatees, then further questions can and should be asked on whether insights of my analysis could apply to financial regulation as a polycentric regime. Focusing on the relationship between regulators and the regulated industry, it is interesting to note that part of the financial industry demands digitalisation and may even appear to be prepared to part with the opportunity to have a voice on how specific rulebook content is to be understood in supervisory visits or other interactions with regulators. However, from this it does not follow that algorithmic financial regulation would be legitimate if human interpretation were to be progressively eradicated as an aspect of regulatory law. On

---

<sup>82</sup> For instance, making rules almost always includes consultation with stakeholders; whereas applying rules (eg, in order to assess compliance or enforce the law) requires regulatees to co-operate with the regulator and the regulator to make a transparent and *intra vires* decisions. My conception of human interpretation as constructive deliberation is not meant to strictly correspond with either of these processes. Rather, it abstracts from their differences to bring attention to a feature that both processes share in common namely that of constructive deliberations through which regulatory law is interpreted in order to be brought into life.

<sup>83</sup> On current concerns about automation and how it would undermine Rule of Law principles, if algorithms were to substitute human judgment and discretion, see notably Hildebrandt (n 16) 1; and Pasquale (n 36) 17-43. For a proceduralist account of the Rule of Law, see J Waldron, 'The concept and the rule of law' (2008) 43(1) *Georgia Law Review* 1. For a discussion of formal and substantive accounts of the rule of law, see P Craig, 'Formal and Substantive Conceptions of the Rule of Law [1997] *Public Law* 467.

the one hand, advocates of digitalisation and automation form an interest group which is hardly representative of the entire financial industry let alone other stakeholders and, on the other hand, here the question I am asking is question of principle and not of empirical fact.<sup>84</sup>

With these clarifications in mind, one might be tempted to answer the question of legitimacy in the affirmative by appealing to the regulators' superior expertise to provide a solution to a coordination problem -all courtesy of technological advancement.<sup>85</sup> Specifically, they might say that, if regulatees would do better -say, in complying with reporting rules or in meeting capital adequacy requirements, by following metarules than by working out what to do on their own, then the authority of financial regulators would be legitimate. Despite its plausibility, this *substantive* test is flawed, because it fails to account for the procedural dimension of legitimacy in public governance.<sup>86</sup>

Several practices are structured by roles which have the attribute of authority in the sense that participation in those practices generates relationships which involve

---

<sup>84</sup> The participation of leading industry players in flagship initiatives like the FCA pilot of DRR suggests that there is a constant demand for automation from some part of the industry. At the same time, it raises concerns of regulatory capture. For a classic exposition, see notably GJ Stigler, 'The theory of economic regulation' (1971) 2(1) *Bell Journal of Economics and Management Science* 3; GS Becker, 'A theory of competition among pressure groups for political influence' (1983) 98(3) *Quarterly Journal of Economics* 171; S Peltzman, 'Towards a more general theory of regulation' (1976) 19(2) *Journal of Law and Economics* 211; RA Posner, 'Theories of Economic Regulation (1974) 5(2) *Bell Journal of Economics* 335; and ME Levine and JL Forrence, 'Regulatory capture, public interest and the public agenda: Toward a synthesis' (1990) Special Issue 6 *Journal of Law Economics, and Organisation* 167. For a critical overview, see D Carpenter and DA Moss (eds) *Preventing Regulatory Capture: Special Interest Influence and How to Limit it* (CUP 2014).

<sup>85</sup> J Raz, *The Morality of Freedom* (Clarendon Press 1988) 53 (providing a classic exposition of this view).

