Department for
Business, Energy
& Industrial Strategy

THE USE OF EMERGING TECHNOLOGIES FOR REGULATION

Main findings and roadmap

BEIS Research Paper Number 2020/041

October 2020

Authors

Kate McEntaggart, Julien Etienne, Simona de Paolis

Acknowledgements

Janne Fillet, Beatrice Marinoiu and Helen Quaresma contributed to the data collection for this report. The participation of interviewees to inform the case studies is gratefully acknowledged. The study team is also grateful to previous work carried out on data-driven and algorithmic regulation for BEIS by Dr Panos Panagiotopoulos and the Better Regulation Executive.

Any enquiries regarding this publication should be sent to us at:
enquiries@beis.gov.uk

Contents

Executive summary

Objectives

Regulators are increasingly making use of new and emerging technologies, for example using automation or data science techniques to improve the efficiency and effectiveness of regulation. However, the extent to which different regulators are currently adopting and developing such technologies is unclear. Regulators themselves could benefit from learning more from the experience of their peers, but to date there has been very little detailed information available on the work regulators are doing in this area.

In this context, the aim of this research was to help address this gap by gathering evidence on regulators' use of emerging technology to improve the efficiency or effectiveness of regulation in a transformational way. Using this evidence, it has looked to identify best practices and practical considerations for regulatory applications of emerging technologies.

This report builds on interviews conducted with around 20 regulators in 2019 by the Better Regulation Executive (BRE) and Dr Panos Panagiotopoulos to gather evidence on technology adoption and application in regulation, and the issues affecting it.

Method

The first phase of the project involved a desk-based review of existing documentation regarding the use of emerging technologies by regulators. This scoping phase identified a long list of 45 examples of emerging technologies being used for regulation.

The second phase of the study involved conducting 12 case studies selected from the initial long list. The final shortlist was chosen to cover a representative range of technologies and regulatory functions as well as different levels of maturity, complexity and working styles.

Case studies were primarily based on interviews with stakeholders. Each case study involved 1 to 3 interviews with members of the project teams who developed the solution. In some cases, a group interview was held. Interviews took place over the phone or through video conferencing and lasted between 45 and 60 minutes.

Case study evidence was then coded using qualitative data analysis software and findings were synthesised by theme. Cases were compared against one another to identify the extent to which interviewees faced similar challenges and to identify where approaches to address those solutions had been identified. This has formed the basis for a series of good practice recommendations, presented in this report.

Findings

The 12 case studies discussed within this report cover a range of regulatory functions, sectors and technologies. These are listed in Table 1 below.

Table 1: Case studies

Example	Technology function
The Intellectual Property Office (IPO) and Machine Learning (ML)	The IPO explored the use of machine learning (ML) tools to improve processes in the registration of patents and trade marks. ML has been used to help examiners identify similar trade marks and, for patents, prior art, more efficiently than this was usually done.
The Rural Payments Agency (RPA) and CROME	RPA used a type of supervised machine learning to develop a crop map of England (CROME) based on satellite images and image classification methods. This can then be used to better assess/allocate rural payments and avoid fraud.
A consortium of private firms and public sector bodies' use of sensor technologies for mosquito classification (the Vectrack Project, EU)	The Vectrack system uses sensors and machine learning to collect and classify mosquitos according to several criteria such as gender, provenance and risk of disease in real time.
Ofcom and the use of blockchain for a common numbering database	Ofcom has piloted the use of blockchain to support the development of a common database of landline numbers in the UK.
The Finnish Government's use of a Virtual Agent Network	The Starting Up Smoothly project explores the use of chatbots to provide an integrated approach to customer service between different organisations. It focused on the test case of linked chatbots to address the needs of foreign entrepreneurs looking to start businesses in Finland.
The Ministry of Justice (MoJ)'s development of an intelligent search tool	The MoJ developed an intelligent search tool for prison reports using natural language processing (NLP) techniques. The tool brings together reports that were previously dispersed in different locations and allows users to conduct searches that return contextually similar results in addition to exact matches.
North Tyneside Council's use of RPA	North Tyneside launched a project to apply robotic process automation (RPA) to repetitive tasks within the Council's services.
The Financial Conduct Authority (FCA)'s Data Strategy	This FCA adopted a new Data Strategy, rather than a specific tool, involving several organisational change projects: bringing in infrastructure, hiring internal data science teams, improving data management practices and promoting an innovation-friendly culture.
Natural England's use of eDNA	Natural England employed the use of environmental DNA (eDNA) for sampling great crested newts, a protected species under UK and European Law.
The Environment Agency's use of webcams	The Environmental Agency employed the use of webcams and image differentiation technologies to monitor culverts and respond to flood hazards.
The Driver and Vehicle Standards Agency (DVSA)'s Intelligent Risk Rating Assessment tool	The DVSA has developed a tool that uses ML algorithms for pattern identification on MOT Testers. This provides risk ratings facilitating targeted MOT inspections based on past behaviour, allowing resources to be focused on garages with the highest risk scores.
The City of Chicago and WindyGrid	WindyGrid is a platform that aggregates City data from across departments. The platform has been used as a basis to develop features for predictive analytics, which leverage this data to improve city services.

Several distinct themes emerged from the case studies, around which the study findings were organised.

Thinking about problems and solutions

The definition of problems or needs and the identification of a potential solution constitute the first steps in the development of innovative uses of emerging technologies for regulation. Through the case studies, the following good practices have been identified:

1. Take a problem-focused approach rather than letting preference for a given solution shape the process. This means identifying problems and considering solutions suitable for those problems, rather than identifying a technological solution and looking for ways to apply it.
2. Consider a range of options and acknowledge that emerging technologies may not offer the best solution in every situation.
3. Regularly monitor for the emergence of new technologies and market trends as off-the-shelf solutions may become available.
4. Learn from the experiences of others, including other regulators.
5. Consider the implications of scaling up prototypes, developing an operational tool, and maintaining that tool over time.
6. Consider starting small or choosing relatively “easy” first examples of implementation to build skills, confidence, and ensure buy-in from internal and external stakeholders.
7. Build on existing projects: adding new features or developing new uses for existing tools can offer a good approach to developing more sophisticated solutions.

Infrastructure and tools

For many of the technologies explored in this study, innovating to apply them to regulatory work has required having at hand or building up / acquiring technical infrastructure or tools. The case studies show that these were often pre-conditions to the successful development and roll-out of innovations using emerging technologies. The case studies identified the following elements of infrastructure or tools which served as pre-conditions to innovation projects:

1. Have cloud-based systems in place, particularly to facilitate data-heavy or machine learning projects.
2. Develop centralised data management and storage systems, for projects that seek to leverage existing data.
3. Develop a bespoke in-house analytical platform for data science teams.
4. Consider bringing in off-the-shelf tools that support data analytics.

In-house teams and external contractors

The case studies demonstrated different approaches to building project teams: some relied largely on in-house resources, others worked alongside external contractors. Many took mixed approaches that reflected the project needs and the skills available. Through the case studies, the following good practices have been identified:

1. Consider the project's needs carefully: what technical and non-technical skill sets are needed, and the extent to which existing teams can provide these.
2. If there is capacity to conduct work in-house, consider what challenges this might bring and the extent to which the use of an expert external contractor may add value to the project.
3. If there is a need to hire new staff, address this as early as possible and plan for challenges.
4. If there is a need to engage external contractors, consider and mitigate against the risks of becoming dependent on those contractors, or of ending up with costly solutions.

Organisational buy-in and staff engagement

Another set of challenges commonly cited across cases was engaging and educating their own staff about the proposed technological solution and achieving buy-in from the wider organisation. Through the case studies, the following good practices have been identified:

1. Collaborate with a wide range of staff early-on in the development process.
2. Have external contractors host informational events for staff.
3. Illustrate the functionality of the tool at an early stage in the process (e.g. showing rather than telling).
4. Rely on operational staff (first users) to explain the tool and its benefits to other staff.
5. Pilot tools internally with staff (including where staff are not the intended end users).
6. Upskill staff and, where possible, involve staff in the development process.

