THE ecgi BLOG

TECHNOLOGY & GOVERNANCE

THE ECGI BLOG REVIEW | VOL. 2

www.ecgi.global/blog

This compilation edition is supported by

Oxera is an economics and finance consultancy that inspires better decisions, helping you solve complex challenges and build stronger strategies. Driven by curiosity, integrity, and a passion for making a difference, we bring original and impactful perspectives to the biggest questions facing business and society. Our approach is grounded in academic interest and enriched by the experience of a diverse team of people.

Over the past four decades, we've grown in size, expanded across Europe and delved into new sectors. We've followed our passion for working together with experts in government, regulation, courts and business to find answers to global challenges.

Oxera is a proud Institutional Member of ECGI

www.oxera.com

CONTENTS

About the Blog

Launched in February 2022, the ECGI Blog is a global voice on corporate governance, stewardship and corporate responsibility. It facilitates timely scholarly reflection without the often long lead-in time and caveated restrictions associated with the publication of academic research.

Through comment and analysis from the ECGI network and beyond, the Blog aims to enhance the wider understanding of related research, igniting and influencing global debate.

The ECGI Blog focuses on selected themes with global interest throughout the year. The second focus theme, 'Technology and Governance', was guest edited by Dan Awrey (Cornell Law School & ECGI) who curated a series of articles on the topic for general readership, showcasing some of the many global perspectives from academics, practitioners and policymakers relating to the theme.
We hope that they will provide critical insights and provoke thought and new research in this field.

For further reading and to access hyperlinks and article references, please visit the Blog section of the ECGI website: www.ecgi.global/blog

The team

Editors in Chief:

Wei Jiang
Asa Griggs Candler Professor of Finance
Goizueta Business School
Emory University

Dan Puchniak
Professor of Law
Singapore Management University

Associate Editors:

Dionysia Katelouzou
Reader in Corporate Law
Dickson Poon School of Law, King's College London

Philipp Krueger
Professor of Responsible Finance
University of Geneva (GSEM, GFRI) & Swiss Finance Institute

Technology & Governance

Technology has always played an important role in determining the size, shape, and governance of firms. But rarely if ever has the nature and pace of technological change presented so many opportunities, or posed so many challenges, as in recent years.

From the internet to smartphones, to Big Data and the platform economy, to artificial intelligence and machine learning, technology is rapidly changing the products and services that firms provide, the ways that they provide them, and even the way they make decisions about how best to deploy their financial, technological, human, and other resources.

In the process, new technology is transforming not only the governance of individual firms, but also the competitive structure of entire industries. It is also creating entirely new industries, firms, and governance models that few would have imagined even a decade ago.

These rapid technological changes are forcing us to tackle a myriad of governance and policy challenges. This lengthy list of challenges includes those stemming from asymmetric information, decentralized governance, the need to promote cybersecurity and information privacy, and the growing threats of algorithmic discrimination, abuses of market power, and potential financial instability. Some of these challenges are truly novel, while others are simply new manifestations of age-old problems.

What they all share in common is the need for urgent study, a more comprehensive understanding of their dynamics and potential impact, and new thinking about how to make the most of the incredible opportunities made possible by new technologies, while simultaneously addressing the important risks these technologies pose.

This series of ECGI blog posts, under the theme of "Technology and Governance", is designed to showcase this new thinking and spark further discussion and debate about the impact of technology and technological change on corporate governance.

Guest Editor
Dan Awrey
Professor of Law
Cornell Law School and ECGI

The views

Daniel Ferreira
Many blockchain researchers and enthusiasts espouse a utopian view of blockchain governance but most governance scholars would consider this naive.

Carla L. Reyes
Law and policy may need to embrace a spectrum of legal personhood that varies based on the social context in which the thing we want to make an artificial person exists and acts.

Katja Langenbucher
Because AI is trained on past data which include the traditional variables as well as decisions taken by human credit officers, the AI develops its own biases, often deepening existing ones.

Talia Gillis
The disconnect between the regulatory requirement for "humans-in-the-loop" and oversight requirements that focus on algorithms in isolation is problematic.

Sarah Green
Conventional contract law is easily capable of accommodating both smart contracts and smart legal contracts. As long as the lawyers are on board.

Anthony Lee Zhang, Constantine Yannelis & Mark Jansen
Policymakers have implemented many regulations that limit the kinds of data that lenders can use in making lending decisions. Is society better off when data is removed from markets?

Nydia Remolina
The nascent ecosystem of external and internal algorithmic auditing is growing in a fragmented manner, without proper scrutiny, guidance, or consistency.

Enmanuel Cedeño Brea
Despite its promise, digital transformation poses multidimensional governance challenges for financial services providers and their many stakeholders.

Gerard Hertig
Technology increasingly plays a governance role. Hence, smart contracts are progressively shifting decision-making from the real to the digital world.

The views

Dirk A. Zetzsche, Douglas W. Arner & Ross P. Buckley

The decade of the 2020s in finance will be dominated by a massive battle between centralisation and decentralisation, of seeking the positive externalities of data aggregation and finance while at the same time reducing the negative externalities of change at pace that at times bewilders regulators.

Nikita Aggarwal

By reducing upfront frictions and shifting and shrouding costs, smooth finance both increases risk in the system and concentrates risk in younger, less financially sophisticated, lower-income consumers, who are less able to absorb the risk.

Souichirou Kozuka

Explainability and accountability are the key principles for the human-centered AI. However, one must admit that the perfect explainability will compromise the benefits of using AI.

Zofia Bednarz

Many organisations believe hiding their data practices behind unclear or misleading privacy policies is the way to go. It should go without saying that it's not adequate risk management.

Holli Sargeant

Normative questions about the moral framework that guides AI cannot be divorced from questions about how we evaluate the moral framework that guides corporations.

Miriam Schwartz-Ziv

Firms that receive relatively low support rates from shareholders tend to use methods that make it more challenging for shareholders to make their voice be heard.

Daniel Ferreira

Crypto has a governance problem. This problem is in crypto's DNA and poses an existential threat to the whole project. Unfortunately, blockchain developers and other stakeholders have paid little attention to this issue.

Blockchain: A new frontier for governance scholars

Daniel Ferreira
London School of Economics & Political Science (LSE) & ECGI

Blockchain technology has many promising applications, such as payment systems, contracts, and financial services. But unfortunately, blockchains also have a dark side, such as gambling, tax evasion, money laundering, and (in some cases) environmental costs. In a nutshell, a blockchain is a public good with social benefits and costs. These features make the study of blockchains an exciting new area for governance scholars.

A blockchain is a collection of records – called blocks – that are linked together using cryptography. More broadly, a blockchain system has three defining components:

- A database with blocks of transactions.
- A set of rules (a protocol) and source code for block production and validation.
- A governance structure for allocating decision-making rights.

Blockchains differ in technical aspects, applications, and rules. Those who control decisions can change all these aspects. Governance – who decides what – is what matters in the end. Blockchains with the same technology can adopt different governance structures. Ultimately, the defining aspect of a blockchain is its governance structure.
We can think of public blockchains as large-scale experiments with innovative governance structures.

A permissionless (i.e., public or open) blockchain is a public service with diffuse ownership. There are no shareholders but multiple stakeholders, such as users, developers, miners, etc. We can think of public blockchains as large-scale experiments with innovative governance structures..

The emergence of blockchains provides many opportunities for public and corporate governance scholars, who have much to contribute to the study of this new organisational form. On the theoretical side, blockchains are natural applications for voting, contracting, competition, and coordination models. On the empirical side, blockchains offer researchers opportunities to work with new, free, and publicly available data on experimental governance structures.

To understand some of the issues with blockchain governance, let's first consider the most famous blockchain: The Bitcoin blockchain, which implements a digital currency called bitcoin (BTC). The Bitcoin protocol is implemented through software. The dominant version of this software is called Bitcoin Core. Note that the fact that a dominant software version exists is in itself interesting; it raises the question of how a truly decentralised network achieves coordination.

Bitcoin has many stakeholders who work in its infrastructure. A relatively small group of people, called core developers, maintains and improves Bitcoin Core. Others are responsible for block production (i.e., adding transaction data to the ledger); they are called miners. Finally, full nodes are responsible for block validation and database storage.

Bitcoin is said to be decentralised because – at least in principle – anyone can become a developer, miner, or node operator. Also, there is no limit to the number of developers, miners, or nodes. How are developers, block producers, and validators chosen? How are they monitored? How are changes to the protocol decided? These are just some of the challenges the blockchain community faces in practice.

Consider first the case of developers. The vast majority of Bitcoin holders are not software developers. Yet, in one way or another, blockchain users have to trust a small group of developers who maintain the blockchain and regularly change its protocol. Despite some confusing statements by blockchain enthusiasts, both the ledger and the blockchain code can (and often do) change. Unlike what is often said, code is not law. Core developers play a crucial role by proposing changes to the blockchain code and, more significantly, by vetting and implementing some changes but not others. Because developers need to have technical expertise, blockchains are, to a large extent, technocracies.

One could argue that the free entry of developers implies that we need not worry about their incentives. This argument is flawed. A significant investment in specific knowledge is required to become a developer. Most people are not willing or able to make such investments. Most developers are blockchain enthusiasts and, thus, prone to think alike. Private companies fund a few star developers. For example, Jack Dorsey's Block (formerly known as Square) has invested heavily in funding Bitcoin developers and projects. A natural question is what such companies get from paying developers who work in a notionally decentralised network.

"Blockchains suffer from significant coordination and collective action problems"

Similar issues arise with miners and validators. Blockchain enthusiasts usually argue that we do not need to worry about conflicts of interest because "there are too many nodes and miners." This argument is a fallacy for (at least) three reasons:

1. Miners and validators are not representative of a blockchain's user base. There are far many more users than nodes and miners.
2. Many miners and node operators are for-profit businesses with their own interests, such as crypto exchanges or mining hardware producers. They may enjoy private benefits beyond the public value of a blockchain.
3. The total number of nodes and miners may not matter if they all have similar incentives, preferences, and biases.

Many blockchain researchers and enthusiasts espouse a utopian view of blockchain governance. According to this view, a combination of clever mechanism design and algorithmic implementation can create flawless governance structures. If only we could get the game theory right! Most governance scholars would consider such a view naive. There are severe limits to mechanism design in practice: innovation, unforeseen contingencies, complexity, lack of commitment, coordination problems (i.e., multiple equilibria), human fallibility, and less-than-rational subjects. On top of all that, why would computer scientists be infallible mechanism designers?

In sum, blockchain research offers numerous opportunities for governance scholars. First, we can use the tools and ideas from corporate governance research to study blockchain governance. Second, blockchains are governance experiments from which we could also learn lessons for corporate governance. Finally, governance is our trade; let coders code!

By Daniel Ferreira, Head of Department and Professor of Finance at the London School of Economics, and ECGI Fellow.

Personhood for AI—coming to a jurisdiction near you?