<sup>86</sup> Substantive accounts of legitimacy differ from proceduralist accounts in that they assess legitimacy by focusing on the content of the decisions that are communicated as instructions, orders or commands. By contrast, proceduralist accounts focus on procedural aspects of decision-making. Hybrid theories of legitimacy combine insights of both accounts. On this point, see S Hershovitz, 'Legitimacy, Democracy and Razian Authority' (2003) 9 *Legal Theory* 201, 212.

a right to rule for certain members and an obligation to obey for certain others.<sup>87</sup> The relationship between teachers with their students is a notable example.<sup>88</sup> Suppose that I am receiving lessons from a piano tutor hoping that one day I will become a virtuoso of Arnold Schoenberg's compositions. It is fair to say that I am more likely to succeed if I follow their instructions. My piano tutor has a claim that I follow their instructions because of their superior technique, specialist knowledge and expertise.<sup>89</sup> Crucially, the justification of their authority over me does not turn on how they decide what is best for my tuition. As long as I am progressing well, it makes little difference to me, for example, if my tutor reads all the relevant literature on the twelve-tone compositional structure of Schonberg's pieces or instead consults someone else on this issue in preparation for my piano lessons.

Even though the same could be said for the authority of doctors vis a vis their patients and others whom we regard as experts in their field, the nature of the authority of financial regulators is different. We care a lot about *how* financial regulators reach decisions which will then communicate as orders, instructions or commands or *how* they see to their proper application.<sup>90</sup> Even though regulators are not required to obtain the consent of the regulatees to make, apply or even enforce the rules they make, we expect the regulators' decisions to be the outcome of a participatory and inclusive process of constructive deliberation and that those at the receiving end of commands should have a degree of (or at least an opportunity for)

---

<sup>87</sup> S Hershovitz, 'The role of authority' *Philosopher's Imprint* Volume 11 No 7 (March 2011) University of Michigan School of Law 11; and *ibid* 210 (noting that the division between *rulers* and *subjects* is not to be exaggerated given the prominence of (quasi-) decentred schemes of decision-making and governance).

<sup>88</sup> Other noteworthy examples are the relationship between doctors with patients under their care and the relationship between athletes and their fitness instructors.

<sup>89</sup> Hershovitz (n 86) 212-213.

<sup>90</sup> *ibid*.

engagement in this process. These procedural aspects matter to us for several reasons. Chief amongst them is our commitment to principles of equality and personal autonomy.<sup>91</sup> While the former commands that regulatees have participation in the making of decisions that are highly consequential to them, the latter requires that they are empowered to take control of their own affairs and projects. We also believe, that as moral agents and rational beings, regulatees should be allowed the space for exercising judgment and discretion when circumstances so demand. A participatory process is also valuable to us because rigorous debates help us reach better decisions and because it promotes industry enrolment which is crucial for the more effective delivery of public policy objectives.<sup>92</sup>

These observations make plain that, unlike scientific experts, the authority of financial regulation is *practical* and not *epistemic* in nature.<sup>93</sup> An implication of this is that the obligation of regulatees' to obey is grounded on the presence of a participatory process of constructive deliberations which we regard as worthy of the regulatees' acceptance because it treats them with equal concern and respect, or because it is seen as instrumental to the delivery of an intrinsic good, for example, that of a shared interest in being able to autonomously control ones' projects. Accordingly, the criterion for assessing legitimacy is first and foremost *procedural* and not substantive in nature. The fact that the legitimacy of financial regulators can hardly ever be reducible to claims about their technocratic expertise and effectiveness is further verified by the range of attributes that we identify as indispensable to any

---

<sup>91</sup> *ibid* 214.

<sup>92</sup> Hershovitz (n 86) 213-214.

<sup>93</sup> On the distinction between practical authority and other forms of authority, see S Darwall, *Morality, Authority and Law* (OUP 2013) 135.

form of public governance with a credible claim to legitimacy.<sup>94</sup> This is not to say that expert input and effectiveness are not essential but rather to make the point that a legitimacy test that focuses only on the regulators' superior expertise and capabilities is liable to miss out other important considerations namely our expectation that decisions ought to be the outcome of a dialectic process that allows space for input and contestation when rules are made as well as when they are applied.<sup>95</sup>