User testing

Tools intended for widespread staff or consumer use require some degree of user testing. In a few cases — notably those conceived as prototypes — interviewees felt that they had not done as much user testing as they should have due to a lack of time, resource or access. Testing is an important part of engaging staff, but it also ensures that the design is relevant and functional. Through the case studies, the following good practices have been identified:

1. Include user testing and feedback early-on in the design process.

2. Plan for the extra time user testing is likely to take.
3. Identify and engage likely end-users as soon as possible.

Data-driven technologies

Many of the case studies involved data analytics or machine learning applications, which required the collection, acquisition and/or processing of large amounts of data. Based on the case studies, some good practices for data use include:

1. Set strict data validation requirements within tools.
2. For external data, consider potential intellectual property issues early in the process.
3. Consider the implications of using personal or sensitive data.
4. Where collecting new data, expect additional challenges and costs.

Measuring impacts

None of the case study examples included in this report have been the subject of a formal impact evaluation, although some either will undergo or have undergone another type of assessment or evaluation.

For both quantitative and qualitative measures of impact, **organisations might benefit from taking a more robust and formal approach** to these. By setting out likely benefits early on and devising methods for evidencing or measuring these, the full value of such projects can be better understood. In turn, this can help innovation teams improve their processes and improve their business case for future projects.

Conclusions

Regardless of regulatory function, size, scope and budget associated with these cases, many of the organisations involved have faced similar issues. For example: choosing the right solution, building appropriate teams, engaging staff and/or users in the process and considering impact are important elements of applying emerging technologies to regulation. Figure 1 below sets out a potential roadmap for applying the best practices identified in this report.

Figure 1: Roadmap of best practices for the use of emerging technologies for regulation

Introduction

BEIS's white paper on Regulation for the Fourth Industrial Revolution[1] notes that changes in technology are occurring at a "scale, speed and complexity that is unprecedented". Emerging technologies such as artificial intelligence (AI), machine learning (ML), blockchain and robotic process automation (RPA) mean likely disruption across economic sectors. From the use of chatbots and other automated systems changing customer service practices, to the development of potential future technologies such as connected and autonomous vehicles, such technologies contribute to the development of new products, services and ways of working.

Those changes will and have already begun to pose challenges for traditional approaches to regulation. However, emerging technologies have been also increasingly considered and used by regulators; harnessing these emerging technologies could help improve regulatory processes, for the benefit of regulators themselves but also of those (individuals, businesses and other organisations) who interact with them.

The use of such technologies can help improve regulatory processes in several ways: allowing for greater process efficiency and cost-cutting (e.g. through the use of RPA or the application of ML to automate operations), improving how the enforcement and monitoring of regulation is targeted (e.g. through the use of data analytics, AI/ML to develop risk models) and through the development of new capabilities (e.g. by using technologies such as sensors to collect data at a scale previously unfeasible). For many public sector organisations, however, making use of emerging technologies comes with challenges, such as finding sufficient resources or the internal technical capacity.

This report was commissioned to support the take-up of emerging technologies within UK regulatory agencies, government departments and local authorities. It presents the findings from a suite of case studies of agencies, ministries, and local authorities in the UK and abroad, which have successfully developed and, for most of them, implemented new solutions that incorporate emerging technologies. Through a review of documentation and interviews, this study has aimed to extract key learning points for the benefit of others within government who may also consider experimenting with AI, ML, RPA or other emerging technologies.

The body of the report presents the background to the study as well as the methodology followed, and then discusses the findings as they emerged from across the suite of case studies. The case studies themselves are presented in an annex.

[1] BEIS (2019). Regulation for the Fourth Industrial Revolution. Available online at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/807792/regulation-fourth-industrial-strategy-white-paper-web.pdf

Background

This section presents the background and context to this study.

The aim of this research was to gain evidence on regulators' use of emerging technology to improve the efficiency or effectiveness of regulation in a transformational way. Using this evidence, it has looked to identify best practices and practical considerations for regulatory applications of emerging technologies.

Experimentation with emerging technologies to date has been slower in the public sector than in the private sector. The experience of public sector agencies has also been characterised by comparatively lower sophistication than observed in much of the private sector, as noted in a recent report on the use of AI in US federal administrative agencies[2]. Several scholars have expressed concern about the fact that regulators were not adopting technologies that could help them regulate activities that grow more complex as a result of those same technologies[3] [4].

The literature discusses several obstacles to explain the slower pace of technology adoption in the public sector. These include:

- A lack of organisational capacity, appropriate resources and understanding of potential applications to transition to a greater use of new technology effectively. Beyond the challenge of transitions themselves, the application of emerging technologies is expensive (or perceived to be). It is also by definition not too well-tested, and therefore perceived to be risky.

- Concerns around data quality, as well as privacy and security have also prevented the greater use of data analytics technologies at regulatory agencies[5] . More generally, ethical concerns have been voiced about the use of algorithms to support or implement public policy[6], which several regulators might have viewed as additional risks to their reputation and legitimacy. For example, algorithms can incorporate unintentional biases based on the data they are fed, and for certain types of machine learning, it can be difficult to ascertain what rules were followed within an algorithm and how decisions have been made.

- Algorithms require a precise and clear decision-making logic, and this has been seen to clash with the ambiguity of many policies and the discretion and tacit knowledge that has often characterised regulatory work[7]. Therefore, using algorithms to support regulatory decision-making requires setting out rules and firming-up principles that may not have been set so clearly in the past. In some cases, this will mean pushing civil

[2] Engstrom, D.F., Ho, D.E., Sharkey, C.M., Cuellar, M.-F. (2020) Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies. Report to the Administrative Conference of the United States. https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf

[3] Schwab, K. (2017) The Fourth Industrial Revolution. Crown Publishing Group, New York.

[4] Engstrom, D.F., Ho, D.E., Sharkey, C.M., Cuellar, M.-F. (2020) Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies. Report to the Administrative Conference of the United States. https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf

[5] Griffiths, A (2019) The Practical Challenges of Implementing Algorithmic Regulation for Public Services, in Yeung & Lodge (eds) Algorithmic Regulation, Oxford University Press.

[6] Yeung K, Lodge M (2019) Algorithmic Regulation, Oxford University Press.

[7] Veale M and Brass I (2019) Administration by Algorithm? Public Management Meets Public Sector Machine Learning. In Yeung & Lodge (eds) Algorithmic Regulation, Oxford University Press.

servants toward defining those rules and setting an agenda in a way that goes beyond their remit and may go beyond the instruction provided by legislation[8].

Together, these challenges may have contributed to delaying the adoption of and experimentation with new technologies in the regulatory world.

The literature also highlights how some regulatory areas are inherently difficult to automate or augment with given technologies that might work well in other areas. First, to 'train' algorithms to identify and then predict non-compliance/poor performance among regulated entities, it is necessary to have data on a sufficiently large population of entities. While many regulators may have data on a large population of entities (e.g. education, food safety or health and safety), those operating in more concentrated sectors will not necessarily have sufficient data.

Second, for an algorithmic approach to work, regulatory outcomes should be clear-cut (e.g. 'compliant' or 'non-compliant') or at least scored consistently and clearly. This may be challenging for regulators operating where a range of outcomes are considered together and human regulators may make trade-offs between them (as is the case, for example, in the assessment of healthcare quality).

Finally, if an algorithmic approach is being used to support regulators, it is important that the data being used has a meaningful relationship to the outcomes. For example, Griffiths (2019) explains how the Quality Assurance Agency (QAA) in the UK — responsible for assuring quality standards in higher education — used machine learning to help develop a risk-based approach to monitoring higher education institutions. Despite having significant data, they were unable to find a relationship between that data and the outcomes of QAA reviews. This could help to explain why some solutions that have been developed in some sectors have not been replicated elsewhere, or have been replicated but with unsatisfactory results.

Current policy is seeking to address these issues in a variety of ways. The UK's white paper on Regulation for the Fourth Industrial Revolution[9], published in June 2019, sets out commitments to improve the funding of regulators and local authorities to support technological innovation. The Regulators' Pioneer Fund (RPF) is one vehicle for funding projects that help regulators to support innovation in the private sector and, where this provides benefits to innovating businesses, to innovate within their own processes. The GovTech Catalyst fund has also supported innovation in regulation by providing funding for suppliers to solve public sector problems using innovative digital technology. Other support has come through the initiatives of the Government Digital Service, working to develop capacity to digitalise government work and sharing information throughout its network across departments and agencies. By commissioning this report, the Better Regulation Executive (BRE) aims to provide information to regulators to address gaps and variations in knowledge about emerging technologies across regulation, particularly in relation to the experiences of other regulators and insights gained from them.