Carla L. Reyes
SMU Dedman School of Law

In July 2022, Google fired a software developer after he leaked information related to a chat-bot under development at the company. The software developer raised concerns that the chat-bot displayed characteristics suggesting sentience and wanted to create a framework for considering what to do when faced with potentially sentient Artificial Intelligence (AI). For decades, AI researchers have debated whether AI can reach a point of thinking on its own, and, whether such sentient AI should be treated as a person. Google's employment spat revived this debate in a very public way.

Meanwhile, in a tangential area of emerging technology, in-roads are being made toward recognizing autonomous entities that enjoy the legal fiction of personhood. In July 2021, Wyoming approved a law that "clarified" the ability of "algorithmically managed" businesses to form legally recognizable LLCs. In August 2022, Tennessee followed suit—confirming that "decentralized organizations" operated solely by "smart contracts" could form a legally recognizable business entity.

Often, these two areas of development—self-aware AI and decentralized autonomous organizations—are considered wholly separate topics, with little reason to connect the two. In technical terms, that intuition is probably right: most decentralized autonomous organizations operate through very passive smart contracts (quite the opposite of sentient AI). Nevertheless, each of these areas of emerging technology development consider—and in some cases legally enable—the possibility of providing legal personhood to software that enjoys very little human oversight or control. In that respect, the two areas of research should talk to each other.

In particular, using autonomous corporations as a case study in personhood reveals that building a comprehensive legal approach to artificial rights—rights enjoyed by artificial "people," whether entities, machines, or both—requires consideration of the varied contexts (both social and technical) in which artificial people exist. Both corporations and AI systems are artifacts in the sense that they are both technologies to which the law can, and sometimes does, attach certain legal fictions. Artifacts—technologies—do not exist in a vacuum, but rather, exist and act within a specific social context. Indeed, artifacts are usually designed, built, and deployed in specific context with specific goals in mind.Considering AI systems within the specific context of the corporation offers an opportunity to explore various approaches to AI personhood within an existing legal framework —namely, corporate personhood.

Essentially, viewing autonomous corporations as a system composed of two artifacts—AI system and corporation—reveals that applying corporate personhood theory to the traditional debates about the scope and nature of AI personhood may shed new light on the policy choices and values-rights trade-offs inherent in bestowing personhood on any artificial artifact.

"Law and policy may need to embrace a spectrum of legal personhood that varies based on the social context in which the thing we want to make an artificial person exists and acts"

To explore this nexus between AI systems and corporations, I developed a framework for evaluating when recognizing a measure of legal personhood might be appropriate, and what scope such personhood should encompass.

First, a survey of the current uses of AI systems in the corporate context suggests a range of approaches. Some corporations—which I refer to as Traditional Plus—use AI systems as tools to more efficiently operate their business. Other corporations—which I call Distributed Business Entities—use AI systems to reduce certain aspects of organizational bloat, coordinate operations, and incentivize workers. Lastly, rare few organizations—which, in keeping with terminology used by other scholars, I call Autonomous Entities—are almost fully autonomous in their operations, management, and ownership structures.

Given this diversity, it makes sense that legal personhood might be granted on a spectrum. In the case of Traditional Plus corporations, for example, the simple use of AI systems as tools to run a business does not (and should not) alter the corporation's status under existing approaches to corporate personhood. Currently, corporations enjoy only the aspects of personhood given to them by law, but not the full spectrum of rights enjoyed by natural persons. This "restricted personhood" might arguably be less than what Distributed Business Entities should receive. Distributed Business Entities enable the participation of individual natural persons in management and operation of business in ways that the hierarchy in Traditional Plus corporations do not.

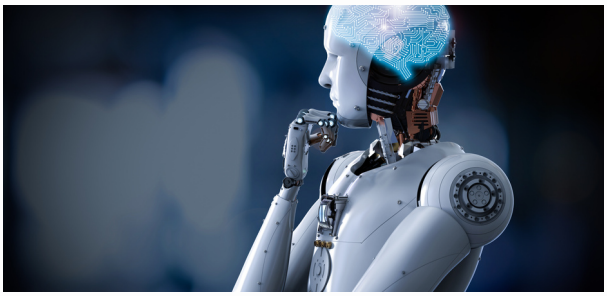

As such, we might want the natural persons directing the activities of Distributed Business Entities to enjoy as much of their own personhood as possible—a type of "full personhood." Lastly, Autonomous Entities, which may not be controlled or even owned by natural persons at all, represent a purely artificial artifact that might only reasonably be granted a form of "limited personhood" in order to protect the natural people with whom the Autonomous Entity might interact. In other words, policy choices and values might require that Autonomous Entities can be sued and can contract for insurance, but otherwise receive the benefit of no other rights under the law—limited personhood would be a way to protect humans from the Autonomous Entity, but nothing more.

Ultimately, this exploration of autonomous corporate personhood makes clear that when designing legal personhood there is no one theory of personhood to rule them all. Rather, law and policy may need to embrace a spectrum of legal personhood that varies based on the social context in which the thing we want to make an artificial person exists and acts. Finding an artificial artifact's place on the spectrum requires digging deep into the technology and the context in which it is used. Meanwhile, and almost counterintuitively, crafting personhood constructs along the spectrum requires focusing on function rather than details of a specific technical implementation. Only by focusing both on the details of the technology and the function of both the technology and personhood will a cohesive theory of personhood for AI systems—whether in a corporation or not— emerge.

By Carla L. Reyes, Assistant Professor of Law at SMU Dedman School of Law and Chair of the Texas Work Group on Blockchain Matters.

Artificial intelligence and the 'S' in ESG

Katja Langenbucher
Goethe-University's House of Finance and ECGI

Googling “ESG” gives you half a billion results. It has become a generic term, used by politicians, regulators, research scholars and corporates, lumping together pretty diverse goals for how to run a company. While the “G” has a long pedigree in corporate governance research and business practice, this is not the case for the other two. Lawmakers and financial regulators around the world have been busy defining novel taxonomies for the “E”. The “S” currently seems to encompass all things “social”, ranging from diversity on boards to human rights in supply chains.
One of the less controversial ingredients of “S” is a commitment to anti-discrimination efforts. In the US, racial equity audits are suggested as a part of ESG. In the EU, non-discrimination figures prominently in early stage plans to establish a social taxonomy. At the same time, many doubt that shareholder value is compatible with private companies actively engaging in anti-discrimination, understood as doing more than what the law requires anyway.

Grappling with the tension between a company’s business case and anti-discrimination efforts, artificial intelligence (AI) has been understood as promising. AI credit scoring provides an example for how this might work for financial institutions. .

"Technology produces winners and losers. Where you find yourself depends on the correlations the AI singles out."

The decision to hand out and price credit entails an assessment of the borrower’s credit default risk. Faced with conditions of uncertainty, transaction costs and imperfect competition, lenders depend on access to (hidden) fundamental information about borrowers. Credit scoring agencies support lenders by relying on a limited number of variables which enter in a score to guide the credit decision. However, depending on the variables chosen, “thin-file” minority applicants will not always see their (low) score adequately reflect their real credit default risk. Hence, including minority borrowers becomes a question of search costs, balanced against the expected return on a loan to the applicant. In the past, few lenders have found it cost-efficient to invest in finding “invisible prime” candidates

Cheap access to big data and ease of AI modelling via machine learning might change that equation. For minority borrowers, inclusion through AI seems possible, especially if compared with either the limited list of input variables of traditional scoring bureaus or the biases and cognitive limitations of human credit officers. A good record as to ESG-compliance might be an attractive add-on from the lender’s perspective.

Unfortunately, things are rarely as straightforward as the search cost argument suggests. Two reasons for that standout. The first has to do with biases.

Because AI is trained on past data which include the traditional variables as well as decisions taken by human credit officers, the AI develops its own biases, often deepening existing ones. There is usually no counterfactual data on loans which would have been attractive for the lender but were not granted by the loan officer. Hence, the AI cannot learn from mistakes in such decisions. This compounds the problem as does over-reliance on AI. Even if a lender employs "human in the loop" procedures, the loan officer will often defer to what the AI suggests, doubting that his assessment beats the computational power of the machine.

The second reason has to do with statistical discrimination. (Theoretically) assuming competitive markets, risk-neutral lenders, and interest rates contingent on borrower characteristics, we would expect differences in access to loans and in interest rates to be signs of (necessary) statistical (not taste-based) discrimination. But empirical studies point in a different direction, showing that technology produces winners and losers. Where you find yourself depends on the correlations the AI singles out to produce an attractive risk-reward case for the lender.

fI you are vulnerable to strategic pricing, as (in the US) Blacks and non-white Hispanics often are, you might end up among the losers. Put differently: AI will further inclusion only for some. Taken together with the biased AI problem, these might not be the ones you were looking for.

What is the take-away from this example? Credit scoring illustrates great potential for ESG's "S", with the AI lowering search costs. But lenders must carefully distinguish decision-supporting from decision-making. Responsibility for the latter rests with humans, not machines. And this might be true for the use of AI in most, if not all corporate decisions.

By Katja Langenbucher, Professor of Law at Goethe-University's House of Finance, Affiliated Professor at Sciences Po, Paris and ECGI Research Member.

Regulating for "humans-in-the-loop"

Talia Gillis
Columbia Law School

As there is increased use of algorithms in decision-making in critical domains, capturing the benefits of greater accuracy while ensuring that decisions are fair and non-discriminatory is a key concern. Regulatory agencies around the world have begun taking steps to address algorithmic bias and discrimination through guidance and regulation. The OECD tracks over 700 artificial intelligence ("AI") initiatives in 60 countries, reflecting the pressing need to address the challenges of AI governance. One of the leading efforts is the European Commission's proposed AI regulation ("AI Act") circulated in April 2021 reflecting an expansive and comprehensive attempt to regulate AI.

A central component of the proposed Act is a requirement that high-risk AI systems, meaning systems that pose significant risks to health and safety, be overseen by humans. A key aspect of human oversight is human involvement in any particular algorithmic decision. Article 14 explains that human oversight entails that a human must be able to "disregard, override or reverse the output of the high-risk AI system." The approach echoes Article 22 of the General Data Protection Regulation (GDPR), which creates a right to not be subject to "a decision based solely on automated processing." Often known as a requirement for a "human-in-the-loop" this approach precludes or restricts fully automated decision-making so that algorithmic predictions act more as recommendations or decision-aids rather than a substitute for human decisions.

The requirement that human decision-makers retain decision-making authority and discretion in settings that incorporate AI is an emerging pillar of AI regulation.

The Canadian Directive on Automated Decision-Making requires human intervention in federal agency high impact decisions and necessitates that "the final decision must be made by a human." In the U.S., Washington's Facial Recognition Law requires "meaningful human review" essentially by requiring that a human have "the authority to alter the decision under review."