It is difficult to see how the legitimacy of the regulators' authority could be established if human interpretation were to be overtaken by algorithmic decision-making and automation. Consider for instance liquidity reporting requirements -a highly technical cluster of rules which the Bank of England is currently seeking to convert into code.<sup>96</sup> The fact that they involve metrics and complicated calculations does not make these rules less consequential for regulatees. They feed into reports about the financial health and soundness of the regulated firms with micro-prudential and macro-prudential implications. Regulatees have every reason to want to have some control over their liquidity data and how it is used to ensure that it is correctly interpreted and that the conclusions to be drawn are valid.<sup>97</sup> None of this is possible

---

<sup>94</sup> For a classic account of legitimacy in regulatory studies, see Baldwin and al (n 17) 26-31.

<sup>95</sup> Here I am assuming that the authority of financial regulators is better accounted for according to a mixed model. David Estlund's epistemic proceduralism is a prominent case in point as it attempts to strike a sensitive balance between epistemic and non-epistemic considerations. See D Estlund, *Democratic Authority: A Philosophical Framework* (Princeton University Press 2009); and for a criticism, E Anderson, 'An epistemic defence of democracy: David Estlund's Democratic Authority' (2008) 5(1) *Episteme* 129, 135 (arguing that the case against 'epistocracy' is first and foremost based on non-epistemic grounds and values).

<sup>96</sup> The PRA Liquidity Monitoring Metric (LMM) tool is an algorithm published in Excel for the calculation of liquidity metrics for reporting purposes. At present, this tool is for guidance only and firms are not allowed to use it for the submission of regulatory returns as required by the rules. However, this may change in the future. As part of the Bank of England's data transformation plan, the LMM currently figures amongst the Bank's 'use cases' the upgrading of which is hoped to lead to the delivery of 'instructions as code' for a more flexible reporting. PRA, *Supervisory tools: Liquidity Tools* <<https://www.bankofengland.co.uk/prudential-regulation/publication/2013/supervisory-tools-liquidity-tools>> accessed 19 August 2021; and Bank of England transforming data collection plan (n 3) para 6.2.

<sup>97</sup> Bank of England transforming data collection plan (n 3) para 4.3 (offering an insight into industry concerns).

without space for a dialectic interaction -typically, during supervision or during investigations when the regulatees' compliance is called into question.

Undoubtedly, the task of interpretation is burdensome, and humans may exhibit predictable and often irrational behaviour,<sup>98</sup> but they remain moral agents capable of self-reflection and of taking responsibility of their acts and omissions.<sup>99</sup> These virtues of human discretion, judgment and agency are deeply embedded in the UK style of financial supervision. The supervisory approach of the FCA and the PRA is judgment driven.<sup>100</sup> Prima facie, financial regulators exercise administrative discretion over the interpretation of available factual evidence to assess things like, for example, the magnitude of emerging risks and the compliance of a specific type of behaviour with the financial rulebooks. In reality, the exercise of judgment is diffused. Following the long British tradition of self-regulation in financial markets governance, members of the industry are not passive recipients of regulatory commands.<sup>101</sup> They are expected to exercise discretion as they engage in a process of self-reflection on how they ought to run their business. Regulatory interventions are on the menu to communicate findings, interrogate business culture and require a particular course of action. However, these interventions are not meant to treat regulatees as mere objects to be controlled but as subjects capable of ruling themselves and of being accountable for their actions. To sum up, the practice of constructive deliberation in the use of rules which lies at the heart of judgment-led

---

<sup>98</sup> D Ariely, *Predictably Irrational: The Hidden Forces that shape our Decisions* (Harper Collins 2009).

<sup>99</sup> Lodge and Yeung (n 12) 31.

<sup>100</sup> [Omitted for reviewing purposes]

<sup>101</sup> On the historical evolution of financial regulation in the UK, see LCB Gower, 'Big Bang and City Regulation' (1988) 51(1) MLR 1 (describing the eventual mutation of self-regulation into statutory regulation); and J Black, 'Regulatory styles and supervisory strategies in M Moloney E Ferran and J Payne (eds) *The Oxford Handbook of Financial Regulation* (OUP 2015) 217, 219-221.

supervision betrays an enduring commitment to a participatory and constructive process of interpretation without which the legitimation of the authority of financial regulators would be questionable.