The purpose of this study has been to explore further the journey through which some public organisations have successfully developed and implemented solutions incorporating emerging technologies. There are numerous success stories to be told, whether in the UK or elsewhere. Through a more detailed investigation of 12 examples, this study has aimed to identify lessons and elements of best practice that could be applied usefully by others.

[8] Ibid.

[9] Available online at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/807792/regulation-fourth-industrial-strategy-white-paper-web.pdf

Methodology

This section summarises the methodology used to complete this project.

The first phase of the project involved a desk-based review of existing documentation on the use of emerging technologies by regulators. This scoping phase identified a long list of 45 examples of emerging technologies being used for regulation. Most of the desk research focused on grey literature, blogs, newspapers and websites, as relevant academic literature was less likely to provide information on the most recent developments. These types of sources tended to provide limited details on projects, meaning that the initial desk research contributed only minimally to the findings in this report.

Of the 45 examples identified, 26 were UK-based and 19 were in third countries. This reflects the fact that UK-based examples were prioritised in the search. However, scoping discussions with academics in this area also suggested that the UK is one of the countries leading the way in the use of emerging technologies in government, meaning that many of the most relevant examples are also UK-based.

Most examples identified appeared to have been developed between 2016 and 2020. The most mature example in the list was launched in 2010. This reflects the fact that the search focused on emerging, rather than mature, technologies.

The examples identified were a mix between applications developed largely in-house and those developed in cooperation with contractors. A small number of the examples were cases where a contractor or start-up had developed a technology which was then taken-on by a public sector organisation. In the course of the research, several examples of technologies developed by contractors or start-ups were identified, but in many cases it was not clear whether these had ever been adopted by a government and, hence, these were excluded.

Based on the initial long list, a short list of 12 examples to bring forward as case studies was agreed with BEIS. These were selected based on:

- the apparent impact/transformative nature of the use case;
- the likely transferability of the application to other regulators/areas of regulation; and
- the extent to which documentation could be reviewed and stakeholders consulted to explore the process of developing the solution and its impacts.

Both UK and foreign examples were selected. The final shortlist was chosen to cover a representative range of technologies and regulatory functions (the final list of case studies can be found at Table 1 in the following section), as well as different levels of maturity, complexity and working styles (e.g. both projects conducted entirely in-house and those with external contractors).

Potential interviewees were identified for each shortlisted case study. In some cases, organisations were contacted without a specific case study in mind, based on indications from previous research that their organisations were making innovative use of emerging technologies.

Each case study was based on 1 to 3 interviews with members of the project teams who developed the solution. In some cases, a group interview was held. A semi-structured topic guide was developed and agreed with BEIS. Interviews took place over the phone or through video conferencing and lasted between 45 and 60 minutes.

The topic guide (available in Annex 3 – Topic Guide) was divided in several main areas of interest: a general introduction to the case study, how the technology was developed (including a description of potential legal barriers, the use of resources and skills within the organisation), the project timeline, a description of use and access to data, the technology's maintenance, the impact of the technology and lessons learnt.

Case studies were written up based on the findings of those interviews and, in some cases, additional documentation provided by interviewees, following a common format.

The completed case studies were then sent back to the interviewees for review and to address remaining gaps.

Case studies were then coded using qualitative data analysis software and findings were synthesised by theme. Cases were compared against one another to identify the extent to which interviewees faced similar challenges and to identify where approaches to address those solutions had been identified. This has formed the basis for a series of good practice recommendations, presented in this report.

Discussion

The 12 case studies discussed within this report cover a range of regulatory functions,[10] sectors and technologies. The case studies have differed significantly in scope and budget, but all feature regulators and public sector organisations who are seeking to develop new tools that will improve the efficiency of their services and/or help them to develop new capabilities. Many show governments looking to make better use of the data they already collect. In others, the technology has focused on developing better ways of collecting data. Some are not focused on data, but on developing tools to directly automate or improve processes normally carried out manually. A full overview of these cases is provided in Table 2.

Table 2: Summary of case studies

Example	Timeline	Functions	Technology function	Impact
The Intellectual Property Office (IPO) and Machine Learning	Started in January 2019, work ongoing	Adjudication / Public services and engagement	The IPO explored the use of ML tools to improve processes in the registration of patents and trade marks. ML has been used to help examiners identify similar trade marks and, for patents, prior art, more efficiently than this was usually done.	Trade mark applications are expected to improve both in the quality of applications received and a more efficient examination process. It is estimated that this will lead to a 20% reduction in the time examiners take to review applications. ML is expected to accelerate the patent examination process and, through the time saved, to improve the quality of the examination.
The Rural Payments Agency (RPA) and CROME	Started in 2015, fully operational since 2017.	Enforcement / adjudication / monitoring and analysing risks	RPA used a type of supervised machine learning to develop a crop map of England (CROME) based on satellite	This supports the accuracy and efficiency of field checks the RPA needs to conduct when assessing Common Agricultural Policy (CAP) payments. It was estimated that the adoption of the

[10] This report uses the set of functions that Engstrom et al (2020) have used to distinguish different uses of AI in US federal administrative agencies: enforcement, adjudication (of benefits and privileges), regulatory research, monitoring and analysis, public services and engagement, internal management.

			images and image classification methods. This can then be used to better assess/allocate rural payments and avoid fraud.	CROME crop map as part of the enforcement regime has saved the RPA approximately £400,000 per year.
A consortium of private firms and public sector bodies' use of sensor technologies for mosquito classification (the Vectrack Project, EU)	Started in November 2019, work ongoing	Regulatory research, analysis and monitoring	The Vectrack system uses sensors and machine learning to collect and classify mosquitos according to several criteria such as gender, provenance and risk of disease in real time.	This tool replaces manual inspections with the use of sensors, automating and accelerating the data collection process, which takes hours rather than days. The data is then processed using algorithms/machine learning, providing more and better information on mosquito-borne disease risk. The use of these technologies helps to significantly reduce costs of monitoring and allows for more rapid intervention.
Ofcom and the use of blockchain for a common numbering database	Started in 2018, expected to go live in 2022	Regulatory research, analysis and monitoring	Ofcom has piloted the use of blockchain to support the development of a common database of landline numbers in the UK.	Potential benefits include: greater transparency and oversight of phone numbers in the UK and an associated reduction in nuisance and spam calls, more transparent and effective mechanisms for moving phone numbers between providers and mitigating the risks posed by the transfer to internet protocol (IP) technology.
The Finnish Government's use of a Virtual Agent Network	Started in March 2018,	Public services and engagement	The Starting Up Smoothly project	The chatbot network was intended to

	went live in May 2019.		explores the use of chatbots to provide an integrated approach to customer service between different organisations. It focused on the test case of linked chatbots to address the needs of foreign entrepreneurs looking to start businesses in Finland.	improve the efficiency and quality of the service. However, the network created was for a niche example, so no widespread impact has been felt. The main purpose of the project was to test the feasibility of the idea and offer a model for the creation of an automated, user-centred solution across government departments. Follow-on research is now exploring this wider application.
The Ministry of Justice's (MoJ's) development of an intelligent search tool	Started in September 2017, deployed in December 2018	Regulatory research, analysis, and monitoring	The MoJ developed an intelligent search tool for prison reports using natural language processing (NLP) techniques. The tool brings together reports that were previously dispersed in different locations and allows users to conduct searches that return contextually-similar results in addition to exact matches.	This tool has given staff efficient access to a large amount of data. Data would have previously been overly time-consuming to bring together. Access to this data allows staff to identify patterns across prisons and use this evidence to inform a more strategic approach to policy.
North Tyneside Council's use of RPA	Started in 2016, phased out	Internal management / adjudication	North Tyneside launched a project to apply robotic process automation (RPA) to repetitive tasks within the Council's services.	The use of RPA helped to speed up processes and improve the response times for customers. For example, the time to process benefit claims decreased from 36 days in