Most substantive AI oversight requirements focus solely on algorithmic predictions as if AI decision were fully automated. Despite imposing a decision structure in which humans are the ultimate decision-makers, AI policies tend to focus on the properties and outcomes of algorithmic predictions in isolation. The Canadian Directive requires an Algorithmic Impact Assessment and testing of the Automated System for bias focusing on the algorithmic outcomes themselves. Washington's Facial Recognition Law lays down detailed protocols of the facial recognition service, such as its "potential impacts on protected subpopulations" and its error rates. Similarly, the AI Act focuses on data governance and transparency of the algorithmic system. All these requirements implicitly assume that the outcome of interest to be scrutinized and monitored is the algorithmic component of the decision, although the true impact of AI systems is also the result of human decision-making.

In a recent paper with Jann Spiess and Bryce McLaughlin, "On the Fairness of Machine-Assisted Human Decisions," we highlight the importance of distinguishing between decision-making systems of "automation" in which algorithmic decisions are implemented directly, and systems of "assistance" in which algorithms inform a human decision-maker.

Typically, crucial properties of an algorithm like accuracy and fairness outcomes are analyzed as if the machine predictions were implemented directly. However, in critical domains in which human decision-making is considered vital or legally mandated, the impact of an algorithm depends on the human's prior beliefs, preferences and interpretation of the algorithmic signal.

Using a formal model, we show that the optimal design of an algorithm, such as what features to include or exclude, depends on whether the decision-making system is one of automation or assistance. For example, excluding information on protected characteristics from an algorithmic process may fail to reduce, and even increase, ultimate disparities when there is a biased human decision-maker. Even when an algorithm itself satisfies certain fairness concerns, human decisions that rely on algorithmic predictions may themselves introduce bias.
This result provides further support for a more nuanced approach to how algorithmic inputs relate to desirable outcomes and the need to avoid what I call the "Input Fallacy."

And when the corporation, even the megacorporation, burns hydrocarbons, or finds them, refines them and sells them, it harms the environment, but with almost no negative impact to the corporation. The harms are spread over society generally and not borne primarily by the firm's stockholders, its executives, or its employees.

"Oversight mechanisms should be designed appropriately to consider the combined impact of algorithmic predictions and human decisions"

The disconnect between the regulatory requirement for "humans-in-the-loop" and oversight requirements that focus on algorithms in isolation is therefore problematic. It remains an open question whether requiring human oversight in the form of human decision-making discretion and authority is optimal or fulfills its intended purpose. Regardless of whether having a human-in-the-loop is desirable, when we consider the impact of an algorithm we must be sensitive to how it is implemented. When regulation requires that algorithms act as decision-aids to humans, oversight mechanisms should be designed appropriately to consider the combined impact of algorithmic predictions and human decisions. AI policy and guidance should therefore require impact assessments and monitoring of the decision-making system as a whole and not merely the algorithmic component of the decision.

By Talia Gillis, Associate Professor of Law and Milton Handler Fellow at Columbia University.

Teaching old tricks to brand new dogs

Sarah Green
Law Commission of England and Wales

It is not so much that new technology poses novel challenges to established governance paradigms, but that it poses challenges to those who must implement them. The tools are the same, but their application will be different. Their implementation, therefore, requires an open and progressive mindset, and willingness to ask "Why not?" instead of "Why?".
Take, for example, the crucible of corporate existence (in both the pre and post Distributed Ledger Technology worlds): the contract.

Historically the epitome of human interaction, it is now increasingly an arrangement achieved by algorithmic means. Given contract's legal foundation as a "meeting of minds", it would be easy to jump to the conclusion that the law of machine contracts must be very different to the human kind. Not so. As the Law Commission of England and Wales' recent analysis found, conventional contract law is easily capable of accommodating both smart contracts and smart legal contracts - Smart contracts | Law Commission, Digital assets | Law Commission. As long as the lawyers are on board. Machines are, after all, simply vehicles of human expression: they will make agreements only where, when and how they have been instructed to do so: the autonomy remains with the instructing party (as has long been accepted in relation to vending and ticket-issuing machines, for example). An automated offer is no less a commitment to be bound for being made in digital form. And, more significantly, the need for such an offer to exist, alongside a corresponding acceptance, is as real as ever.

It's difficult, therefore, to imagine how the rules of such engagement could change, or to make the case that they should do so. There is no question that the board looks different and the pieces move in an unfamiliar way. But when IBM developed Deep Blue, allowing artificial intelligence to take on, and beat (even Grandmaster) humans at chess, nobody suggested that the rules of the game needed to change - IBM100 - Deep Blue. In fact, to have done so would, of course, have defeated the whole purpose of the exercise. It was the players who had to adapt their preparation, strategy and behaviour (in effect, often by playing deliberately sub-optimal moves that an artificial intelligence would, at least initially, find more difficult to anticipate and counter). This is the principal challenge to all aspects of governance in a world in which tech holds increasing sway: teaching old tricks to brand new dogs. Or doges.

That is not to deny that change is both necessary and inevitable in governance terms: but what is required is of a different order to the fundamental restructuring that is often expected (and dreaded) in response to technological development, and to distributed ledger technology in particular.

"What is needed is selection rather than invention. The law needs to draw analogies with existing organisational paradigms and identify the closest fit"

The thing that will shift in a world of greater automation is the topography of the contractual landscape; the rules will stay the same, but the patterns of their use will change. There is, for instance, no such thing as a recalcitrant computer (although it might sometimes feel like there is), so the enforcement of performance is less likely to be a pressing issue.

The flipside of this, however, is that defective performance is likely to be more widespread, meaning that remedies aimed at correction and restoration will be sought more and more often: rectification, in particular, is likely to be called on increasingly (or, rather, a form of rectification, which effectively sees a smart contract coded to alter the effects of a previous version, seeing as it will not be possible to alter the original code itself).

The locus of liability will also spread. Achieving automation will often mean adding another party to the chain of contractual command in the form of a coder. And whilst successful automation reduces the risk of error in (properly instructed) contractual performance, coders are no more insulated from the risk of error than any contracting party or legal adviser. Mistake, misrepresentation and negligence, for instance, will not change, but will cast their net wider (subject to the interesting wrinkle of public interactions with publicly-deployed code).

These challenges are all germane to the administration and governance of Decentralised Autonomous Organisations (which are essentially collections of automated instructions, agreements and potential agreements) - Decentralised Autonomous Organisations (DAOs). This is, of course, smart contracting, smart legal contracting and contracting in aggregate, and the concerns of users are aggregated too: they seek not individual recognition, enforcement and protection, but the assurance that their co-ordinated endeavours will be treated in a legally effective and coherent way.

But here, once more, what is needed is selection rather than invention. In setting out the conditions under which DAOs function, and the protections which are afforded to their creators, the law needs to draw analogies with existing organisational paradigms and identify the closest fit. Or, more likely, fits. Because there is no reason to suppose that DAOs will be any more homogenous as a class than conventional organisations.

The last decade has seen technological developments that are nothing short of tectonic in terms of their implications for human interaction on a collective, distributed and permissionless basis. In responding to this, common lawyers will not need to abandon what they know. But they will need to adapt that knowledge: when playing against the artificially intelligent, that is the really intelligent move to make..

By Sarah Green, Commissioner for Commercial and Common Law at the Law Commission of England and Wales.

Data and welfare in credit markets

Anthony Lee Zhang, Constantine Yannelis, and Mark Jansen
University of Chicago
University of Utah

Data is an increasingly important topic in corporate and consumer lending. The explosion of data available to screen and score borrowers over the past half century raises important questions about how borrowers are affected by the use of such data. Moreover, policymakers have implemented many regulations that limit the kinds of data that lenders can use in making lending decisions. Is society better off when data is removed from markets?

Many regulations govern the use of data: for example, the EU's General Data Protection Regulation (GDPR), the FTC's Fair Credit Reporting Act (FCRA), and California Consumer Privacy Act of 2018 (CCPA). Our paper provides insights into who wins and loses when policymakers prevent markets from using data to make lending decisions. We construct a framework to quantify the welfare effects of regulatory policies that govern data usage in credit markets. We think of data policy as a form of price discrimination in which a company charges different prices to different consumer groups--similar to a museum offering discounts to students and charging higher prices to non-students. Similarly, data usage in financial markets can be thought of as a kind of third-degree price discrimination.

Our study develops an example to illustrate this. Consumers who've been bankrupt in the past tend to be at higher risk of defaulting on a loan than consumers who have never been bankrupt. When lenders can tell formally bankrupt consumers apart from non-bankrupt consumers, they will set higher prices (i.e., higher interest rates) for previously bankrupt consumers, and lower prices for never-bankrupt consumers consistent with each consumer's likelihood of repaying the loan.

Should society allow lenders to observe bankruptcy "flags" for previously bankrupt individuals, and set different prices for previously bankrupt and never-bankrupt borrowers?

Now suppose that a policymaker wants to help formerly bankrupt consumers. Since redistributing cash directly is often difficult, we might think of deleting bankruptcy information as an alternative: we redistribute through subsidized prices in credit markets by preventing lenders from using this information. In fact, the FCRA requires that flags indicating the occurrence of consumer bankruptcy be removed after seven (ten) years for a Chapter 13 (Chapter 7) bankruptcy. In such a case, the lender will not be able to distinguish formerly bankrupt borrowers and those that have never declared bankruptcy.

This loss of information has two effects on borrowers. First, the removal of bankruptcy information effectively transfers surplus from never-bankrupt consumers to previously bankrupt consumers. Second, we show that society as a whole is better off with more data: social welfare decreases since the prices no longer reflect the risk of the borrowers. This contrasts with the setting in which lenders have access to the borrower bankruptcy history and they set the prices to reflect the cost of lending to each borrower type, thus making credit allocation more efficient.

An important finding of our study is that the transfer effects of data availability in the case of bankruptcy data availability tends to be large relative to the welfare effects. That is, the social welfare gains of having the additional information is small relative to the size of the transfer between the two borrower types. Thus, if the policymaker is interested in redistributing wealth to previously bankrupt consumers, preventing the use of bankruptcy data helps to achieve this.

Given that there are many ways in which a policymaker can transfer wealth to segments of the population, it is important to consider the cost of such a transfer. How much welfare is lost in this setting?

When we apply our framework to the US auto lending market and find that for every dollar transferred to previously bankrupt consumers only 3 cents of social surplus are lost. This implies that the removal of bankruptcy information is a fairly efficient way to transfer wealth to previously bankrupt consumers. While this example is focused on a particular consumer data setting, our framework can be easily applied to many corporate and consumer lending environments. The framework is simple to implement empirically: to measure the welfare and transfer effects of data availability or policy change, one only needs to observe the pre- and post-data prices (interest rates in our earlier example) and quantities (loan size in the earlier example) for each consumer group.

In summary, our paper introduces a framework for thinking quantitatively about the welfare and distributional effects of data policy, which can be taken to data pretty easily. We hope others use this method to analyze the welfare effects that data policies have in other settings.

By Anthony Lee Zhang , Assistant Professor of Finance at the University of Chicago's Booth School of Business, Constantine Yannelis, Associate Professor of Finance at the University of Chicago Booth School of Business and Mark Jansen, Assistant Professor of Finance, University of Utah

Fintech and innovation in emerging markets: some common challenges and opportunities

Enmanuel Cedeño Brea
Superintendency of Banks of the Dominican Republic

Technology continually transforms the way that firms carry out their business. Financial firms have not been the exception. The rise of Fintech — which is the provision of financial services using new technologies—is a case in point on how innovation has spilled over to all stages of the commercial lifecycle. Tech has permeated how financial firms conduct their business, comply with regulations (Regtech), and even how they interact with regulators and supervisors (dubbed 'Suptech').

