

Feature

KEY POINTS

- Compliance software and artificial intelligence is active within the government sphere.
- Technology has been used to prosecute insider dealing, monitor orders and executions in the securities markets and protect governments against cyberattacks.
- Despite this trend courts have stepped in to protect the privacy rights of individuals.

Authors Christopher Murrer and Caleb Sainsbury

Governmental use of technology for compliance monitoring and enforcement

In this article, the authors analyse governments' increasing use of AI and other technologies to enhance their compliance monitoring capabilities and to conduct investigations.

INTRODUCTION

Governments around the world grapple with the challenge of keeping up with technology. Most of the time, that struggle centres around developing a regulatory framework for that technology to protect consumers or the government's own financial interest. One sees this, for example, in the way the US' IRS has developed regulations and rulings for the appropriate treatment from a tax perspective of cryptocurrency or in the way certain countries have approached regulations for self-driving vehicles.

What is less well known, however, is the way in which governments harness new technologies for their own means. Or, in other words, the way in which governments use technology to enhance tax and compliance monitoring capabilities, conduct investigations, and capture criminals. This article, by no means comprehensive, addresses a few important areas in which governments have started to harness investigations and monitoring software and artificial intelligence technology. In particular, this trend is taking off in sectors that generate large amounts of data or particularly complex types of data. It may be premature to predict which regulatory agencies will adopt artificial intelligence as part of its day-to-day operations, however, it is clear that compliance software and artificial intelligence is active within the government sphere and will continue to grow.

THE SEC'S TRADING INVESTIGATION SYSTEM, ARTEMIS

Stating the obvious, the volume of data generated in today's world is significantly

larger than it was twenty years ago. This presents both a challenge and opportunity for governments. The opportunity means that there are many more footprints in the data available for governments to track catch criminals and punish behaviour that were not available before. However, the amount of data is both a blessing and a curse. The curse is similar to trying to find a needle in a haystack, but the haystack grows every day. So, governments need to either adapt or write this additional data off as a lost cause.

One agency that has adapted to the new influx of data is the United States Securities and Exchange Commission (SEC). The SEC's mission is to "protect investors, maintain fair, orderly, and efficient markets, and facilitate capital formation".¹ One of the more well-known aspects of enforcing this mission is the SEC's work around prosecuting insider trading. In years past, the most reliable way for the SEC to catch wrongdoers in this area was to investigate tips or information provided by known informants. However, in recent years, the SEC has created a program that is designed to look for irregularities in trading as part of its investigatory toolkit. The program is called the "Advanced Relational Trading Enforcement Metric Investigation System" or "ARTEMIS". Former SEC chairman Mary Jo White stated about the program:

"In insider trading investigations, for example, the Division uses sophisticated software to identify and assess suspicious trading. One very successful program, called ... 'ARTEMIS', analyzes patterns and relationships among multiple traders

using the Division's electronic database of over six billion electronic equities and options trading records."²

So, tapping into its massive database, the SEC uses ARTEMIS to analyse the data to look for irregular trading patterns. For example, say a certain company has recently announced a round of layoffs indicating and posting lower quarterly earnings. The program would look at trades surrounding that event and identify if an individual's trading activity of that stock deviated from their normal trading pattern. If those trades do not match the history, then the SEC's enforcement officers would then have evidence of irregularities and cause to look further into that particular individual.

The SEC openly announces its use of data analytics in enforcement actions. For example, in 2018 the SEC brought an enforcement action against a defendant for a "cherry-picking" scheme.³ In this particular scheme, the defendant would purchase stocks with an omnibus account which included client funds, wait to see if the value went up or down, and routinely assign picks that increased in value to his account or accounts of family members and assign losing picks to client accounts. In its press release on this case, the SEC stated it had "uncovered the alleged fraud with data analysis used to detect suspicious trading patterns".⁴ So, the use of data analytics to mine vast quantities of data is now a reality for those regulated by the SEC.

CONSOLIDATED AUDIT TRAIL IMPLEMENTATION BY SECURITIES AND EXCHANGE COMMISSION

The SEC also established a plan in 2012 to create a consolidated audit trail system (CAT). The CAT is intended to improve regulatory oversight of the securities market. The impetus behind the CAT was

the 2010 “flash crash”. On 6 May 2010, a US\$1trn stock market crash occurred and lasted approximately 36 minutes. In a 2010 joint report between the SEC and the Commodity Futures Trading Commission (CFTC) an investigation into the crash “portrayed a market so fragmented and fragile that a single large trade could send stocks into a sudden spiral”. The report espoused a theory of the reason for the crash, which received significant criticism and there remains disagreement over the true cause. Accordingly, the CAT was envisioned to facilitate cross-market oversight and analysis with the aim of enhancing investor protection and market integrity.