## **6. Designing the rulebooks of the future**

In view of the difficulties discussed above, the idea of designing algorithmic rulebooks to allow full automation in place of human interpretation should be abandoned. Instead, a more cautious approach is required namely one which will aim to harness interpretation by helping human rulebook users take advantage of both their general intelligence while also benefiting from the specialist intelligence of machines.<sup>102</sup> So, if algorithmic decision-making is to co-exist with human decision-making, how are we then to design the financial rulebooks of a system of governance that aspires to be data-driven?

A good starting point here is to decipher the nature of the relationship between the natural language which underpins human decision-making and the algorithmic language that enables the execution of various tasks by machines. There is no doubt that the use of code promises to lift barriers, gaps and other obstacles that currently inhibit procedural efficiency. However, Frank Pasquale notes, that '[w]hile computer code and human language both enable forms of communication, the affordances offered by each are distinct and, in many respects, mutually exclusive. Code seeks to eliminate the forms of ambiguity and flexibility characteristic of much language,

---

<sup>102</sup> M Lodge and A Mennicken, 'Reflecting on Public Service Regulation by Algorithm' in K Yeung and M Lodge (eds) *Algorithmic Regulation* (OUP 2019), 178, 180 (viewing algorithmic regulation as an extension of existing regulatory systems but of different nature).



including legal language.’<sup>103</sup> In the first instance, if both forms of communication were to co-exist, the tension would be inevitable.

To be sure, the tension between linguistic precision and vagueness is always present in financial regulatory law.<sup>104</sup> The open-ended language of the statutory objectives of the FCA and the PRA and of a good number of rulebook provisions coexists with the more precise formulation of a plethora of more detailed rules and non-binding guidance.<sup>105</sup> However, the proliferation of metarules is set to push the trend towards greater linguistic granularity even further as this will be necessary to promote standardisation, consistency and accuracy. If this trend is left unattended, it will progressively close down the necessary linguistic space for interpretive judgment. Moreover, the emerging data-driven architecture will place rulebook users under growing pressure to lean towards a particular outcome or towards a particular way in which a task is to be carried out.<sup>106</sup> To understand how this might happen, suffice is to consider the impact of the automation of algorithmic decision-making on what might be described as the incentives’ problem in judgment-led supervision.

The fact that regulators and regulatees are allowed the necessary linguistic space to exercise judgment (at least for the time being) does not mean that they are

---

<sup>103</sup> Pasquale (n 36) 3.

<sup>104</sup> The choice of legal form and ultimately the design of financial rules is a matter of balancing a series of competing policy considerations. For a classic discussion, see Diver (n 19) 73-74.

<sup>105</sup> On the statutory objectives of the FCA, see FSMA 2000, section 1B(2) (defining the FCA’s ‘strategic objective’ as ensuring that ‘markets function well’). On the FCA ‘operational objectives’, see FSMA 2000, sections 1C (consumer protection objective), 1D (market integrity objective) and 1E (competition objective). On the Bank of England’s statutory objectives when acting as the PRA see, FSMA 2000, section 2B (defining the authority’s ‘general objective’ as ‘promoting the safety and soundness of PRA-authorised persons’), section 2C (insurance objective) and section 2H (on the authority’s secondary competition objective). In addition, the FCA and the Bank of England are tasked with a series of resolution objectives. See Bank of England Act 2009 section 4(3) (on relevant authorities), section 4(3A) – section 4(10) (on resolution objectives).