				2014 to 27 days in 2016. By reducing the time staff needed to spend on routine work, staff and resources could support tasks that improved the quality of Council services.
The Financial Conduct Authority (FCA)'s Data Strategy	Started in September 2019, is expected to take 3-5 years	Internal management	The FCA adopted a new Data Strategy, rather than a specific tool, which involves several organisational change projects: bringing in infrastructure, hiring internal data science teams, improving data management practices and promoting an innovation-friendly culture.	There are specific impacts related to the different activities linked to the strategy. More generally, the intended impact is to improve the FCA's management and use of its data. This is expected to lead to improvements in efficiency, as well as new capabilities.
Natural England's use of eDNA	Started in 2012	Regulatory research, analysis and monitoring / adjudication	Natural England employed the use of environmental DNA (eDNA) for sampling great crested newts, a protected species under UK and European Law.	This tool improves the time efficiency of the methods (reducing sampling time for one pond from 48 hours to 2 hours) as well as the quality of the detection, reducing the uncertainty in the results. These improvements in efficiency have facilitated the development of a new licensing regime for great crested newts.
The Environment Agency's use of webcams	No information provided	Regulatory research, analysis and monitoring / enforcement	The Environmental Agency employed the use of webcams and image differentiation	The technology increased the efficiency of the monitoring and intervention for culvert inspections, thus allowing a more efficient

			technologies to monitor culverts and respond to flood hazards.	allocation of staff time.
The Driver and Vehicle Standards Agency (DVSA)'s Intelligent Risk Rating Assessment tool	Started in 2017, live since end of 2018	Regulatory research, analysis and monitoring / enforcement	The DVSA has developed a tool that uses ML algorithms for pattern identification on MOT Testers. This provides risk ratings facilitating targeted MOT inspections based on past behaviour, allowing resources to be focused on garages with highest risk scores.	The use of this tool has allowed for more targeted inspections: previously all garages were assessed every 1-3 years. Resources have been allocated toward those at highest risk, meaning that there has been an increase in disciplinary actions. Examiners' prep time for enforcement visits has fallen by 50%, increasing time and cost efficiency of the service.
The City of Chicago and WindyGrid	Started in 2012, has been live since	Regulatory research, analysis, and monitoring / enforcement / public services and engagement	WindyGrid is a platform that aggregates City data from across departments. The platform has been used as a basis to develop features for predictive analytics, which leverage this data to improve city services.	The tool makes sense of data that previously would have been in silos. Bringing this data together and allowing for analytics has had specific impacts on both efficiency and capability across a range of government services.

The following section sets out the key findings and learnings from the case studies. The section is organised into themes emerging from the case studies.

Thinking about problems and solutions

The definition of problems or needs and the identification of a potential solution constitute the first steps in the development of innovative uses of emerging technologies for regulation. Through the case studies, the following good practices have been identified:

1. Take a problem-focused approach rather than letting preference for a given solution shape the process. This means identifying problems and considering solutions suitable for those problems, rather than identifying a technological solution and looking for ways to apply it.

2. Consider a range of options and acknowledge that emerging technologies may not offer the best solution in every situation.

3. Regularly monitor for the emergence of new technologies and market trends as off-the-shelf solutions may become available.

4. Learn from the experiences of others, including other regulators.

5. Consider the implications of scaling up prototypes, developing an operational tool, and maintaining that tool over time.

6. Consider starting small or choosing relatively “easy” first examples of implementation to build skills, confidence, and ensure buy-in from internal and external stakeholders.

7. Build on existing projects: adding new features or developing new uses for existing tools can offer a good approach to developing more sophisticated solutions.

1. Problem-focused approaches

Across many of the case studies, interviewees advocated for a **problem-focused approach** to selecting a technology or tool. In organisations, decision-making often works more with solutions looking for problems rather than the other way around[11]. Particularly for emerging technologies, many interviewees noticed a tendency for teams to focus on how a given technology could be used, rather than on the problems that needed to be resolved and what features a solution to those problems would need to have. A lack of focus on the problem and how to resolve it may lead ultimately to either a suboptimal outcome, one that is poorly adapted to the need or the problem, or to no solution at all. Mapping out problems first, having a clear conception of the end goal and exploring options (including solutions that would not necessarily involve reliance on emerging technologies, as discussed below) were cited in several case studies as tactics to mitigate this risk. Many of the case studies echo principles of user-centred development (which are at the core of the Government Digital Service's Service Standard[12]). They show how successful development of new solutions have been driven by the teams or departments within organisations whose needs were being considered. For example, in the City of Chicago case study, new analytics applications were developed alongside City departments, focusing on their most relevant needs. In several other examples (including the IPO case studies, the DVSA case study, the Finnish government’s Starting up Smoothly

[11] Cohen, M.D., March, J.G. and Olsen, J.P. (1972) ‘A garbage can model of organizational choice’, Administrative Science Quarterly 17: 1-25.
[12] Available online at: https://www.gov.uk/service-manual/service-standard

project and Natural England's use of eDNA), user testing of the tool or technologies was conducted early-on to help the development team ensure that the solution fits user needs.

In the examples reviewed for this report, a wide range of problems and needs have been at the origin of the use of emerging technologies. They include:

- a lack of staff resources to perform regulatory functions with the required frequency or on the required scale;
- concerns about a lack of consistency in the manner in which staff were performing tasks;
- the need for essential data streams being centralised to facilitate the management of a high-stakes and high-risk event; and
- a need to improve the efficiency of daily functions.

Even when the availability of a technology creates an opportunity to automate tasks or parts of tasks – as was the case in North Tyneside council, where RPA was implemented to automate a number of repetitive tasks – it was not the technology itself that defined the change. Rather, realising the change required working out a solution that would be a close fit for the needs of the users (in that case, the council and individuals within the community who use the council's services) and use of the technology was limited to applications where efficiency gains were clear.

2. Considering a range of options

Some interviewees noted that, for the public sector in particular, **emerging technologies would not always offer the most feasible solutions** to the problems that need solving. Part of taking a problem-focused approach also requires considering feasibility and looking at a range of solutions, which may include both emerging and non-emerging technologies.

There are drawbacks to the use of emerging technologies. While they can offer highly innovative solutions, emerging technologies are often costly while the impact of using them might be uncertain. The decision to make use of such technologies will therefore be influenced by resource availability and by the appetite for risk, as well as by a consideration of what is most suitable for the problem. For some regulatory applications — particularly those related to enforcement, adjudication or legally-mandated monitoring — certain technologies would need to be tested extensively and rigorously before they can be considered suitable to replace existing approaches. Not all public organisations will have the resources to invest in the required development and testing. For those reasons, some organisations may choose to wait until the technologies they would consider using become somewhat more established. As a result, they may be more affordable and tested, which means that later adopters can learn from the challenges that (more resourceful and/or less risk-averse) early adopters have faced.

Some examples within the case studies — such as the use of chatbots by the Finnish government or the use of webcams by the Environment Agency — show organisations applying established technologies and off-the-shelf solutions to their problems. In the case of the IPO's use of ML for trade mark applications, the solution includes a mix of emerging and established techniques to address different aspects of their problem. In these cases, established technologies still offered significant efficiency gains. These cases also show, however, that incorporating established technologies is not without challenges: such technologies still require resources for development and are not always an "easy" option. Ensuring that a solution is user-friendly and fit-for-purpose is necessary, regardless of how established the underlying technologies or principles might be. As noted by an interviewee in the IPO case study, it is important to ensure that established aspects of a solution also receive adequate resources for development and are not ignored in favour of emerging aspects.

3. Assessing the suitability of existing techniques

In many cases, existing techniques or off-the-shelf platforms or algorithms were used to support the development of tools. Part of considering a range of options can involve **assessing the suitability of existing techniques.** This can help to make the subsequent development process more efficient. **Regularly monitoring new relevant technologies** can help to build a basis for this, and the availability of existing tools, especially for emerging technologies, will likely change over time. As technologies become more established, it is possible that more off-the-shelf solutions become available. For example, as explained by the interviewee in the City of Chicago case study, the in-house approach used by the City's team in 2012 would be less likely today, because the market for relevant off-the-shelf solutions has expanded significantly since. Within other case studies, interviewees noted that they regularly monitor and look out for new techniques. This may take the form of keeping up-to-date with networks and industry trends, such as in the Environment Agency, or making use of away days, workshops and holding events hosted by external experts, as was mentioned by both the FCA and the MoJ.