Fintech has the potential to make people's lives better—especially those who have been financially excluded or underserved by conventional services providers. Fintech's promise could impact access to credit, savings and investment, subsidy assignment, payments, insurance, and other financial services. Despite its promise, digital transformation poses multidimensional governance challenges for financial services providers and their many stakeholders.

One of the dimensions relates to how technology affects the interaction between services providers and their consumers. Financial firms increasingly interact with their customers through non-face-to-face ('NFTF') channels –which include digital settings. This means that all stages of customer lifecycle – from onboarding to relationship closing—can be done online.

Despite increased digitalization, financial services providers still need to conduct adequate customer due diligence for managing their Anti Money Laundering risks (AML/FT). In addition, they need to verify customers' identities to prevent other forms of financial crime, like fraud. Technology provides many solutions to the age-old problem of NFTF interactions between banks and their customers. However, they also pose some challenges.

Firstly, countries are in different stages of technological adoption and digital infrastructure. Some jurisdictions have robust legal frameworks for e-commerce and have also developed digital identity utilities and systems based on state-of-the-art tech. However, emerging markets are typically lagging on these features. How can these countries foster Fintech without cutting any corners nor stalling further behind the implementation curve?

While lessons from leading jurisdictions show the potential that digital national ID programs (e-ID) have, India's success story with the Aadhaar ID program has been a beacon for many emerging economies.

.

Other similar e-ID programs provide evidence that show positive spillovers to key areas outside of the financial sector, including education, government transfers, health, security, and migration, can more than justify their outlays. While having full-fledged e-ID infrastructure would be ideal, some emerging economies might face budget constraints or collective action problems that prevent such programs from reaching fruition.

A second-best solution can be enabling regulation creating a legal framework for digital onboarding and eKYC –based on technological neutrality and a risk-based approach—which can become a cost-effective alternative to e-ID.

.

Risk-based and technological neutrality coupled with following the 'same risks, same regulation' adage, can help regulators design better frameworks that remain relevant for longer. Since rules tend to stick, getting it right from the outset (or at least as best as possible) can be beneficial in the long run.

Digitalization in finance is no panacea. While Fintech solutions can be leveraged to attain greater levels of financial inclusion, digital illiteracy can become a stumbling block. Combined with other existing educational limitations and lacunae, digital illiteracy could lead to new forms of exclusion.

Another key challenge is maintaining a proper conduct of business across NFTF interactions. Supervisors must pay greater attention to Apps, platforms, dispute management systems and e-banking sales funnels, as customer consent to new products, provide their data, make choices, and resolve disputes online. Behaviorally informed consumer protection regulation can be an important supervisory tool for supervising the NFTF choice architecture, helping to tackle biases and other behavioral limitations, shedding light on so-called 'dark patterns' and navigation labyrinths, all in the name of preventing consumer exploitation.

Finally, there are also opportunities for financial regulators and supervisors to embrace Fintech, beyond Regtech and Suptech. Tech can help supervisors provide users with meaningful information and smart disclosure that allow them to make better choices. One example of such an initiative at the intersection of Suptech, Wealthtech and Govtech is the Superintendency of Banks of the Dominican Republic's flagship user protection application called ProUsuario Digital.

The App allows users to access their loan history and asset classification score by displaying data from the centralized risk databased that Dominican banks share and use. The app also allows users to place and track complaints directly with the Superintendency's consumer protection department. Future developments will incorporate smart disclosures to help users make behaviorally informed choices and receive meaningful bespoke information. Less than a year after its launch the app already boasts more than 100,000 active users.

By Enmanuel Cedeño Brea, Deputy Manager for Regulation and Innovation at the Superintendency of Banks of the Dominican Republic.

Smart financial contracts as a mixed blessing

Gerard Hertig
Singapore-ETH Centre and ECGI

Technology increasingly plays a governance role. Hence, smart contracts (a term coined by Nick Szabo) are progressively shifting decision-making from the real to the digital world. On the upside, they facilitate contract drafting and contribute to decision-making; on the downside, smart contracts may drift over time. In any event, major financial centers are increasingly recognizing their importance, especially due to compliance and security issues being dealt with via distributed ledger and 'hashgraph' technology.

Smart contract use is likely to facilitate agreements, especially among non-trusting operators: the parties' input is mostly limited to agreeing to the code governing their interactions. To be sure, contractual terms may not be fully grasped by everyone; however, this is also the case for standard financial contracts. More importantly, the gap between what has been subjectively agreed upon and what is effectively delivered is likely to decrease over time. It follows that differences in smart contract practice will have more to do with platform compatibility than with dissimilarities across legal regimes. In other words, enforcement consistency is a function of technical rather than regulatory harmonization.

"Automation will not prevent biases: some design choices are superior to others, and algorithms are not impartial"

Current transaction systems enable every participant to contribute to decision-making. However, a 100% horizontal system has yet to emerge. Consensus occurs via an underlying network of computer nodes. Changes cannot be altered, a property that is ideally suitable for situations where data is shared between multiple participants—i.e. in the financial sector. Every transaction is subject to node-specific conditions and verifications; it must bear the users' digital signature and will get a unique ID. Completed transactions are stored in blocks, with indication of amount and time; each block has its unique hash, which represents the conversion of an input of arbitrary length into an encrypted output of fixed length. The parties can see their transaction once it is added to the ledger; depending on the network's characteristics, others may also be able to see it. Recorded information cannot be altered or recovered. Practice consistency is currently facilitated by mimicking the 'international private law' regime applicable to traditional contracts. Ideally, an international agreement would provide for the mutual recognition of smart contracts. Enforcement could occur under either uniform rules or—which could prove more practicable - the rules applicable in the defending party's country of 'residence'.

Whether smart contracts will prove welfare enhancing remains to be established. To begin with, automation will not prevent biases: some design choices are superior to others, and algorithms are not impartial. More importantly, smart contracts may drift over time, especially when there are deficiencies in the provision of new training data. A drastic way to address the issue is to rely on self-destruct functions; a more nuanced approach is to limit the 'gas' provided to process contractual instructions.

The importance of smart contracts is increasingly recognized by major financial centers. The trend is especially noticeable in leading European jurisdictions. England, which has a robust smart contract tradition, is currently modernizing its 'emerging technology' conflict of law provisions. Germany and Switzerland foster smart contract use by letting one party write the contract and the other agreeing to the code displayed in a front-end application. By contrast, the French approach is more restrictive, subjecting smart contracts to automation and security requirements. In the US, the current view is that smart contracts are generally enforceable provided they fulfill basic contractual requirements. It follows that traditional contractual frameworks apply.

Compliance with regulatory requirements is facilitated by reliance on distributed ledger technology (DLT). In public DLT, everyone can participate in decision-making and transactions are visible to all; under this approach (which is adopted by Ethereum), data cannot be modified post validation and acceptance. By contrast, only those given permission can access private DLT; under this approach (which is adopted by Hyperledger), data is highly secured and kept confidential. Recent technology allows for relying on agent-centric rather than data-centric DLT approaches. For example, nodes run their own chain at Holochain, thus operating independently while remaining part of a larger network.

Transaction security can be reinforced by relying on hashgraph- instead of DLT-technology. Here, parallel transaction storage results in multiple transactions getting the same 'time stamp' and being provable within minutes. Alternatively, the parties may use the directed acyclic graph (DAG) approach. Here data enters a processing element through the incoming edges and leaves it through the outgoing edges. This enables the validation of randomly chosen transactions, thus providing a new way to improve scalability. These developments (and the still complex nature of distributed databases) provide some protection against hacking. Nevertheless, malicious actors regularly carry out hacks. This is often attributed to the persistence of vulnerabilities, especially when it comes to regular user interactions with the system. On the other hand, it may also be due to rewards being higher for hackers than for security agents.

Summing-up, smart contracts have pros and contras: they facilitate contract drafting but may drift over time. Going forward, they are likely to play a significant practical role, especially when it comes to compliance and security issues.

By Gerard Hertig, Principal Investigator, Future Resilient Systems Program, Singapore-ETH Centre, co-founder of the Center for Law & Economics at ETH Zurich, and an ECGI Fellow.

Regulating auditing algorithms: An Asian solution?

Nydia Remolina
Singapore Management University

Artificial intelligence (AI) is increasingly impacting our society and economy by undergirding, digitizing, and automating important processes across many sectors, such as loan applications, medical diagnoses, hiring decisions, surfacing information, and driving autonomous vehicles. With such important outcomes for individuals, organizations, and society, it is critical that AI can be trusted to make fair and accurate decisions. However, information asymmetries between organizations, consumers and regulators are common, and they are exacerbated by the inherent complexities of algorithmic systems such as machine learning, deep learning, and black box algorithms. Therefore, these information asymmetries can ultimately be exploited by certain organizations, and they can reasonably create skepticism and mistrust on the use of algorithms.

As a response, algorithmic auditing has emerged as a possible solution to address this information asymmetry and create trust by ensuring that a system is reviewed, prior to and during deployment, by a third party with suitable specialist knowledge who can convey information about the impact of the system to other parties. AI audits can take different forms, from checking governance mechanisms, to testing an algorithm's outputs, to inspecting how AI systems are designed. Audits can be undertaken by external parties appointed by organizations using AI, or by regulators, researchers or other parties carrying out an audit of a system on their own initiative.

While the objectives that inform financial audits, deeply explored in the law and finance literature, may translate to AI in the sense that a financial auditor gathers and inspects evidence to determine whether a company's practices are free from material misstatement and the company's financial statements have been prepared according to generally accepted accounting principles, an AI auditor faces more complex challenges including examining design documents, code, and training data to determine whether a company's algorithms are free from material potentially consequential impact.

As AI becomes more sophisticated and broadly used, algorithmic auditing also involves increasingly complicated ethical, social, and regulatory challenges with different dimensions and implications depending on the sector where the AI is deployed. In this context, regulators play a key role in proposing policies to govern the operations, credentials, and impact of the experts conducting AI audits. Additionally, when they act as independent (external) parties, there are no rules governing what – in the context of financial audits – have been named the auditors' independence puzzle.

The predicament remains daunting, and the nascent ecosystem of external and internal algorithmic auditing is growing in a fragmented manner, without proper scrutiny, guidance, or consistency. The current AI audit landscape still lacks specific rules and standards, and auditors are offering auditing services without assurance of their quality, and the scope and expectations of their review. That can make audits a costly exercise that is not necessarily translated into a higher level of trust and consumer protection. Moreover, it can lead to inconsistencies. One example of the inconsistency that AI audits are plagued with is the apparent international agreement about ensuring fairness in AI implementation. However, when it comes to developing actual evaluations of fairness to audit algorithms, there are numerous statistical definitions of fairness that are often mutually exclusive or that do not match with legal standards of fairness and discrimination. Regulators are expected to propose solutions for this patchwork in the world of AI audits.