<sup>106</sup> Oswald (n 64) 16-17; and M Nielsen, ‘How Computers Are Changing the Way We Explain the World’ *Wired Magazine* 8 (August 2015) <<http://www.wired.com/2015/08/computers-changing-way-explain-world/>> Accessed 20 February 2020.

always willing to exercise it. Cognitive limitations, heuristics, biases and ulterior motivations are bound to shape attitudes.<sup>107</sup> For members of the financial industry, the exercise of judgment is often associated with the risk of interpretive error which they would rather avoid due to concomitant costs.<sup>108</sup> For regulators, judgement-led supervision quite often implies greater exposure to blame in case of failure. Undoubtedly, financial regulators do (and should) have the power to overrule machines but, if their judgment is to be increasingly perceived as relying on ‘personal views’ instead of the ‘science’ of machine-outputted micro-directives, their future confidence in their judgment should not be taken for granted. For example, if at some point the bureaucratic culture within their internal organisation commands that an algorithmic prediction shall be followed as a matter of best practice, it will be difficult to insist on the value of judgment. To be sure, successful judicial review proceedings might challenge that emerging practice on the grounds of failure to consider relevant factors or due to improper delegation of power to an algorithm amongst others.<sup>109</sup> However, the potency of administrative law to counteract this trend is not a reason to be complacent, as the desirability to accommodate emerging ‘best practices’ in public administration may eventually instigate a shift of paradigm in that field of law with

---

<sup>107</sup> Zamir and Teichman (n 57) 393-399; CR Sunstein, ‘Nudges.gov: Behaviorally informed regulation’ in E Zamir and D Teichman (eds) *The Oxford Handbook of Behavioral Economics and the Law* (OUP 2014) 719, 721-725; J Rachlinski and C Farina, ‘Cognitive Psychology and Optimal Government Design’ (2002) 87 *Cornell Law Review* 549; RB Korobkin and TS Ulen, ‘Law and Behavioral Science: Removing the Rationality Assumption from Law and Economics’ (2000) 88 *California Law Review* 1051; N Rangone, ‘Making Law Effective: Behavioural Insights into Compliance’ (2018) 9 *European Journal of Risk Regulation* 483, 484.

<sup>108</sup> Creative compliance may stand at odds with the imperatives of judgment-led supervision and the regulators’ aspiration to foster a culture of reflective compliance but from the cost-efficiency point of view it is a rational approach to follow.

<sup>109</sup> For an insightful discussion of how UK administrative law might react to the use of automated decision-making in public administration, see n 67, 635-655. On the treatment of improper delegation by UK administrative courts, see notably *Ellis v National Dock Labour Board* [1953] 2 QB 18; and *H Lavender & Sons v Minister of Housing and Local Government* [1970] 1 WLR 1231 (taking advice from others does not amount to improper delegation as long as public officials did not have the decision dictated to them); and *R v Home Secretary Co Ltd v Minister of Technology* [1971] AC 610, 625. On the treatment of irrelevant considerations, see *R v St Pancras Vestry* (1890) 24 QBD 371, 375 and *R v Home Secretary ex p. Venables* [1998] AC 407.

potential repercussions on future perceptions about appropriate industry attitudes to automation and concomitant business culture.<sup>110</sup>

With the propagation of metarules, the incentives' problem is set to worsen due to the growth of automation bias and the concomitant phenomenon of deskilling. Automation bias is the unfounded but nevertheless strong belief that -compared to humans- computers are more rational and objective in their decision-making.<sup>111</sup> Deskilling is intertwined with automation bias.<sup>112</sup> Decisions about the use of legal rules involve normative reasoning skills including the capacity to sense a degree of social connection, critical judgment, empathy and moral imagination. As rulebook users will have less of an opportunity to develop those skills, their ability to engage in normative reasoning and to appreciate the moral choices of their actions will