4. Learning from others

Considering a range of options and assessing the suitability of existing techniques can also be supported by learning from the examples and experiences of others. Several interviewees emphasised the importance of **learning from others** more generally, in terms of how they faced similar problems or had similar needs, and how they resolved them. In some cases, interviewees mentioned benefitting from collaborating with other regulators or observing tools other regulators had applied (e.g. the Environment Agency learned from the experience of Network Rail, the IPO learned from the experience of the World Intellectual Property Office and of its Australian counterpart, and other cities in the United States have learned from the experience of the City of Chicago). This did not mean readily transposing a solution developed elsewhere, but rather adapting it to the needs and ways of working of the organisation.

5. Considering the implications of scaling up

When using an emerging technology for the first time or when off-the-shelf tools are not available, many organisations choose to start small, with a feasibility study, proof-of-concept or prototype. This is an important first step for establishing the potential of a tool and identifying likely challenges before additional funding is invested in developing something further. Many of the interviewees explained, however, that moving from these types of projects to an operational tool is challenging. In many cases, **scaling up requires different skillsets from those held within data science or innovation teams**. Prototypes are generally developed on limited budgets, meaning that certain features are not developed, and less resource is available for aspects of development such as user testing. For example, in the MoJ case study, the interviewee noted that their intelligent search tool had been designed as a prototype and therefore it had not been designed to handle large amounts of data. For prototypes, resources often are not set aside for maintenance and monitoring after development. It may therefore be helpful for development teams at the outset to **consider what the resource and design implications of developing an operational tool would be**, even if this is not the initial intention. Where prototypes are successful, it may be worth having a plan in place to scale up.

6. Starting small

In some cases, the emerging technology case developed was intended to serve as a model or pilot for future similar work. By starting with a pilot, organisations can learn and develop a broader approach through an initial example of implementation. For example, the Starting up Smoothly case in Finland was meant as an “experiment” to see whether a network of chatbots could be used to bring together government customer service functions. Another example is the use of eDNA by Natural England: although this project led to a fully operational use of eDNA for conservation and licensing of great crested newts, the technology and associated licensing regime are setting a model for the use of a similar approach for other species in the future. One interviewee highlighted the value of **choosing a relatively easy example for these initial models**. This allows the development team to work out basic challenges in design before expanding. Such examples illustrate that successful cases of adopting, adapting and implementing emerging technologies at regulatory bodies are often the result of experimentation and organisational learning. The process, as documented in the case studies, is characterised by trials (and failures), obstacles of various kinds along the way (e.g. staff concerns towards the technology, or lack of engagement with the change), and project teams imagining solutions, testing them, improving them along the way.

7. Building on existing projects

Many of the case studies involved **the development of an initial tool, and then subsequent projects that added to or amended that tool**. For example, the Environment Agency’s use of webcams was improved by the addition of image differencing techniques to automatically alert staff of changes. In the City of Chicago WindyGrid case, the initial tool bringing together data from different city departments was then built on over the years as new data sets were added and new analytical capabilities were developed. In many cases, follow-up projects were not yet ongoing but they were anticipated by interviewees. Building on existing projects helped in those cases to develop more sophisticated tools and to ensure that the functionality of existing tools remained relevant and up-to-date. Moreover, building on something existing is in many cases much easier than starting from scratch.

Infrastructure and tools

For many of the innovations explored in this study, applying them to regulatory work has required having at hand or building up / acquiring technical infrastructure or tools. The case studies show that these were often pre-conditions to the successful development and roll-out of innovations using emerging technologies. Through the case studies, the following good practices have been identified:

1. Have cloud-based systems in place, particularly to facilitate data-heavy or machine learning projects.
2. Develop centralised data management and storage systems, for projects that seek to leverage existing data.
3. Develop a bespoke in-house analytical platform for data science teams.
4. Consider bringing in off-the-shelf tools that support data analytics.

1. Cloud-based systems

One of the most common points noted amongst interviewees was the importance of having a **cloud-based system** in place. Using cloud-based systems means that rather than owning hardware, organisations rent hardware from a cloud provider. For many organisations, access to cloud services will therefore mean access to greater computational capacity and/or access to specialised hardware without the need to purchase this and store on site. This can be particularly beneficial for applications that would involve storing and processing large amounts of data and for machine learning (ML) applications. ML requires high processing speeds and storage capacity not generally available through traditional information technology infrastructure. In some cases, ML will also require specialised hardware — such as graphics processing units (GPUs) — and this may only be needed briefly (e.g. to train a model). In such examples, renting hardware may be a more cost-effective approach than purchasing additional hardware. In the CROME example from the RPA, the interviewee noted that certain future deep learning applications would not be possible unless there was a switch to a cloud-based system.

2. Centralised data management

Centralised data management and data storage systems have also enabled data-based innovation for some of the organisations interviewed. Having data in separate formats and in different places adds significant time to the process: data must be cleaned, accessed and brought together before it can be used. This was the case, for example, for the MoJ in the development of their intelligent search tool for prison reports. Prison reports were stored in different places and in different formats, and this proved an obstacle for development. The FCA's Data Strategy involves the centralisation of data management and storage as a crucial first step to overcome similar problems and to enact more ambitious data projects.

3. In-house analytical platforms

Some organisations interviewed noted that the development of **an in-house analytical platform** (e.g. MoJ's Analytical Platform, the FCA's Data Science Workbench) has been an important first step to the development of future tools. These platforms provide a basic, tailored software infrastructure that data science teams can then use to develop and experiment with new ideas.

4. Off-the-shelf tools

Off-the-shelf tools that support data analytics can also be brought in to support staff from across an organisation making better use of data. Such tools may not be bespoke or specific to an organisation's needs, but in some instances may provide a lower cost and tested solution for analysing data. For example, as part of its Data Strategy, the FCA has started to use Tableau™, a tool that allows for the creation of interactive dashboards of data, and is providing training to staff to use that tool. Interviewees at the FCA noted that the use of such tools can help staff from across the organisation become more comfortable with data science. The City of Chicago's WindyGrid tool, which combines a map of the city with city data and analytical tools, was similarly intended to be used directly by staff, and while this was not an off-the-shelf tool, the interviewee noted that the market for such tools has expanded in recent years and an off-the-shelf solution might have been preferable in the current context. Off-the-shelf tools may also form the basis of other solutions: benefits and challenges of this are discussed in more detail in the section on in-house teams and external contractors below.

In-house teams and external contractors

The following section discusses how in-house teams at regulatory and other organisations interviewed for the case studies contributed to the development of emerging technologies for regulation, and the benefits and challenges to different approaches to building project teams. Through the case studies, the following good practices have been identified:

1. Consider the project's needs carefully: what technical and non-technical skill sets are needed, and the extent to which existing teams can provide these.
2. If there is capacity to conduct work in-house, consider what challenges this might bring and the extent to which the use of an expert external contractor may add value to the project.
3. If there is a need to hire new staff, address this as early as possible and plan for challenges.
4. If there is a need to engage external contractors, consider and mitigate against the risks of becoming dependent on those contractors, or of ending up with costly solutions.

1. Considering project needs and 2. Considering likely challenges

The case studies illustrate a range of team structures for developing emerging technologies. Some examples were led and carried out entirely in house. Others were carried out mostly by external contractors. Others used a mix of in-house and external expertise.

Interviewees shared various considerations on the benefits and challenges of different approaches.

Projects that were carried out entirely in-house corresponded to cases (a) where investments had already been made in building up the relevant technical teams or (b) the scope of the projects was limited, and therefore all the work could be carried out internally. Interviewees mentioned several benefits to leading and/or carrying out work in-house. These include:

- Being able to **approach projects and data in a longer-term or more continual way.** Often, contractors are employed for short-term pieces of work, and there may not be an opportunity to envision something longer-term, or the knowledge to build on previous cases. In-house teams, by contrast, are more easily able to follow up on previous projects or design projects with future work in mind.
- **Mitigating or preventing issues of data privacy and security that might otherwise arise should a private third party be involved.** This is especially pertinent for organisations that handle data on consumers or businesses, and is echoed in research in the US: "in-house expertise promotes AI tools that are (…) more likely to be designed in lawful, policy compliant and accountable ways"[13].