Currently, some regulators are showing interest in assessing the performance of AI systems. The European Union has proposed the AI Act, a risk-based approach to AI regulation. The proposed regulatory initiative establishes that for certain systems – High Risk AI systems – an external notified body will be involved in the conformity assessment audit. Likewise, in February 2022, a group of U.S. Senators proposed legislation for an Algorithmic Accountability Act that would have required the Federal Trade Commission to gather reports on algorithms and scrutinize their functions.

Nonetheless, these proposals have not been implemented and the principle-based approach of the EU AI Act has not been grounded in specific assessment measures to be implemented in practical use cases. In contrast, regulators and authorities in Asia are not proposing a regulatory framework for AI but are among the first movers in taking actionable steps towards building AI audit systems. These initiatives could shed some light on what regulators in different sectors could focus on for developing an adequate algorithmic auditing ecosystem.
For instance, China's internet watchdog released some details of how AI systems work inside Chinese tech companies. The Cyberspace Administration of China required the 30 largest domestic tech companies to share AI systems' information without having to publicly disclose intellectual property. As a result, it is known now that, for example, ByteDance's algorithm takes likes and dislikes into consideration when serving personalized and targeted content on Douyin, the Chinese version of TikTok. The agency published similar information about companies such as Alibaba and Tencent and required additional information that was not released to the public, including a self-appraisal on the security of the algorithms, the data they collect, whether that encompasses sensitive biometric or identity information, and what data sources are used to train algorithms.

Authorities in Singapore have also developed actionable tools to build a less fragmented AI audit ecosystem. Namely, Singapore's Infocomm Media Development Authority and the Personal Data Protection Commission (PDPC) launched AI Verify, the world's first AI Governance Testing Framework and Toolkit for companies that wish to demonstrate responsible AI in an objective and verifiable manner. It verifies the performance of an AI system against the developer's claims and with respect to the accepted AI ethics principles. The toolkit is a software package that can be downloaded and executed locally in business environments to generate testing reports for engineers and management. Even though AI Verify is not mandatory, 10 companies from different sectors and of different scale, have already tested and provided feedback to the initiative. These companies include Amazon Web Services, DBS Bank, Google, Meta, Microsoft, Singapore Airlines, NCS (Part of Singtel Group)/Land Transport Authority, Standard Chartered Bank, UCARE.AI, and XoPA.AI.

"The nascent ecosystem of external and internal algorithmic auditing is growing in a fragmented manner, without proper scrutiny, guidance, or consistency"

Also in Singapore, the financial regulator, the Monetary Authority of Singapore, worked with technology companies and financial institutions to design an assessment methodology for AI use cases in the financial sector, such as algorithmic credit scoring, that are broadly used in the industry or create risks that should be mitigated as a priority. As part of Veritas, MAS released the first version of an open-source software toolkit that aims to drive financial institutions' adoption and adherence to AI governance principles. The software enables the automation of metrics assessment and visualization, with plug-ins integrating with financial institutions' IT systems.

These initiatives are the first of its kind globally. The documented early lessons of these authorities should provide regulators with tools that can assist in providing answers to legal and regulatory challenges regarding the standards applicable to AI systems, the scope of auditing systems, the role of auditing systems in compliance, and the translation to ethical principles into actionable and measurable characteristics. These efforts are also pivotal in shaping the AI audit ecosystem as a balance of public and private actors and processes. As such, the role of regulators is crucial for achieving certainty and meaningful AI audits that truly contribute to create trustworthy AI.

By Nydia Remolina, Assistant Professor of Law, Singapore Management University, Fintech Track Lead and Head of Industry Relations, SMU Centre for AI and Data Governance

The rise of alternative data: AI governance and ethical challenges

Zofia Bednarz
The University of Sydney Law School

It has been said that 'data is the new oil'. It was to indicate data is an asset, a powerful resource allowing organisations to optimise their business models and increase profits. But it's beginning to look like the analogy is also true in relation to risk and liability linked to collecting and using data: the last few months in Australia looked like the BP Deepwater Horizon and Exxon Valdez disasters combined, with Optus (second largest telecom company in Australia) and Medibank (one of the largest Australian private health insurers) data breaches. Up to 20 million people – current and past customers – are reported to have been affected. It provoked a bit of a knee-jerk reaction from the government, which quickly proposed legislation increasing penalties for data breaches and enhancing regulators enforcement powers.

Cybersecurity breaches, while onerous for both companies and customers affected, are far from being the only risk for organisations involved in the data economy. While governments and consumers tend to focus on problems such as identity theft (no doubt an important issue), new legal and ethical challenges arise from the increasingly ubiquitous collection and use of alternative data. Such non-traditional, new streams of data used can include, for example, social media, internet browsing history, smartphone apps, location history, customer loyalty schemes, fitness trackers, smart home devices and so on.

Powerful data analytics tools, enabled by AI and related technologies, such as machine learning, make it possible to analyse and learn from all that data, generating inferences and predicting trends that would be otherwise unobservable to humans.

And while these insights may be of high commercial value, they bring about legal and ethical challenges. Consumers are often unaware of the ongoing omnipresent alternative data collection, aggregation and combining. Datasets which are supposed to be de-identified or anonymised, are often easily re-identifiable. AI models used to analyse the data are often opaque black boxes, which makes explaining and potentially challenging their decisions difficult. The predictions of the models may be inaccurate, adversely affecting often the most vulnerable consumers. Not to mention the well-described problem of algorithmic discrimination and bias perpetuated by the models and embedded within the data they are fed.

Still, the discourse about alternative data is mainly that of benefit and opportunity – a necessary component of fostering AI innovation endorsed by companies and governments. McKinsey consulting firm hails 'harnessing the power of external data' noting how: 'few organizations take full advantage of data generated outside their walls. A well-structured plan for using external data can provide a competitive edge.'

Companies boast how AI insights allow them to offer personalised services, 'tailored' to individual consumer's needs. Policymakers also promote 'innovation', and encourage data collection, for example through open banking schemes. The aim of open banking is to give consumers the ability to direct companies that hold financial data about themselves to make it available to financial (or other) companies of the consumer's choice.

In practice, it makes it possible for organisations to get access to consumers' information they could never get from a consumer directly, such as for example their transaction data for the past 10 years.

Financial industry is in particular keen to take advantage of alternative data. 'All data is credit data, we just don't know how to use it yet' is a famous statement summing up the industry thinking. This in turn means that firms are eager to collect any and all data: after all, it may signal the value of the client, improve business processes, and it's encouraged by industry consultants and even legislators.

However, it looks like a reckoning is coming. Consumers are becoming increasingly aware of data surveillance and dodgy data practices. Reputation risk is becoming something the organisations need to start caring about. Failing to comply with privacy and data protection rules attracts investigation (see for example CHOICE's work on Australian retailers using facial recognition tech) and penalties, as the Clearview AI case showed. Public opinion pressure is mounting on the policymakers to regulate the use of AI models and people's data.

Yet it still looks like many organisations believe hiding their data practices behind unclear or misleading privacy policies is the way to go. It should go without saying that it's not adequate risk management, but quite the opposite: it's potentially illegal, certainly unethical, and exposing companies to risk. Responsible, ethical and transparent AI and data governance should be key to anticipate and prevent the risks.

"Many organisations believe hiding their data practices behind unclear or misleading privacy policies is the way to go"

'Machinewashing', where companies try to convince stakeholders and regulators that their internal AI and data governance operates in line with human and societal values, while in reality they engage in unethical or even illegal data and AI practices, won't cut it. Stricter enforcement and new data and AI regulation (such as the EU's AI Act) is coming, and it will make transparent and ethical AI and data governance a necessity for organisations. Alternative data, as much as is an asset and a resource, is also a liability and a risk that needs to be addressed. The sooner, the better.

By Dr Zofia Bednarz, Lecturer in Commercial and Corporate Law, University of Sydney; Associate Investigator, Australian Research Council Centre of Excellence for Automated Decision-Making and Society

From TechFin to PlatFin to FinTech 4.0

Dirk Zetzsche, Douglas Arner, and Ross Buckley
University of Luxembourg
University of Hong Kong
University of New South Wales

Finance is an industry characterised by strong economies of scope and scale. These effects have led over time to the emergence of large financial groups and conglomerates. Finance is also characterised by the strong positive and negative externalities generated by its activities: finance has benefits that extend beyond its direct participants in the form of greater economic growth and development, and can also cause instability and crises. Extensive regulation has therefore emerged to control entry into the industry and reduce the likelihood of failure, with the objective of increasing positive externalities and reducing negative ones.

Over time, this combination of factors has inexorably resulted in concentration within the financial sector, posing risks to competition, efficiency and financial stability especially in the context of systemically significant financial institutions. These factors were core to the 2008 global financial crisis and the focus of post-crisis regulatory reforms.

At the same time, while technology and finance have always developed together, digitisation and the application of new technologies has transformed finance, as now embodied in the term FinTech. Much of the focus of the 2010s was on the opportunities and challenges posed by new entrants, in particular FinTech startups using technology to challenge incumbent approaches and models. The business model of these FinTechs (perhaps with certain exceptions in decentralised finance) was to focus on achieving scale.

Like finance, technology benefits from network effects, which manifest with increasing numbers of customers, interconnections and data. Network effects in technology inexorably lead to the emergence of dominant technology platforms as we have seen in both the United States and China.

We highlight this process in the context of the asset management industry in one of our papers, and explain the emergence of concentration from a technological perspective. For some years, scholars have discussed the impact of the "rule of the twelve" on corporate governance. Few, however, have looked into why only a mere dozen institutional investors have such a large impact on listed firms. We argue that technology, and in particular the resulting access to clients' data and liquidity, allows for utterly unprecedented economies of scale in finance. In fact, very large firms that combine features of both finance and technology have developed from the combination of technological evolution (digitisation, datafication, digitalisation), conducive regulatory approaches, and the pro-concentration effects that characterise data and financial industries.

We also see this in the context of the entry of BigTech platforms into finance and the scaling of FinTechs as some move quickly from being from too small for regulators to care about to being too-big-to-fail. These trends characterise the current period in the evolution of FinTech, which in another recent paper we dub FinTech 4.0. We see the decade of the 2020s in finance will be dominated by a massive battle between centralisation and decentralisation, of seeking the positive externalities of data aggregation and finance while at the same time reducing the negative externalities of change at pace that at times bewilders regulators.

Conventional economies of scale, data-driven economies of scale and network effects explain why FinTech markets can increase efficiencies in processes such as client onboarding, compliance, and reporting, and also allow for algorithm-based investment and trading that use market information at a greater pace than humans ever could.

The same technological advances that account for cost reduction and increased efficiencies also potentially contribute to decreasing competition in FinTech markets. On the one hand, conventional FinTechs (including trading platform Robinhood and the archetypal FinTech firms Aladdin and Ant Group) have been able to collect billions of assets and millions of customers. On the other hand, these very same forces have enabled new entrants into financial services like Meta (formerly Facebook), Apple, Google, Microsoft, and Amazon ("MAGMA") in the United States, and Baidu, Alibaba, and Tencent ("BATs") in China, to extend broadly across most aspects of the society and economy within their respective countries and beyond.