The infrastructure buttressing the CAT will require self-regulatory organisations (SROs) to create a comprehensive surveillance database that will collect and store all securities activity from every broker and investor. The existing system allows SROs to use their own, separate audit trail systems to track information and the requirements of each system vary among markets – leaving regulators to obtain, merge and reconcile data in disparate forms to assess issues in the market. Thus, the SEC posited that there was no single database of comprehensive and readily accessible data regarding orders and executions, and that the CAT was necessary to rectify this deficiency.

SEC Commissioner Hester Peirce, however, published a dissent over the SEC’s decision to pursue CAT. She describes the to-be-created CAT as, a “gigantic database, housing all this information in a single place, ... accessible to thousands of people at the Commission and the SROs, who will be able to watch investors’ every move in real time”. She asserted that the CAT was unlikely to fulfil the intended purpose and would have unintended consequences. Commissioner Peirce explained that the CAT system effectively creates a surveillance system in that trading activity is not simply a financial mechanism, it can provide insight into the moral convictions, values and beliefs of companies, markets, products, and the economy at large. Furthermore, trading

activity can reveal business strategies. Likewise, it presents an opportunity to question any particular activity in a vacuum or against broader trends in trading activity. She also raised the concern of security breaches of this information either by hackers who gain unauthorised access or by the numerous SEC employees and contractors who will have authorised access.

Notwithstanding Commissioner Peirce’s dissent, the plan to establish the CAT continues.

PROTECTING GOVERNMENTS AGAINST ATTACKS WITH DARKTRACE

Perhaps one of the more interesting aspects of AI and machine learning is the intersection between private companies and government agencies. An example of this is a company called Darktrace which has business relationships with many governments around the world. Darktrace is different from ARTEMIS in that it is not designed to prosecute individuals *per se*, but rather to protect governments against attacks. It has developed a software called “Enterprise Immune System” designed specifically for this purpose.⁵ Per the company’s description:

“The Enterprise Immune System uses unsupervised machine learning and AI to understand all about your organization. Observing your users and devices, cloud containers and workflows, it learns ‘on the job’ what is normal for your organization.”

There are many examples of governments using Darktrace’s products in interesting ways. According to a press release from Darktrace a city government in the US employed Darktrace’s software and identified a vulnerability in its system that exposed the personal data of city residents.⁶ In another instance, a provider of educational services to 70 school districts and 60 charter schools in Texas used the software to protect student data.⁷ The use of the software is not limited to lower level school districts and cities. For example, the

UK government has hired Darktrace to protect “crucial public services and citizens’ data”.⁸

Though this is an example which differs from an enforcement type program, the fact that governments now actively use AI in their day-to-day operations suggests this direction will be the new reality. There are some tricky legal and ethical issues with outsourcing access to sensitive data to private parties; however, as of now having a solid cybersecurity defence carries more weight.

USE OF ARTIFICIAL INTELLIGENCE AND AUTOMATED DECISION SYSTEMS IN GOVERNMENT ADMINISTRATION

The confluence of Big Data, advances in computing speed, and more sophisticated algorithms has paved the way for expanded and new applications of AI. Indeed, a report by IBM indicated that of insurance companies that outperform their competition, the vast majority used technology providing advanced and predictive analytics. To maintain a competitive advantage or to simply keep up, AI and advanced and predictive analytics will become increasingly mainstream. Governments have also utilised AI in the insurance context, with mixed results.

The Michigan Unemployment Insurance Agency provides temporary income replacement to workers who become unemployed through no fault of their own. Workers who quit their jobs or who were fired for misconduct are not eligible for the benefits. Additionally, continued benefits are subject to the individual making every effort to find full-time, suitable work. Historically, the Agency audits and examines data to ensure that the benefits an individual receives matches their prior employment wages and there is otherwise no evidence of fraud.