[13] Engstrom, D.F., Ho, D.E., Sharkey, C.M., Cuellar, M.-F. (2020) Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies. Report to the Administrative Conference of the United States. https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf p. 7

- **Facilitating the participation of internal staff and creating opportunities for them to learn new skills.** Developing new skills in staff can help not only in carrying out future work in-house, but can also facilitate the development of new and innovative ideas.
- Helping teams **better understand their own processes and how to optimise them**. The development of tools and technologies often requires careful reflection and can provide new insights on processes. Engstrom et al. (2020) note that developing tools in house means "better-tailored" solutions than if those solutions were developed by contractors.

On the other hand, many of the organisations consulted in this study made some use of external contractors to address the lack of in-house skills or capacity. Several interviewees across cases noted that the use of external contractors with expertise in a specific technology or approach had benefits, including:

- Implementing a **more creative approach to development**, stemming from expertise and experience that the contractor would have gained in a variety of other sectors.
- Encountering **fewer challenges in development,** as contractors will often have completed similar work in the past and will be better able to mitigate risks.
- **Accessing tested off-the-shelf solutions and processes**.
- **Learning from external contractors**, helping to upskill in-house teams.

3. Developing in-house teams

Despite the noted benefits of in-house teams, getting to a point where using in-house teams is a viable approach was difficult for many of the organisations interviewed.

Many indicated challenges in hiring or retaining staff, due to **high competition for talent and the inability to match private sector salaries**. For example, Ofcom noted that although they benefited from London's significant technology and fintech sectors when building a team for their blockchain project, they found it difficult to compete for that talent alongside private firms. Similarly, in North Tyneside's work with robotic process automation (RPA), they trained existing staff in RPA rather than recruiting externals, but then found that some of the staff who had undergone RPA training ended up leaving for higher paid roles elsewhere.

In cases where organisations had developed significant in-house teams, they noted that this required upfront investment. The FCA's Data Strategy includes an aim to build out and develop their data science team using a hub-and-spoke model, with a central "hub" incorporating the FCA's central data services team and the advanced analytics team and "spokes" made up of eight data science units with their own managers, technical specialists and a number of graduates. Interviewees noted that hiring into and building these teams **required investment and the support of senior management**. For existing in-house teams, some interviewees noted the importance of away days and training events to upskill staff and promote the exploration of new ideas.

Some interviewees noted that, in addition to having technical teams in place, cultural changes were also important to fostering innovation. Interviewees at the Environment Agency noted that, within the organisation, technologies that are "emerging" are perceived by some as costly, complicated to implement and their efficacy uncertain. According to interviewees at the FCA, one important component of their Data Strategy is **creating a more innovation-friendly**

environment: allowing more time for innovation, encouraging people to try new things and accepting that some innovation projects will fail. An interviewee at the IPO noted that there was also a challenge related to how business benefits are defined, particularly for ML. Because it is difficult to know from the outset how well an algorithm will perform, it is difficult to guarantee a business benefit in terms of hard figures.

Where using in-house teams for development, it may also be important to consider **what non-technical skills are needed**. In the Finnish government's development of a chatbot network, the interviewee noted that although they were able to upskill staff with the required technical skills to create content, they did not initially anticipate that staff would also require training in copywriting to produce the content.

4. Using external contractors

Using external contractors also came with challenges for several of the organisations interviewed.

In some cases, the public body identified external contractors through a competitive procurement process. Although many achieved success with this approach, the experience with procurement discussed in the Environment Agency case serves as a counterexample: the initial procurement competition led to a solution that was more costly than anticipated and the project subsequently stalled. When a second procurement competition was later set up, the interviewee noted that significantly more time and effort was taken to identify a more appropriate contractor than had been anticipated.

Many of the organisations worked with external contractors with whom they had **existing partnerships or collaborations that might have been unrelated to the project in the first instance**. In some cases, these were private firms, in other cases academic organisations. For example, Natural England had ongoing collaboration with conservation groups who supported their work on great crested newts. For those, reliance on an external contractor for the development and deployment of the innovation was therefore contingent upon these pre-existing relationships rather than an assessment of the pros and cons of running the project internally. For some examples (e.g. the North Tyneside case and the Vectrack project), projects were conceived and led by the contractors themselves. In the North Tyneside case, the interviewee noted that the contractor's decision to pursue the technology was driven in part by commitments made to North Tyneside Council to increase the efficiency of public services.

For some examples, external contractors also provided software or platforms that acted as the basis for development. In two cases where this was mentioned, the contractors also provided training to staff for the use of that software. In such cases, there is **a risk that organisations become dependent on that software or platform**. In Ofcom's exploration of blockchain, this was cited as a reason to avoid using an external contractor. Developing their own software or platform, even if this would take longer, allowed them to future-proof their tool against licensing or dependency issues. The issue of dependency may also relate to any software the organisation might have been using before implementing the change required: this occurred in the North Tyneside Council case, where changes to existing software used for council services meant that the automations they had developed no longer functioned. This issue was also mentioned in the Environment Agency example, where the project stalled due to an external contractor going out of business.

Organisational buy-in and staff engagement

Another set of challenges commonly cited across cases was engaging and educating their own staff about the proposed technological solution and achieving buy-in from the wider organisation. Through the case studies, the following good practices have been identified:

1. Collaborate with a wide range of staff early-on in the development process.
2. Have external contractors host informational events for staff.
3. Illustrate the functionality of the tool at an early stage in the process (e.g. showing rather than telling).
4. Rely on operational staff (first users) to explain the tool and its benefits to other staff.
5. Pilot tools internally with staff (including where staff are not the intended end users).
6. Upskill staff and, where possible, involve staff in the development process.

1. Engaging with staff early

An initial hurdle for many projects was convincing staff that tools were reliable and valuable, particularly where proposed projects would be altering the way staff work. Where this was the case, development teams needed to support staff to understand the purpose of the tool and how it may change their way of working. **Having staff engaged early in the process** was an important first step. For cases where the use of a new technology will lead to significant change in regulatory approach, additional engagement may be necessary to address concerns specific to the regulation. For example, in the Natural England case, the technology developed has led to a new licensing approach for great crested newts, and additional work had to be done to engage with environmental and conservation NGOs on these changes.

2. Holding informational events

In some cases, interviewees also mentioned a challenge in managing staff's expectations: particularly for tools based on AI or machine learning, some interviewees noted instances where staff felt the technology would be able to deliver more than it could (e.g. expecting AI to mean full automation, rather than partial automation of specific processes). Literature on perceptions of AI also suggests that in many cases, the general public's response to AI is often rooted in anxiety and many assume that AI means a complete replacement of human labour with computers or "robots".[14] Managing expectations about the use of emerging technologies can therefore go both ways: addressing concerns as well as correcting unrealistic expectations. An interviewee from the IPO referenced **a general session on AI** for staff held by their contractor. According to the IPO interviewee, this session was well received and helped many in the organisation better understand what AI could address and what it could not.

[14] Cave, S., Coughlan, K., & Dihal, K. (2019, January). "Scary Robots" Examining Public Responses to AI. In Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (pp. 331-337).

3. Demonstrating the functionality of the tool

In some cases, interviewees noted that providing a demonstration of the functionality of the tool to staff was important to managing expectations and collecting feedback. For example, in the MoJ case, the interviewee noted that having an early model of the tool to show staff the usage and potential helped to build engagement and interest.

4. Relying on staff to share information

Where staff have been engaged and the value of a tool has been demonstrated, this can help spread interest and understanding across an organisation. For example, in the City of Chicago case, successful use cases of data analytics applications helped publicise the potential value of analytics across city departments.

5. Internal pilots with staff

Internal pilots are part of user testing, discussed in more detail below, but they can similarly be important for demonstrating the use and benefit of a tool internally. This may also be the case where staff are not the intended users. For example, in the Finnish government's use of a Virtual Agent Network, the tool developed was meant for agency customers, but initial testing was done with staff internally. According to the interviewee, this helped to gather initial feedback, as well as improve understanding of what the tool could and could not do and achieve buy-in across the wider organisation.