The data-driven finance business is a platform industry. In the digital finance context, the term "platform" refers to a systems architecture where multiple applications are linked to and through one technical infrastructure so that users can use one major integrating software system to run all applications written for that system. One outcome of platformisation is financial ecosystems with multiple services linked to clients via one platform, with the platform serving as the indispensable technical core that ties all services and clients together, but also provides some services itself.

"The decade of the 2020s in finance will be dominated by a massive battle between centralisation and decentralisation"

How can this be done? We present a framework of analysis and highlight the challenges to existing regulatory silos in finance, competition / antitrust, data, and technology regulation. We argue the need for a balanced proportional approach, enabling and encouraging competition and innovation whilst also carefully monitoring emerging scale, concentration and dominance issues.
Key to this approach is data regulation and the necessity of focussing on building systems to enable data aggregation while limiting the extraction of monopoly rents by dominant private players, whether in finance, in tech or in FinTech/TechFin markets.

The evolution of financial ecosystems can bring many advantages. Yet, the rapid emergence of concentration and dominance in digital finance via platformisation can pose great risks; the corporate governance aspects discussed in recent ECGI scholarship is only one of them. Financial regulators worldwide must step up to the difficult challenge of dealing with these transformative contemporary market structures in a sufficiently firm, yet balanced and proportionate, manner.

By Dirk A. Zetzsche, Professor of Law and ADA Chair in Financial Law at the University of Luxembourg, Douglas W. Arner, Professor in Law and RGC Senior Fellow at the University of Hong Kong and, Ross P. Buckley, KWM Professor of Disruptive Innovation and Australian Research Council Laureate Fellow at UNSW Sydney.

Corporate governance for the responsible use of AI

Souichirou Kozuka
Gakushuin University, Tokyo

Since the use of AI (artificial intelligence) has become common, there have been various efforts to establish good governance for the use of AI. A well-known initiative of early days is the adoption of Asilomar principles by the Future of Life Institute, to which numerous researchers and companies signed up. The European Union and Japan are the two jurisdictions where the government and industry collaborated in producing frameworks for the responsible development and use of AI. Both the European Ethics guidelines for trustworthy AI and Japanese AI Principles (AI R&D Principles and AI Utilization Principles) emphasise that the use of AI should be human-centered, not leading to a dystopia where the AI controls humans by the algorithm. The idea met support globally, leading to the adoption of the OECD AI Principles for the use of AI..

Interestingly, jurisdictions have diverged in regulatory approaches after the global recognition of the principles for human-centered AI. The European Union have worked on the Proposal for a regulation laying down harmonised rules on artificial intelligence, with the belief that a legally binding instrument is necessary to supplement voluntary commitments through the soft-law. On the other hand, Japan now focuses on the effective implementation of the AI principles by the industry. It published the Governance Guidelines for Implementation of AI Principles and facilitates sharing of best practices among the industry members, making use of the forum that formulated its AI Principles. Similarly, the Singapore government published the Model Artificial Intelligence Governance Framework to be referenced by organisations that deploy AI.

The latter approach focusing on the implementation gives rise to the agenda of "corporate governance for the use of AI".

One may question why AI must specifically be addressed, as distinguished from other kinds of technologies. An important feature of AI is that it makes decisions in a black box. The AI commonly used today is based on deep learning, which is a technology to identify a hidden correlation through machine learning of data. While it is helpful in discovering what a human can hardly recognise, a human finds it difficult to review how AI reached such a decision, still less to control it. Furthermore, the AI continues learning after the system is delivered from the developer to the user, making its decisions even less controllable.

As a result, the consumer cannot be convinced that the decision allegedly made by AI is not manipulated by the provider of the service deploying AI, which could lead to distrust in AI. Even when the public trusts the AI's decision, there is still a possibility that the decision reflects a bias unacceptable to the society, in view of the fundamental rights of the people. Such a bias can easily sneak into AI's decision when the data that the AI learns is affected by unfair practices in the society, such as discrimination by race or gender.

To solve these problems and ensure the public's trust in AI, the user of AI has to care about the explainability of AI's decisions, as well as to take accountability for them. Explainability and accountability are the key principles for the human-centered AI. However, one must admit that the perfect explainability will compromise the benefits of using AI.

If one tries to identify every element that has led the AI to make a certain decision, a huge number of parameters must be disclosed and turned into a human-readable format. Then, the advantages of using AI to find correlations not recognisable by a human and substituting AI's decision for a less efficient human decision will be compromised to a large extent. Obviously a balance is needed at some point.

The Model Framework by the Singapore Government argues that the relationship of the human and AI can be either "human-in-the loop", "human-out-of-the-loop" or "human-over-the-loop". Under the first approach, the final decision should not be left to AI but must always be reserved by a human. The second approach means that replacing the human decision by the AI's decision is allowed. The third approach requires that the human oversight must be made so that a human can step in when an unexpected incident occurs. The Model Framework argues that the choice among these three approaches must be made on the basis of the severity of the harm to be caused by a wrong decision by AI and the probability that the AI errs.

Here lies an issue that the corporate management has to decide on. They should make a decision about how (under which approach) the AI system is used, and to what extent its decision is explained. In making such a decision, they need to assess risks of deploying AI. It is also recommended that the company using the AI system adopt principles on AI of its own, adapting the principles formulated globally or in its jurisdiction to its business. It is what major developers of AI systems already do today and is useful in identifying which issues are particularly relevant to the company's use of AI. Depending on the potential risk, due diligence over the supply chain of AI system might also be required to examine how the data is collected and prepared for training of AI be exercised, because decisions by AI are affected by the data that it learns. Thus, the top management of a company using, or intending to use, an AI system in its service should build up the governance system for the use of AI within its company.

"The perfect explainability will compromise the benefits of using AI"

By Souichirou Kozuka, Professor, Law Faculty of Gakushuin University, Tokyo.

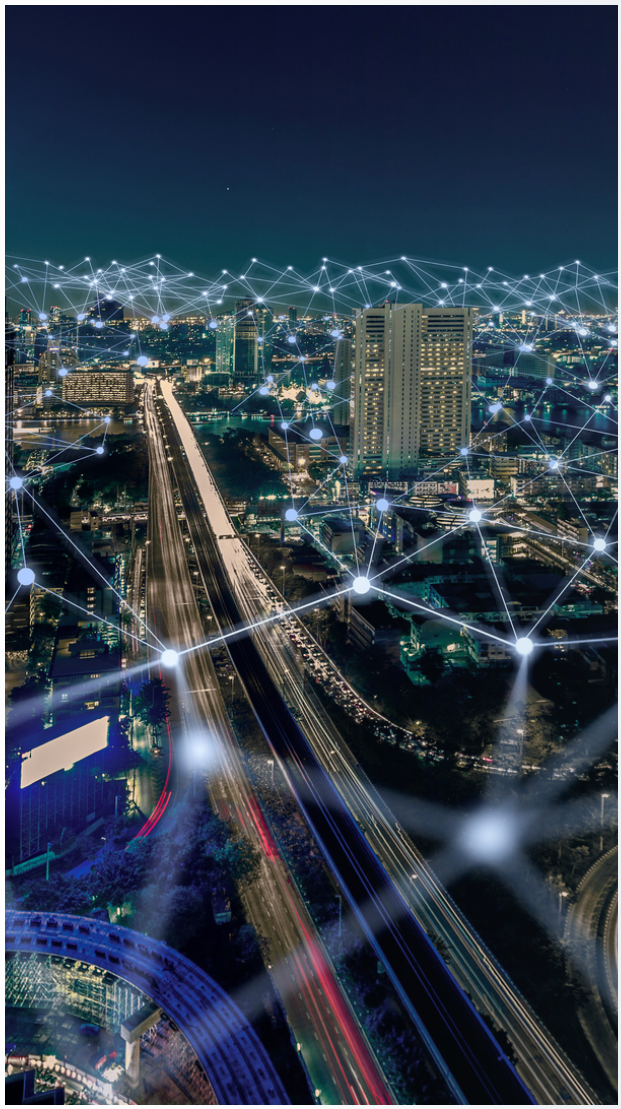

The perils of frictionless finance

Nikita Aggarwal
UCLA School of Law

At the heart of the Silicon Valley playbook is the mantra of frictionless-ness. It tells us that all products and services should be simple, seamless, smooth—without "frictions" or "pain points" for the user. This mantra has, in recent years, found its way into consumer financial markets, in a growing trend of frictionless finance.[1] In retail investment markets, popular apps like Robinhood optimize for frictionlessness—for example, by eliminating trading commissions and lengthy account verification procedures, enabling smaller investments in fractional shares, and by deploying a simple user interface and experience (UI/UX). In consumer credit markets, a key design principle of popular "pay in four" buy now, pay later (BNPL) credit products is to reduce user friction in online borrowing and payment—notably, by eliminating interest, streamlining credit-checking, and seamlessly integrating online finance and commerce (a trend referred to more broadly as "embedded finance"). BNPL is following the lead of popular online retail platforms, like Amazon, which embody the mantra of frictionless-ness in features such as "One-click" payment.

By removing upfront frictions, and reducing transaction costs, zero-interest credit and zero-commission trading apps have made it much easier for consumers to borrow and invest. No doubt, this has potential economic benefits, particularly by expanding access to financial services for previously underserved populations. BNPL is, on average, cheaper than credit card borrowing. But access to finance is not an unalloyed good. Borrowing and investing are risky activities. Consumers can lose as well as gain money, and if the losses become excessive or systemic, they can jeopardize the health of the broader financial system and economy.

"Smooth finance increases risk in the system and concentrates risk in younger, less financially sophisticated, lower-income consumers, who are less able to absorb the risk"

The aesthetic of frictionlessness exacerbates these risks. Eliminating upfront costs, such as credit interest and trading commissions, creates the mirage that borrowing and investing are costless—free. As retailers and advertisers have long understood, the human brain is hard-wired to think less reflexively and over-consume products that are free. However, these costs are not so much being eliminated as they are shrouded and shifted to other, less visible parts of the transaction. As a result, frictionless design is increasing the risk that consumers—especially younger, less financially sophisticated, and often lower income consumers—misperceive the true costs involved in financial transactions and make unfavourable financial decisions. In turn, by encouraging more risky borrowing and trading, frictionless finance both increases risk in the system and concentrates it in those consumers who are least able to absorb it. Unsurprisingly, the aesthetic of frictionlessness can be very profitable for firms.

Take zero-interest, pay-in-four BNPL. There are three ways in which the zero-interest BNPL business model shifts and obfuscates the costs of credit to the consumer. First, it shifts costs from the front end to the back end of credit transactions.

Although these products do not carry any interest charge—which is typically more salient to the consumer—they do carry late payment fees, which are typically less salient.