---

<sup>110</sup> In the US, this swift of paradigm on perceptions of due process seems to be happening already in relation to criminal law proceedings. See notably *State v Loomis*, 881 N.W.2<sup>nd</sup> 749 (Wis. 2016) in which the sentencing judges referred to an algorithmic risk assessment tool without disclosing the methodology of the risk assessment that was embedded into the tool. The Supreme Court of Wisconsin found that the court's approach did not violate the due process rights of the defendant. It concluded that, in the circumstances, judges made an independent decision as regards the likelihood of recidivism, and that they referred to the algorithmic risk assessment tool only as a source of information. Critics argue that the secrecy surrounding the modelling of the risk assessment tool hampers any meaningful scrutiny of the reasoning of the judges. For instance, it is not possible to tell what data points were selected as relevant for the assessment of risk. See 'Criminal Law - Sentencing Guidelines - Wisconsin Supreme Court Requires Warning before Use of Algorithmic Risk Assessments in Sentencing - *State v. Loomis* 881 N.W.2d 749 (Wis. 2016)' (2017) 130 *Harvard Law Review* 1530; and F Pasquale, 'Secret algorithms threaten the rule of law' *MIT Technology Law* (1 June 2017) <<https://www.technologyreview.com/2017/06/01/151447/secret-algorithms-threaten-the-rule-of-law/>> accessed 1 October 2020. Similar concerns have been voiced in the UK in relation to the so-called Harm Assessment Risk Tool (HART) which is deployed by police to support their decisions notwithstanding the openness of the British police as regards features of the HART's modelling. For an insightful analysis from the point of view of natural justice and administrative law principles, see notably Oswald (n 64) 7.

<sup>111</sup> Contrast here with the non-pejorative and purely technical use of the term 'bias' in probability, statistics and machine learning and social sciences as described by T Scantamburlo, A Charlesworth and N Christianini, 'Machine Decisions and Human Consequences' in K Yeung and M Lodge (eds), *Algorithmic Regulation* (OUP 2019) 49, 57-58.

<sup>112</sup> *ibid* 75-76. On the degradation of cognitive skills, see further, K Volz, E Yang, R Dudley, E Lynch, M Dropps, MC Dorneich, 'An evaluation of cognitive skill degradation in information automation' (2016) 60(1) *Proceedings of Human Factors and Ergonomics Society 2016 Annual Meeting* 191-195.

atrophy.<sup>113</sup> Furthermore, in the absence of any countermeasure, automation bias will show no sign of abating despite the fact that machines will continue to take faulty decisions by law's standards.

To ensure that any problems are kept under control and that no stakeholder becomes worse off as a result of the increasing use of metarules in the future, it is necessary to establish the primacy of natural language for the communication of rulebook content. The following principles are therefore recommended for the design of the financial rulebooks: The first of those principles is the principle of *optimisation*. According to this principle, financial rulebooks should be drafted in a way that helps users benefit from both the general intelligence of humans as well as the specialist intelligence of machines. The remaining principles are mutually reinforcing and introduce qualifications to the principle of *optimisation*. Specifically, the principle of *user-centricity* requires that human rulebook users be treated as active members of the regulatory community rather than passive recipients of explanations about what they are to do or what they did wrong in a given situation.<sup>114</sup> The implication of that principle is that human judgement precedes over machine outputs and that humans remain in control of and ultimately responsible for machine determinations of regulatory content as it befits their moral agency. Next is the principle of *mutual exclusivity* of natural language and algorithmic language as forms of communication. According to that principle, algorithms should not be introduced into value-laden assessments for, otherwise, there is a risk that judgment and discretion are eroded. Finally, the fourth principle is the principle of *non-*

---

<sup>113</sup> SV Shiffrin, 'Inducing Moral Deliberation: On the Occasional virtues of fog' (2010) 123 HLR 1214, 1222, 1244. For a more general discussion, see L Alexander and E Sherwin, *The Rule of Rules: Morality, Rules and the Dilemmas of Law* (Duke University Press 2001).

<sup>114</sup> Oswald (n 64) 7.

*substitutability* of human decision-making, which is here broadly conceived to underscore two things: The crucial function of human decision-making for the detection of errors by law's standards as well as the deliberative character of the process of interpretation that is embedded into the use of rules as a prerequisite of regulatory legitimacy. Taken together, these design principles are desiderata which serve as criteria for assessing how far to go with the coding of rulebook content for a more data-driven scheme of governance.