6. Involving staff in the development process

The North Tyneside Council case provides one example of successful upskilling: subject-matter experts within the Council's teams were trained in robotic process automation (RPA) and were able to develop their own automations specific to their work. The interviewee did explain that this was not always successful, and some subject-matter experts were not able to complete the training. For those that were, however, this approach helped to build engagement and meant that staff could develop simple automations themselves. An interviewee from the Environment Agency also emphasised the **benefit in having innovations developed in part by on-the-ground staff, alongside innovation teams**: this not only builds engagement, but also builds ownership, meaning that staff will be more likely to use these tools.

User testing

Tools intended for widespread staff or consumer use require some degree of user testing. In a few cases — notably those conceived as prototypes — interviewees felt that they had not done as much user testing as they should have due to a lack of time, resource or access. Testing is an important part of engaging staff, but it also ensures that the design is relevant and functional. Through the case studies, the following good practices have been identified:

1. Include user testing and feedback early-on in the design process.
2. Plan for the extra time user testing is likely to take.
3. Identify and engage likely end-users as soon as possible.

1. Early user consultations and testing

Some examples, particularly those where building engagement with users was of high importance, took an early and intensive approach to user testing. In both the DVSA and Ofcom case studies, **early user consultations and testing were core to the design process**. In the case of the DVSA, this meant that there was little need for upskilling when the tool was deployed. This suggests that early testing can help to improve accessibility and reduce training costs, likely increasing actual use.

2. Considering the timelines

In many cases, interviewees noted that **conducting user testing is time consuming**. For emerging technologies, end users may also be unfamiliar with certain aspects, making it harder for them to provide constructive feedback. Ofcom found this to be the case when discussing blockchain with the telecoms industry. To help address this, Ofcom created a special repository where interested parties could download the source code and experiment with it to better understand the functions and accelerate their own innovations.

3. Engaging likely end users as soon as possible

Conducting user testing with actual end users may also be challenging when the end users are not within the organisation. For example, for the Finnish government's Starting Up Smoothly project, the intended end users were foreign entrepreneurs looking to start a business in Finland, and this was a difficult audience to identify and engage. Even where end users are within an organisation, they may have limited time to test and provide feedback. Preliminary testing with available staff may be helpful in such cases to provide initial feedback (and can also help to build organisational buy-in), but testing with end users will provide the most relevant results. It may therefore be useful **to identify and engage likely end users as soon as possible**.

Data-driven technologies

The following section discusses challenges and best practices related to the use of data-driven technologies in regulation. Based on the case studies, some good practices for data use include:

1. Set strict data validation requirements within tools.
2. For external data, consider potential intellectual property issues early in the process.
3. Consider the implications of using personal or sensitive data.
4. Where collecting new data, expect additional challenges and costs.

1. Data validation requirements and high-quality data

Having good quality data is essential for ML or data analytics applications. ML applications in particular require significant amounts of data. Even in examples where data was well-structured and managed in a centralised way, interviewees recalled challenges related to ensuring data quality, particularly where applications relied on large amounts of data. Many of the cases made use of historical data and therefore development teams had little control over formats and structure. For the future, some interviewees noted a need for **stricter data validation requirements to improve the quality and organisation of data being collected**. For example, building systems so that data in certain formats is rejected or which stipulate certain formats. Initiatives such as the FCA's Data Strategy include work to improve data management practices from the outset, so that data is accessible and high quality for future data-based projects.

2. External data sources

Several examples — notably those that involved **mapping** — required **using data from external sources**. In both the Natural England and RPA cases, the use of this data entailed intellectual property rights, while the intention was for the end solution to be open access. In the Natural England case, this was addressed by working through these issues and developing appropriate usage agreements. In the RPA case, this was addressed by identifying an alternative approach to mapping that would be suitable to their work and not subject to intellectual property rights. Coming to appropriate usage agreements can take time and lead to delays in the process, but this may be necessary if suitable alternatives are not available.

3. Personal or sensitive data

All cases included in this report avoided **the use of personal or sensitive data**, meaning there were fewer hurdles in terms of data privacy laws. For applications where this cannot be avoided, it is important to consider what requirements will need to be fulfilled to use data appropriately. This may also mean relying more on in-house teams for such projects: Engstrom et al. (2020), among others, discuss issues associated with sensitive data, including the greater risks of data leaks that may stem from relying on external contractors to develop and manage AI solutions within government.

4. Collecting data for machine learning or data analytics

Examples where **teams collected their own data** for ML or data analytics applications also came with challenges and additional costs as compared to examples where teams were able to use existing data. ML and many data analytics applications require large volumes of data to train algorithms or to produce robust findings. Depending on what type of data needs to be collected, this can be resource intensive. For example, the Vectrack project, which uses sensors and ML to trap and catalogue mosquitoes, required algorithms to be trained with a significant number of live mosquitoes with divergent characteristics. To train the algorithms appropriately, mosquitoes with specific characteristics and divergent morphological traits need to be reared in laboratories. Another example, Natural England's use of eDNA, required rigorous testing of the data collection methodology before it could be used for regulatory applications. However, in both examples, it is the application of emerging technologies (sensors for Vectrack; eDNA for Natural England) that allowed for large volumes of new data to be collected.

Measuring impacts

None of the case study examples included in this report have been the subject of a formal impact evaluation, although some either will be evaluated or have undergone another type of assessment or evaluation. One example, the FCA's data strategy, will be evaluated at each stage as the strategy progresses. Another example, the Finnish government's Starting Up Smoothly project was evaluated, but this evaluation focused less on impact, and more on the project process and implications for future work. Alongside this, in many cases, organisations were collecting relevant quantitative data on impact. For those examples that aimed to automate or augment a process that used to be conducted by humans only (such as the use of RPA by North Tyneside Council or the development of search tools by the IPO and the MoJ), many interviewees were able to provide metrics related to **improvements in processing times**. In some cases, these metrics were calculated early on in development, as part of determining feasibility. For instance, in Natural England's development of eDNA as a sampling method for great crested newts, initial research determined that sampling using traditional methods would take approximately 48 hours per pond, while using eDNA methods would take 2 hours per pond.

Some interviewees explained that impacts related to new capabilities or improved quality of delivery were harder to measure. Many impacts related to an **improvement in the quality of services**: reducing human errors, providing better information to customers, creating a stronger evidence base for policy decisions. These types of impacts are harder to measure quantitatively, and across cases, interviewees were only able to provide anecdotal information on such impacts. In such cases, evaluations could use theory-based methods.

For both quantitative and qualitative measures of impact, **organisations might benefit from taking a more robust and formal approach** to these. By setting out likely impacts early on and devising methods for measuring these impacts, the full value of such projects can be better understood. In turn, this can help innovation teams improve their processes and improve their business case for future projects.

Two examples have **directly impacted the approach to regulation**: the DVSA's risk assessment tool has led to a system of more targeted inspections and Natural England's use of eDNA has facilitated a new approach to licensing. In both cases, the impacts of those changes to regulation have or will be evaluated. In the case of Natural England, the new approach to licensing marks a distinct change from the previous approach, rather than an improvement to the previous approach, and the results of this will have implications for Natural England's future work. Where projects have these types of impact, conducting robust evaluations is of even greater importance.

Several cases have had further impacts on efficiency and capability by **setting a precedent for future work** — and perhaps helping to inspire the innovation-friendly environment that several interviewees indicated is important. For the DVSA case, the interviewee noted that their tool had created a further interest to look at other uses of ML and AI within the DVSA. In other cases, there is a sense that the initial application will be continuously added to, amended and improved. For example, the interviewee at the IPO indicated that potential future developments of their trade mark tool might include automated rejections on absolute grounds or an extension of the approach to designs and patents.

Conclusions

Regardless of regulatory function, size, scope and budget associated with these cases, many of the organisations involved have faced similar issues. For example: choosing the right solution, building appropriate teams, engaging staff and/or users in the process and the need to consider impact are important elements of applying emerging technologies to regulation.

At the same time, the diversity of case studies examined shows that there are a variety of ways to apply emerging technologies to regulation. Technologies have been applied to improve efficiency in internal administrative processes, to improve the efficiency of data collection and to exploit data already held by agencies to support better analysis and risk assessment. The cases studied ranged significantly in size, duration and complexity, but all had a positive impact on improving the effectiveness and/or efficiency of regulatory processes. In many instances, a foundation for future innovation by the regulator was also created.