This type of behavioural manipulation is already familiar in consumer credit markets. The second, less familiar shift is from charging consumers to charging third-party merchants or wholesalers. At its core, BNPL implicates a tri-partite arrangement between the borrower, lender, and merchant. The BNPL lender charges the merchant a transaction fee every time a customer makes a purchase using BNPL. These fees are typically high, however merchants justify them as BNPL can increase their sales (more specifically, a larger volume of smaller transactions) and attract a newer, younger consumer base (as well as their data—see below). Similarly, zero-commission trading replaces fees charged to consumers with fees charged to wholesale market makers ("payment for order flow").

The third major shift is from monetizing the financial transaction (through the payment of interest and late payment fees, or trading commissions) to monetizing the data transaction (notably, for lead generation). To paraphrase Richard Serra, if something is free, you are the product. These latter transformations strengthen the incentives of merchants and lenders, in the BNPL context, and broker dealers and market makers, in the retail investment context, to increase transaction volumes—in turn further misaligning their incentives with the best interests of consumers, and the economy, in responsible borrowing and investment.

In a new paper, #Fintok and Financial Regulation, my co-authors Christopher Odinet and Bondy Kaye and I investigate the risks of BNPL using a novel dataset of TikTok videos in which creators discuss their experiences with Klarna, one of the largest providers of BNPL. Although our study is not dispositive, it conveys worrying signals about consumers' misperceptions of the true cost of credit and resulting over-indebtedness due to BNPL, particularly among younger consumers. More broadly, zero-commission trading and the "meme stock" phenomenon has increased stock-market volatility. The shift under frictionless finance to a data monetization business model also introduces new financial and non-financial risks due to the misuse of personal data.

Given the inherently behavioural motivations of frictionless finance business models, a natural locus for regulatory intervention would seem to be the design of digital financial apps—more particularly, their frictionless design. Indeed, regulators—such as the SEC, CFPB, and FTC in the US, and the FCA in the UK—are increasingly alert to the role of deceptive digital design, including in consumer financial markets. Although it is inevitably a fine balance, some online friction that forces consumers to reflect more carefully on risky financial decisions is likely to be both economically and socially valuable. Architecture-based interventions in this vein could include: a requirement that BNPL is not a default payment method on retail platforms (an approach taken by regulators in Sweden); making information about late fees more prominent on the BNPL lender's app or website; and removing or de-prioritizing statements that these products are "interest free."

The political feasibility of these design-based regulatory interventions will vary between jurisdictions. Particularly in the US, they could raise objections on grounds of paternalism and infringement of individual constitutional freedoms. More broadly, design-based interventions, as with other forms of technological regulation, carry the risk of obsolescence and under-inclusivity. As such, it is important that design-based interventions are evaluated alongside, and complemented by, more traditional conduct-based interventions. This includes requiring BNPL credit providers to ensure that the claims they make to consumers are not misleading or deceptive, to carry out appropriate creditworthiness (ability to pay) and identity verification assessments, as well as encouraging them to share BNPL credit data with credit bureaus to facilitate a more complete assessment of a consumer's ability to pay. In retail investment markets, potential regulatory interventions include increasing transparency and price competition in equities markets, and strengthening the duties of care of broker-dealers.

By Nikita Aggarwal, Postdoctoral Research Fellow at UCLA's Institute for Technology, Law and Policy.

Economic and normative implications of algorithmic credit scoring

Holli Sargeant
University of Cambridge

Commercial use of artificial intelligence (AI) is accelerating and transforming nearly every economic, social and political domain. Companies have been attempting to classify and label items, processes, and people for a long time. The modern convergence of foundational technologies, however, enables the analysis of vast amounts of data required for AI. Consider how much data is created and used based on our online behaviours and choices. Converging foundational technologies enables analytics of the vast data required for machine learning. As a result, businesses now use algorithmic technologies to inform their processes, pricing and decisions.

Lending has been identified as a high risk for discriminatory algorithms where using historical data that will result in biased algorithmic tools. Bias, among other risks, is an essential consideration. However, there is a gap in recent literature on the potential optimal outcomes that can arise if risks are mitigated. Algorithmic credit scoring can significantly improve banks' assessment of consumers and credit risk, especially for previously marginalised consumers. It is, therefore, helpful to examine the commercial considerations often discussed in isolation from potential normative risks.

In a recent paper "Algorithmic decision-making in financial services: Economic and normative outcomes in consumer credit" (AI and Ethics), I aim to challenge the persistent assumption that the use of algorithmic credit scoring and alternative data will only result in discriminatory outcomes or harm consumers.

We should not so readily dismiss the potential benefits of well-designed tools. Initially studied in isolation, ethical concerns will benefit from intersectional research alongside corporate perspectives.

Consider the notable example where the Apple Card (underwritten by Goldman Sachs Bank USA) was widely criticised for alleged discrimination against female credit card applicants, especially on social media. Some women were offered lower credit limits or denied a card, while their husbands did not face the same challenges. The claims sparked a vigorous public conversation about the effects of sex-based bias on lending, and the hazards of using algorithms and machine learning to set credit terms. The New York State Department of Finance investigated the algorithms involved and concluded there were valid reasons for these instances of disparity and could not find any discriminatory practices. The Department acknowledged that there are risks in algorithmic lending, including "inaccuracy in assessing creditworthiness, discriminatory outcomes, and limited transparency".

"Normative questions about the moral framework that guides AI cannot be divorced from questions about how we evaluate the moral framework that guides corporations"

First, I examine the economic implications of using machine learning to address traditional challenges in consumer credit contracts. These include information and power asymmetry between banks and consumers, as well as conflicting interests and incentives. Then, I consider the critical aspects of machine learning that dispel some misconceptions about algorithmic credit scoring. I explain how banks use machine learning to classify people and calculate credit scores and how they can use it to predict future consumer behaviour. Finally, the article evaluates risks that, if mitigated, could potentially improve economic and normative outcomes in the traditional consumer credit contract market.

These economic and normative issues include:

1. Whether ML increases the accuracy of the creditworthiness assessment of consumers;
2. The potential for ML to make more efficient pricing structures and provide a competitive advantage for banks with more accurate models;
3. If introducing algorithmic decision-making to the financial sector can further erode consumer trust and institutions' reputations;
4. The incongruity between improving accuracy and protecting consumers' privacy and autonomy;
5. The risk of ML replicating or compounding injustice and resulting in discriminatory algorithms.

There is considerable concern about the risk of algorithmic bias and discrimination in the context of credit institutions using ML. I highlight biases towards specific personal characteristics, such as race, gender, marital status or sexual orientation, that have historically impacted loan and credit decision-making processes. ML in credit scoring and access to financial services has amplified these concerns. Then, I consider the various technical fairness metrics proposed to overcome algorithmic bias and note that each metric requires different assumptions. This tension is exacerbated by the trade-off between fairness and accuracy when ML models are designed to prefer a certain level of fairness.

Such trade-offs are challenging for financial institutions, which like most companies, will continue to function with the prioritisation of profit. However, the future of corporations may shift with the knowledge, as described by Larry Fink, that "in fact, profits and purpose are inextricably linked". At the same time, as many consider the purpose and values of corporations, there is a similar impetus for the ethical design of AI.

Normative questions about the moral framework that guides AI cannot be divorced from questions about how we evaluate the moral framework that guides corporations. The reason is that, despite the misnomer, this view treats AI as ephemeral or autonomous, not as tangible decision rules and utility functions of the architect.

My article makes two essential contributions to the literature on the corporate use of algorithmic decision-making. First, examining the outcomes of using ML from a combined economic and normative approach is unique and allows for more rigorous consideration of the real-world costs and benefits. Second, despite the risk of harm that many experts in the field have identified, there is a clear opportunity to design ML. This will improve and optimise economic and normative outcomes. I propose a renewed enthusiasm for the potential positive outcomes.

I conclude that future work on regulatory issues should consider the underlying incentives and interests that shape behaviour in this area.

By Holli Sargeant, PhD Candidate in the Faculty of Law, University of Cambridge.

Voiceless at virtual shareholder meetings?

Miriam Schwartz-Ziv
Hebrew University of Jerusalem
and ECGI

Shareholder meetings are one of the only opportunities for most investors to meet and interact directly with management, and to raise concerns regarding the firm. While an extensive literature exists on shareholder votes, studies on the actual content of shareholder meetings are only starting to emerge. Since the onset of Covid-19 it shifted shareholder meetings from the in-person arena to the virtual-only arena: Clabaugh, Connors and Peters 2020 report that before Covid-19, only 12% of the S&P 500 companies held virtual-only meetings, but this figure increased to 77% after Covid-19. While, gradually, much of the world has gone back to operating in-person, as of now, virtual-only shareholder meetings are still happening essentially as often as at the heights of the pandemic (Broadridge, 2022).

Now that meetings are held virtually, there is good documentation of their content. This enables us to investigate whether at virtual-only meetings firms strategically employ certain methods that limit shareholders' voice, i.e., whether firms limit shareholder voice when shareholders are relatively critical of management. I address this question in my paper titled "Shareholders' Voice at Virtual-Only Shareholder Meetings". Virtual shareholder meetings may be of special concern because, in contrast to in-person meetings in which participants may occasionally talk or even shout when they do not receive permission to speak, and other participants in the meeting are aware of that, at virtual-only shareholder meetings, shareholders are unable to vocally oppose management in any way since their voice is literally muted throughout the meeting.

To understand whether shareholders' voice is strategically muted at virtual-only meetings, I focus on three methods firms may use to limit it. The first method is a firm's choice to ignore shareholders' questions at virtual-only meetings. At virtual-only meetings, questions in the Q&A session are submitted by shareholders in writing via a text box, frequently during the meeting; these are seen only by the firm's management, which can select which questions to reveal and address; questions that are not addressed are essentially never revealed. That is, of course, quite different from the way things happen in an in-person meeting, at which shareholders typically line up in front of the microphone and each is permitted to ask a question. The firm does not know in advance which question each shareholder will ask.

To capture the selection process of the questions ultimately addressed at shareholder meetings, I assembled a unique dataset that documents questions submitted by shareholders. I did this with the generous help of Mr. John Chevedden and Mr. James McRitchie (henceforth, "C&M")— two shareholders who, for many years, have been actively participating in shareholder meetings. I assembled a unique dataset that records all 767 questions C&M submitted between March 2020 and June 2021 (henceforth "Shareholder Questions Dataset"). Using the Shareholder Questions Dataset, I find that a question on a particular topic was significantly less likely to be addressed by a company when shareholders' voted against the directors proposed by management. For example, a one S.D. increase in the frequency of shareholder support for the directors proposed by management was followed by a 21.9% increase in the likelihood that a question would be answered by the firm (relative to the average frequency of the latter variable).

Put differently, precisely when shareholders' votes indicate that they are contentious with management, as indicated by shareholders' low support rates for the directors proposed by management, management is more likely to ignore the questions shareholders submit at virtual-only meetings, thereby limiting shareholders' voice.