In 2013, the Agency implemented the Michigan Integrated Data Automated System (MiDAS). MiDAS leveraged AI to make determinations about whether individuals receiving unemployment benefits from the Agency had engaged in

Feature

fraud. The goals of MiDAS were improved customer service, increased data accuracy, improved data security and privacy, reduced operating costs, increased automation, and improved integration of the Agency's functions. If MiDAS determined that fraud existed, it triggered an automatic termination of the unemployment benefits and demand for repayment of previously paid benefits. In practice MiDAS determined that at least 40,000 individuals had fraudulently claimed unemployment benefits between 2013 and 2015. In 2014, MiDAS opened nearly 27,000 cases, which is more than five times the number typically opened during any given year. Several Michigan residents brought a class action lawsuit against the state of Michigan alleging that MiDAS violated their constitutional rights by granting decision-making abilities to an automated fraud-detection system. In some cases individuals were subjected to 400% fines and faced wage garnishment. The plaintiffs alleged further that the system had the power to terminate their benefits without providing them with notice or an opportunity to present evidence countering the conclusion. In 2019, the Michigan supreme court and a court of appeals rendered decisions that permitted the lawsuit to continue. Progress on the substantive issues languish as courts continue to make decisions on procedural issues.

In the meantime, a technology that assisted the development of MiDAS has been used on other projects in South Carolina, New Mexico, Illinois and Tennessee. Automated decision-making systems have been deployed across a range of topics, including, for example, housing matching programs for people experiencing homelessness, monitoring child welfare, and evaluating teachers for termination or bonuses. Complaints with some of these systems have mirrored those in Michigan. For example, in 2014 Rhode Island deployed an automated system intended to streamline the federal and state benefits programs. Thereafter, residents dependent on state aid reported that they lost coverage without adequate notice. This issue led to

a backlog of over 15,000 applicants and two federal class action lawsuits. Idaho implemented an automated decision system to determine Medicaid care budgets for developmentally disabled individuals. The system determined that benefits would cease for thousands of existing recipients. A lawsuit ensued, which revealed methodological issues with the data on which the system relied and its conclusions relative to conclusions given by human reviewers in similar cases.

It is also important to recognise that MiDAS and the other systems were implemented in light of AI's potential to reduce bureaucracy and offer efficient, accurate decisions. Several examples confirm this potential. Of course, other examples indicate that there are still significant lessons to be learned.

LIMITS ON AI USE

Although the trend is towards AI use, there have been examples where the courts have stepped in to prohibit its implementation. One recent interesting case comes from the Netherlands. The Netherlands had developed an algorithm designed to prevent fraud "in the fields of social security and income-related schemes, tax and social insurance contributions, and labor laws".⁹ Essentially, the algorithm linked data from various Dutch agencies, such as tax authorities, municipalities and immigration services, and generated a risk report if a person was suspected of fraud.¹⁰ The use of this algorithm was challenged by several civil rights organisations in the Netherlands.

In reviewing the case, the court looked at the EU General Data Protection Regulation as well as the EU Charter of Fundamental Rights and ultimately found that the algorithm's use in practice did not strike an adequate balance between the Dutch government's stated purpose of preventing fraud and the privacy rights of individuals. There was also concern that the program would discriminate against minority and immigrant groups.¹¹ So, there appears to be limits to the use of AI software. The coming years will likely

bring additional challenges and therefore boundary setting in this area.

GOVERNMENTS LEVERAGING CRYPTOCURRENCY INVESTIGATION SOFTWARE TO FIGHT CRIME

Bitcoin became operational in 2009, which spurred what many refer to as a new asset class of cryptocurrencies. Cryptocurrencies are digital native units that are treated as a medium of exchange and, for the most part, are not government created, issued, or controlled. Transactions involving many of the popular cryptocurrencies are recorded on a blockchain, which is a continuously growing list of records (blocks) linked together chronologically and secured using cryptography. The blockchains on which mainstream cryptocurrencies are based are most often considered public and permissionless, meaning that any individual may set up a computer (referred to as a node) to view that cryptocurrency's blockchain and, consequently, the transactions involving that cryptocurrency.

Numerous cryptocurrency-based services have arisen to serve individuals' access to, exchange and transfer of cryptocurrencies. Cryptocurrency exchange companies, for example, have been created to facilitate the exchange of fiat currency (that is, traditional government issued currency) to cryptocurrency, the transfer of cryptocurrency among various other individuals, and the exchange of one cryptocurrency for another type. These activities serviced by cryptocurrency exchange companies are not usually recorded on the blockchain. Rather, these transactions are recorded on a separate ledger of that particular exchange (referred to as off-chain activity). When an individual directs the exchange to transfer the cryptocurrency to an individual who uses services of a different exchange, that transaction is recorded on the blockchain. Consequently, the blockchain records that the two exchanges were involved in a transaction – although this record may not reveal all activities that happened when a particular cryptocurrency exchange had control over that particular cryptocurrency.

Biog box

Christopher Murrer is an associate in the Fintech, International Tax and Wealth Management practice at Baker McKenzie, based in Zurich.