To make sure that future rulebooks pass the test of all the above four design principles, their content must reflect a clear division of labour between human decision-making and machine decision-making. This can be done through the progressive introduction of two-tiered rulebook content expressed in natural and algorithmic language respectively. The allocation of different rules to humans and machines would not be easy and will most probably require a radical reconceptualisation of the current content of the rulebooks. One possibility might be to split tasks between humans and machines in terms of those aspects of analogical reasoning that humans are known to do better compared to machines and vice versa and then re-write the content of the rulebooks accordingly.

Written in natural language and intended for humans, the first tier would be similar to the existing rulebooks. It would consist of legal rules of varying degrees of linguistic vagueness and precision to accommodate the use of general human intelligence. It would cover the full spectrum of rules and requirements of the existing financial rulebooks to regulate how human decision makers root interpretive determinations in judgments of principle and to enable human oversight over

algorithmic decision-making. The second tier would be for machines. It would be written in algorithmic language to facilitate the use of specialist machine intelligence for the execution of all other tasks of analogical reasoning under the necessary human oversight namely the retrieval of factual information, the identification of the applicable set of rules, the navigation of the history of their past applications, the generation of compliance scores and the mapping of similarities, differences or other correlations.

With this two-tiered structure in place, the penetration of algorithms into the governance of financial markets will improve effectiveness without posing a serious threat to the skills and confidence of regulators in exercising judgment and discretion when circumstances so require. Ultimately, it will be left for them to explain, for instance, why a machine output should be followed, called into question or even overridden in the case at hand. On their part, regulatees will be able to reap procedural and cost-efficiency benefits, but they will still be required to exercise judgment on what is prudent, fair, honest and reasonable where appropriate.<sup>115</sup> More generally, all determinations of regulatory content would continue to be open to public scrutiny, contestation and review by human rulebook users.

At present, technology solution vendors and other private sector technology firms are at the forefront of the conversion of regulatory instructions into code and of the authorship of the relevant protocols and operating manuals. The conspicuous

---

<sup>115</sup> A further advantage of this approach would be to help put an end to the worsening of the incentives' problem in judgment-led supervision especially if it is also supported with a series of other measures. Examples include training to cultivate an aptitude for informed exercise of discretion and investing on innovation and cross-disciplinary synergies to improve the reviewability of metarules and the development of robust governance controls over how the industry makes use of software experts and other external professionals and for what purpose.

presence of the private sector ensures that the latest cutting-edge research feeds into the development of regulatory technology, but there is a downside to this arrangement: For all intents and purposes, the coding of the more specific instructions and templates is left in the hands of a group of stakeholders that neither has the de jure power to clarify the content of regulatory law nor is subject to appropriate accountability requirements. To address this problem, it is further recommended - albeit not without controversy- that financial regulators take direct control over the algorithmic conversion of the rules they make and that they continue to develop synergies with all relevant stakeholders to benefit from expert input.<sup>116</sup> In this regard, the input of lawyers would be crucial in helping computer scientists appreciate the nature of modern regulatory law, the value of constructive deliberations that are woven into the use of rules, and the need for an algorithmic architecture of financial regulation apt to accommodate human action, individual freedom and personal autonomy.<sup>117</sup>

The two-tiered financial rulebooks are expected to be greater in size, volume and complexity compared to the existing ones and they will not come cheap. However, here, I will not be concerned with the question of cost-efficiency on purely methodological grounds: If it can be shown that it is desirable and, in principle, feasible to design rulebooks as recommended above, then additional questions may and should be asked about the cost efficiency implications of what will be required

---

<sup>116</sup> At present, financial regulators provide data specifications and validation. Take the example of the Data Reference Guide (DRG) for data items submitted through GABRIEL, the FCA's online system for the collection and storing of data from firms. This Guide contains XML specifications and details of validations applicable to the data submitted. <<https://www.fca.org.uk/firms/gabriel/data-reference-guides>> accessed 12 February 2020.