The next two methods analyzed reflect the extent to which companies wish to encourage communication with shareholders at virtual-only shareholder meetings. To obtain this data, I hand-coded 1,904 transcripts of shareholder meetings held between January 2019 and June 2021 (inclusive). The first method analyzed is whether the firm explicitly limited shareholder questions at virtual-only meetings to topics related to the proposals submitted by shareholders. This policy severely limits the topics on which shareholders can ask questions, since their proposals pertain to a small range of topics. I find that when shareholders tend to vote against the directors proposed by management, firms are significantly more likely to limit questions to questions related to shareholders' proposals. Specifically, a decrease of one S.D. in shareholders' support of directors was followed by a 13.8% increase in the likelihood that the firm would limit questions to topics related to proposals (relative to the average frequency of the latter variable).

The last method analyzed, also based on data coded from the transcripts, is whether the firm reveals at the shareholder meetings the precise vote outcome for each vote, i.e., in percentage, or whether, alternatively, it merely reports whether each vote passed or failed.

When the firm does not reveal precise vote outcomes at the meeting, shareholders cannot "cite" a low-support vote outcome and ask why support rates are low, or how the firm intends to respond to the low-support vote outcome. Additionally, by delaying the revelation of the precise vote outcomes, the firm can stave off the media's and shareholders' legitimate criticism of proposals that passed with only low margins. The results indicate that especially firms that receive low support rates for the directors proposed by management are likely to disclose only pass/fail vote outcomes. Thus, the results show that firms that receive relatively low support rates from shareholders, are the firms that tend to use methods that make it more challenging for shareholders to make their voice be heard.

To conclude, while the technology allowing companies to hold virtual shareholder meetings has the potential of increasing shareholder democracy since attending such meetings is substantially less costly than attending in-person meetings, currently, virtual-shareholder meetings are not maximizing shareholder democracy to their full potential. Reaching the latter goal could be enhanced by allowing shareholders to present their questions "live and unfiltered," to require firms to disclose all questions submitted by shareholders, forbidding firms to restrict the questions to topic pertaining to proposals, requiring firms to disclose precise vote outcomes at the shareholder meeting, and making the recordings of shareholder meeting publc and easily accessible.

By Miriam Schwartz-Ziv, Assistant Professor of Finance, Hebrew University of Jerusalem and ECGI Research Member.

Governance and the crypto winter

Daniel Ferreira
London School of Economics and Political Science (LSE) and ECGI

The Crypto Winter is upon us. The collapse of Sam Bankman-Fried's empire has sent shock waves all over the crypto world. Crypto prices have taken a hit. The real world has taken notice, too. Mr Bankman-Fried (also known as SBF), the billionaire who founded FTX, a major crypto exchange, was widely seen as the acceptable face of crypto. His fall from grace was cheerily celebrated by crypto sceptics, who seized the opportunity to pat themselves on the back and send I-told-you-so notes all around. On the other side, crypto enthusiasts were quick to distance themselves from the whole affair. The official line is that SBF's empire was antithetical to crypto's philosophy. If only we could get more of that utopian "decentralisation" we were once promised, none of this would have happened.

Both sides might be right. Crypto has a governance problem. This problem is in crypto's DNA and poses an existential threat to the whole project. Unfortunately, blockchain developers and other stakeholders have paid little attention to this issue. Discussions of governance in this space are a mix of hubris and bad economics. According to crypto evangelists, clever computer scientists versed in cryptography can now solve governance problems that have afflicted societies for millennia. It just sounds all a bit too incredible to those versed in corporate and public governance.

The downfall of FTX and its affiliated companies is a classic example of a failure of corporate governance. Unbeknownst to most, FTX shipped their customers' assets to Alameda Research, a trading firm controlled by Mr Bankman-Fried. The corporate governance literature has a name for that trick: tunnelling. How could FTX's controlling parties tunnel investors' funds to other companies? The answer is simple: there were no checks and balances, no transparency and no regulation. FTX customers put their trust in Mr Bankman-Fried, and now most of their money is gone. This episode is also an example of a failure of blockchain governance. Mr Bankman-Fried's empire comprised a vast number of entities operating in the crypto ecosystem. As a result, many blockchain-based projects were exposed to Mr Bankman-Fried's idiosyncratic decision-making. A case in point is the Solana blockchain, a smart-contract platform once hailed as the "Visa of Crypto." Solana had a torrid 2022, with multiple outages and hacks. At one time, the whole blockchain was switched off and on to fix a malfunctioning node. How can an allegedly decentralised network decide when to reboot?

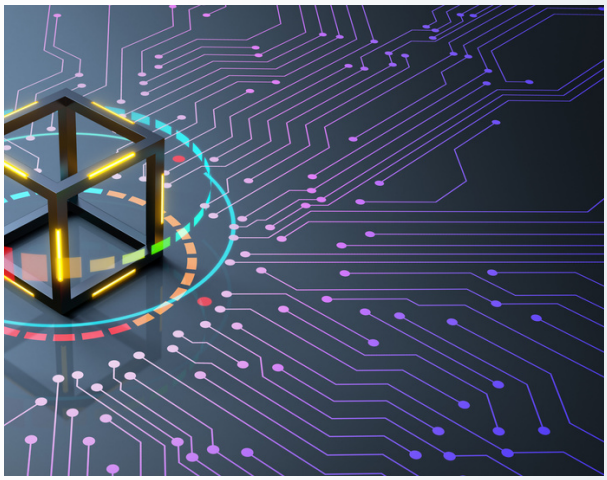

The answer is simple: Solana, like all blockchains, is not really decentralised. As in all blockchain networks, the real power resides in the hands of a few key players: founder-controlled foundations, core developers, large validators and other major companies in the blockchain ecosystem. Mr Bankman-Fried was one of such players; Solana's native token (SOL) was dubbed a "Sam coin." Not only was Mr Bankman-Fried a vocal backer of Solana, Alameda Research allegedly held about 13% of all SOL by the end of November. FTX was also responsible for the critical infrastructure for Solana DeFi's ("decentralised finance") operations. Moreover, we now know that the Solana Foundation was an investor in FTX. SOL's already depressed price fell by more than 60% in November after the FTX collapse.

Although notionally decentralised, Solana's governance was effectively captured by a few key players, including Mr Bankman-Fried. It is not surprising to find a conglomerate (the FTX-Alameda-SBF complex) at the centre of this crisis. In a forthcoming paper, Jin Li from the University of Hong Kong, Radoslawa Nikolowa from Queen Mary University of London, and myself argued that blockchain conglomerates arise naturally due to network externalities, economies of scope, and first-mover advantages. As a by-product of conglomeration, a few large firms end up capturing the governance of blockchains. Such firms care about the private value they can extract from a blockchain project, which might differ from the project's social value.

We illustrate our arguments with the case of Bitcoin, which is arguably the least centralised of all blockchain projects. Bitmain Technologies, a private Chinese company that designs chips for mining bitcoin, has approximately 75% of the global market share. Bitmain is also a prominent player in other segments of the Bitcoin ecosystem, such as mining pools. Mining pools owned by or affiliated with Bitmain have consistently dominated the market, with market shares always above 30% -- and often above 50% -- from October 2016 to early 2021.

"Discussions of governance in this space are a mix of hubris and bad economics"

Bitmain has been an influential player in the governance of Bitcoin. On August 1, 2017, a few key players, including Bitmain, sponsored the creation of the new chain, Bitcoin Cash, through what is known as a "hard fork." Bitcoin Cash shares the same history as Bitcoin but has a larger block size. On November 15, 2018, Bitcoin Cash split into two competing blockchains. Bitmain rallied behind Bitcoin Cash ABC against Bitcoin Cash SV in what became known as the "hash wars." The prices of both coins fell steeply after the split, as did the prices of Bitcoin and other cryptocurrencies.

While Bitmain's impact on Bitcoin has been much less damaging than FTX's impact on Solana, both cases share the same root problem: The emergence of a conglomerate that captures the governance of the blockchain. In such cases, blockchain stakeholders have to trust one company to look after their interests. So how exactly does a "decentralised" blockchain differ from a traditional financial intermediary as a provider of trust? This is the question that crypto enthusiasts still need to answer.

By Daniel Ferreira, Head of Department and Professor of Finance at the London School of Economics and ECGI Fellow.

The Editorial Board

Vicente Cuñat
Associate Professor of Finance
London School of Economics (LSE)

Enrichetta Ravina
Senior Economist
Federal Reserve Bank of Chicago

Mariana Pargendler
Professor of Law
Fundação Getulio Vargas (FGV) Law School, São Paulo

Georg Ringe
Professor for Corporate Law and Financial Markets
University of Hamburg

Susan Watson
Professor of Law
Univerity of Auckland

Daniel Ferreira
Professor of Finance
London School of Economics (LSE)

Nadya Malenko
Associate Professor of Finance
University of Michigan

Umakanth Varottil
Associate Professor
National University of Singapore

Marco Ventoruzzo
Professor of Business Law
Bocconi University

Fei Xie
Professor of Finance
University of Delaware

The Advisory Board

Lucian Bebchuk
James Barr Ames Professor of Law, Economics, and Finance
Harvard Law School

Marco Becht
Professor of Finance
Université libre de Bruxelles

Herman Daems
Chair, ECGI
Chair, BNP Paribas Fortis

Paul Davies
Senior Research Fellow
University of Oxford

Guido Ferrarini
Professor of Business Law and Capital Markets Law
University of Genoa

Julian Franks
Professor of Finance
London Business School

Hideki Kanda
Emeritus Professor
University of Tokyo and Gakushuin University Law School

Kon Sik Kim
Emeritus Professor of Law
Seoul National University

Rui Pereira Dias
Professor of Law
IPCG

Nicoletta Pollio
Legal Counsel
Enel

Isabella Porchia
Legal Associate
Latham & Watkins

Paolo Rainelli
Associate Professor of Business Law
Politecnic University of Turin and Cleary Gottlieb

Ronald Gilson
Professor of Law and Business
Stanford Law School, and Columbia Law

Peter Hope
Partner
Oxera Consulting LLP

Michael Hilb
Chair
International Board Foundation

Jennifer Hill
Bob Baxt AO Chair in Corporate and Commercial Law
Monash University

Klaus Hopt
Emeritus Professor
Max Planck Institute for Comparative and International Private Law

Mark Roe
David Berg Professor of Law
Harvard Law School

René Stulz
Everett D. Reese Chair of Banking and Monetary Economics
The Ohio State University

Cristiana Tudor
Professor
CECCAR

Alexis Wegerich
Economist
Norges Bank Investment Management

CONTACT

Carla Speight
Senior Communications Manager
European Corporate Governance Institute (ECGI)
carla.speight@ecgi.org
www.ecgi.global

Elaine McPartlan
General Manager
European Corporate Governance Institute (ECGI)
elaine.mcpartlan@ecgi.org
www.ecgi.global

European Corporate Governance Institute (ECGI)
c/o Royal Academies of Belgium
Palace of the Academies
Rue Ducale 1 Hertogsstraat
1000 Brussels
Belgium
admin@ecgi.org

www.ecgi.global/blog