Email: christopher.murrer@bakermckenzie.com

Caleb Sainsbury is an associate in the Global Wealth Management and Compliance and Investigations practice group at Baker McKenzie, Zurich.

Email: caleb.sainsbury@bakermckenzie.com

Further, many governments require that cryptocurrency exchanges that operate within their jurisdictions must fulfil customer due diligence and anti-money laundering requirements when onboarding customers and monitoring activity on the exchange – in similar ways that banks and other financial institutions are required to do. In other words, although the transactions involving the exchange are not recorded on the public, permissionless blockchain, the exchange itself should have information about the individuals involved.

Chainalysis is a company that offers cryptocurrency transaction monitoring services. Governments engage Chainalysis to leverage the company's investigation and monitoring capabilities. A government coalition involving the US, UK, South Korea, Germany and Saudi Arabia, The United Arab Emirates, the Czech Republic, Canada, Ireland, Spain, Brazil and Australia, pursued a network of contributors and users of child abuse material websites by tracking cryptocurrency payments involving these individuals. The government coalition was able to trace the transfer of cryptocurrencies using Chainalysis's software both on the blockchain's public ledger and transfers among various cryptocurrency exchanges. When the services of a particular cryptocurrency exchange were implicated, a government with jurisdiction could then request information from the cryptocurrency exchange about the customer which presumably was obtained through the customer due diligence process. In October 2019 the US Department Justice announced that it shut down the largest ever child abuse material website, arrested its owner and operator and arrested more than 227 site users across 38 countries. More importantly, 23 minors were identified and rescued from abuse as a result of the investigation.

The SEC also published a document indicating plans to run through contractors of Bitcoin and an Ethereum full node and nodes on as many as possible of the following blockchains: Bitcoin Cash, Stellar, Zcash, EOS, NEO and XRP Ledger. The SEC is seeking the full ledgers

since inception (ie the genesis block) and all derivative currencies (tokens) for all of those blockchains. The SEC stated that the purpose would be “to support its efforts to monitor risk, improve compliance, and inform Commission policy with respect to digital assets”.

Thus, the commonly stated adage that cryptocurrency transactions are anonymous is a misnomer. Rather, it is a matter of governments acquiring the knowledge and capabilities to monitor and trace cryptocurrency transactions. They have already made great progress on this front.

CONCLUSION

These examples make clear that as governments pursue compliance and investigation software and leverage AI, the governments will become larger and larger repositories of data and more and more capable of monitoring behaviour in real-time. These results may produce the effect intended by governments: for individuals and companies to fulfil their compliance obligations and limit the incidences of non-compliance. These results will also lead to new issues and responsibilities for governments. The large amounts of data will be a target of cybercriminals who look to steal the data, and governments will have to institute protections commensurate with that risk. Governments will also have to consider if and how they limit the scope of their monitoring software so the software's capabilities do not creep into aspects of individual behaviour unrelated to compliance issues and, consequently, infringe personal freedoms. As usual with technological developments, great potential also brings great responsibility and complications. ■

- 1 <https://www.sec.gov/Article/whatwedo.html#:~:text=The%20mission%20of%20the%20U.S.,markets%2C%20and%20facilitate%20capital%20formation.>
- 2 <https://www.sec.gov/news/statement/statement-mjw-040816.html>
- 3 <https://www.sec.gov/litigation/complaints/2018/comp-pr2018-189.pdf>
- 4 <https://www.sec.gov/news/press->

[release/2018-189](#)

- 5 <https://www.darktrace.com/en/products/enterprise/>
- 6 <https://www.darktrace.com/en/blog/trusting-the-cloud-unencrypted-data-upload-by-government-body/>
- 7 <https://www.darktrace.com/en/press/2018/217/>
- 8 <https://www.darktrace.com/en/press/2018/246/>
- 9 <https://www.loc.gov/law/foreign-news/article/netherlands-court-prohibits-governments-use-of-ai-software-to-detect-welfare-fraud/>
- 10 <https://www.loc.gov/law/foreign-news/article/netherlands-court-prohibits-governments-use-of-ai-software-to-detect-welfare-fraud/>
- 11 <https://www.loc.gov/law/foreign-news/article/netherlands-court-prohibits-governments-use-of-ai-software-to-detect-welfare-fraud/>

Further Reading:

- The future of financial transactions: automation and the use of structured data in legal practice (2020) 6 JIBFL 359.
- Automation and blockchain in securities issuances (2018) 3 JIBFL 144.
- LexisPSL: Practice note: The automation of contracting and contract lifecycle management.