<sup>117</sup> Hildebrandt (n 16) 30; and Kroll et al (n 44) 699-705.

for their development. However, if it is not, then the question of cost-efficiency does not arise in the first place.

To ensure that the imperatives of the proposed design principles will not be diluted in practice, it is also essential to be clear about the legal status of all those metarules which will make up the algorithmic tier of the financial rulebooks. Metarules share certain properties with legal rules, but they lack others. For example, like legal rules, metarules exhibit a sort of linguistic structure to the extent in which code can be loosely described as a form of algorithmic language serving as the alphabet and grammar of a data-driven system of financial regulation. They also appear to display a normative dimension in the sense that they have the potential of exerting legal effects in the form of pre-programmed interpretive guidance. Given these similarities, it might be tempting to think of metarules as *sui generis* soft law. Nevertheless, this idea should be resisted at least at present because their making does not correspond to the same procedural criteria that apply to the PRA and the FCA rulebooks. The supremacy of legal rules over their algorithmic equivalent should therefore be acknowledged firmly and unequivocally in primary legislation so that it becomes mandatory to all. This option may lack the appeal of market-based voluntary arrangements, but it merits consideration because it guarantees that the primacy of legal rules will neither hinge on the good will of rulebook users nor on the effectiveness of contractual mechanisms of enforcement.<sup>118</sup>

---

<sup>118</sup> K Yeung, 'Regulating blockchain: The emerging battle for supremacy between the code of law and code as law' (2019) 82(2) *Modern Law Review* 207, 220-224.



## 7. Conclusion

In their attempt to implement a more data-driven approach to regulation, the FCA and the Bank of England are experimenting with the creation of rulebook content that can be read and executed by machines. Re-writing rulebook content into code promises to lift the burden of interpretation of an increasingly complex volume of rules and requirements with profound benefits for the industry and the regulators alike. At the same time, it raises important questions about the future place of human interpretation as an aspect of algorithmic financial regulation and how far to go with automation.

As I argued, we have good reasons to ensure that human interpretation remains an indispensable component of the emerging data-driven governance. To support my thesis, I brought attention to the limited translatability of regulatory content into algorithmic language and to the difficulty of machines to engage with the full spectrum of tasks of analogical reasoning. I further showed that, even if it were possible to overcome these challenges, it would be desirable to preserve human interpretation on procedural grounds pertaining to the legitimacy of the regulators' authority namely on grounds that go beyond considerations of efficiency and evidence-based expertise.

If human interpretation is to remain a feature of data-driven regulation, it is essential to have in place rulebooks which will help their human users take advantage of their general intelligence as well as the specialist intelligence of machines. In view of this, I concluded the discussion by putting forward a series of design principles for the drafting of the rulebooks of the future. These are the principle of optimisation, the

principle of user-centricity, the principle of mutual exclusivity of natural language and algorithmic language and the principle of non-substitutability of human decision-making. The proposed rulebooks will not be perfect and their design may require further refinement but they will preserve judgment and discretion as vital skills for normative reasoning and they will ensure that determinations of regulatory content remain susceptible to public contestation and scrutiny by humans, so that no stakeholder becomes worse off as a result of the penetration of algorithms into the governance of financial markets.

My thesis allows for a more balanced assessment of the benefits of algorithmic financial regulation. It also sheds light on the often-neglected procedural aspects of human interpretation and what they stand for in public governance. To be sure, it does not have all the answers to the problems that we are likely to encounter as the digitalisation of financial regulation will be progressing and possibly expanding beyond information gathering as it's been currently contemplated. However, it merits consideration because, at the very least, it shows that human interpretation is not an impediment but a prerequisite of regulatory effectiveness. It is imperative therefore that it is treated as such. The goal then for the years to come is to develop a digital architecture which -instead of negating- it enhances the inclusive, participatory and constructive character of interpretation as a process while, at the same time, improves the aptitude of human rulebook users for judgment and discretion when circumstances so demand.