Although the case studies included here all show regulators applying technology to their work, they also indicate that regulators are in very different positions when it comes to applying emerging technologies to regulation. Some are approaching this in a top-level way (such as the FCA's development of a data strategy), others are taking a more problem and technology-specific approach. Many of the more ambitious cases have been dependent on improvements in infrastructure and capacity, such as having the right IT and data systems, and having the skills. Where regulators in other cases have struggled, it has often been due to lacking certain tools or skills within the organisation. This suggests that beyond the project-specific practices identified in the findings of this report, many organisations could also benefit from looking at what top-level changes could be made to better support innovation.

Across projects, having access to sufficient time and resources for development was often a challenge. Where time and resources are finite, one way to address this challenge is by having a clear concept from the outset of what steps will be important to the development process. This can help ensure that resources are allocated appropriately and that the project's scope is suitable. Not all of the best practices presented in this report will be relevant to regulators looking to incorporate emerging technologies in all contexts. However, a general best practice approach illustrating the key steps in the development process can be identified. A 'roadmap' bringing together the best practices identified through the cases studied is illustrated in Figure 2 below.

Figure 2: Best practices for the use of emerging technologies for regulation

01 Clearly identify the problem or need to be addressed

Once the problem or need is identified, it can be helpful to map out what this entails and ensure that those responsible for innovation have a good understanding of it, and more generally of regulatory processes.

02 Identify suitable approaches and technologies to address identified problems or needs

This can be done in different ways: engaging external experts, monitoring what others are doing, keeping data science and innovation teams engaged with new ideas through workshops, away days etc. For some tools, an initial scoping phase of the project to develop different possible solutions is helpful, particularly where proven techniques or off-the-shelf options are not apparent.

03 Treat the previous two steps as an ongoing conversation

Not all "problems" will be obviously problems, there may be existing processes that work well but could be enhanced. Sometimes the problem may not be the starting point – but it should be the deciding factor on what to pursue and in what way.

04 Clearly identify resource needs

What expertise is required (both technical and non-technical)? How much effort will this take? Is it possible to carry this out in-house? If it is feasible to carry this out in-house, what would be the benefits of receiving support from externals and what would be the risks? Ensuring that there is a plan for any resources that may be required after development (e.g. for maintenance and monitoring).

05 Consider resource availability

Is the funding available considering those resource needs? If not, are there alternative options for receiving funding (e.g. through other organisations with an interest in the outcomes or through innovation grants)? What aspects of development could be altered to match resource availability (e.g. leaving out certain features, altering the composition of the development team)? What impact would this have on the work?

06 Put together a demonstration to share with internal staff and decision makers

Once the initial idea has been developed, this can help to illustrate the tool's functions. This can provide useful early feedback, help to engage decision-makers and build a business case.

07 Clearly identify the user group and set a plan for engagement as soon as possible

This will be important for testing and improving the tool. Some user groups may be harder to reach than others, and many potential users may not have much time or capacity to engage in testing. Reaching out early can therefore help to mitigate such risks.

08 Clearly identify any likely data needs and challenges

Interviews suggested that teams often underestimate the time and effort that dealing with issues such as licencing agreements takes. Consider this from the outset.

09 Think about future uses and their implications

Could this lead to any connected changes in approach or regulation in future? For example, will a more efficient or timely data collection process mean that data can now be leveraged to improve regulation?

10 Identify desired impacts and how these can be measured

Having a method for evaluation in mind can also help to define what intended outcomes are and how those can best be achieved. Knowing planned evaluation methods also makes it possible to capture any baseline data needed at the start.

References

Cohen, M.D., March, J.G. and Olsen, J.P. (1972) 'A garbage can model of organizational choice', Administrative Science Quarterly 17: 1-25.

Engstrom, D.F., Ho, D.E., Sharkey, C.M., Cuellar, M.-F. (2020) Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies. Report to the Administrative Conference of the United States. https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf

Griffiths, A (2019) The Practical Challenges of Implementing Algorithmic Regulation for Public Services, in Yeung & Lodge (eds) Algorithmic Regulation, Oxford University Press.

Schwab, K. (2017) The Fourth Industrial Revolution. Crown Publishing Group, New York.

Veale M and Brass I (2019) Administration by Algorithm? Public Management Meets Public Sector Machine Learning. In Yeung & Lodge (eds) Algorithmic Regulation, Oxford University Press.

Yeung K, Lodge M (2019) Algorithmic Regulation, Oxford University Press.

Annex 1 – Case studies

These are captured in a separate document.

Annex 2 – Glossary

This glossary defines some of the more technical terms used in this report and in the case studies.

Algorithms – are a set of rules, formula, or a finite sequence that defines and instructs computations.

Application Programming Interface (API) – is an intermediary tool that aids the communication of data between different systems. It allows the transition and translation of requests and responses across different systems.[15] The API takes requests, communicates what to do to the system, and then brings the response back to the users. APIs facilitate the interaction between applications, data, and devices.[16]

Artificial Intelligence – covers technologies that enable machines/computers to do tasks that typically require human intelligence. A system incorporating artificial intelligence will analyse its environment and take actions with some degree of autonomy in order to achieve the goals that were programmed in their code.[17]

Blockchain – is a shared database filled with entries that must be confirmed by all users of the database and encrypted. Blockchain is a type of distributed ledger database formed by a sequence of "blocks" that are added to the chain of transaction records. A distributed ledger technology does not necessarily have this feature.

Brittle (Software Brittleness) – refers to the increasing fragility of some older software which are seemingly operational but fail when presented with unusual or altered data.

Chatbot – a software designed to substitute human verbal interaction with automated processes. This is achieved using artificial intelligence and natural language processing systems in order to simulate the software's conversational skills and understanding of natural language inputs[18].

Clustering - is a technique to organise large amounts of data, setting it into significant groups in which patterns are detected, i.e. "clusters". The model maximises distance *between* clusters, while minimizing distance *within* the same cluster between similar data points. The model attempts to detect natural statistical separation within a dataset. The similarity of data samples is defined by the distance between vectors. Clustering works well on large datasets with multiple features. Output is optimised through adjusting input parameters.

Cloud – is a term used to describe a network of servers that can be accessed remotely regardless of one's location. Clouds can also be defined as online pools of resources which can be used to store and manage data, or to support the running of different applications or services. They replace the local storage space of a personal computer with an online-

[15] Gazarov, Petr (2019) "What is API? In English, please" https://www.freecodecamp.org/news/what-is-an-api-in-english-please-b880a3214a82/

[16] MuleSoft (015) "What is an API?" https://www.mulesoft.com/resources/api/what-is-an-api[17] High-Level Expert Group on Artificial Intelligence, (2019), *A definition of AI: Main capabilities and scientific disciplines,* European Commission, Brussels.

[17] High-Level Expert Group on Artificial Intelligence, (2019), *A definition of AI: Main capabilities and scientific disciplines,* European Commission, Brussels.

[18] The Ultimate Guide to Chatbots – Drift.com Available at: https://www.drift.com/learn/chatbot/#pillar-start

accessible storage space, which can be retrieved by multiple devices with the necessary credentials.[19]

Cloud-based data lake – is a cloud-based repository that allows storage of both structured and unstructured data.

Cloud Computing – refers to on-demand availability of IT services (applications, storage or programmes) through cloud access, and thus renting rather than buying hardware.

Computational Semantics – is a branch in computer science that researches how to automate the process of semantics and human communication. One of the central questions of this discipline is how to represent meaning in ways understandable by computers [20]. Computational semantics aims to perform automated meaning analysis of natural language, and it has a prominent role in the natural language processing (NLP) applications [21].

Data analytics – refers to the techniques used to analyse raw data to recognise patterns and identify behaviours[22]. Data analytics entails the process of extraction, gathering and categorisation of data for inference. The techniques involved in these processes vary depending on the desired end results and structure of the data.

Deep learning – is an area of machine learning. In machine learning, the algorithm is given a set of relevant data to analyse and learn from it in order to generate predictions or to group data objects. In deep learning, the algorithm is given raw data and it is programmed to partially solve feature extraction issues, particularly relevant for unstructured data such as images, speech and text. Deep learning models typically require a substantial amount of data in order to generate accurate predictions. Although all statistical machine learning improves with higher amounts of data, DL data requirements are more substantial. For that reason, smaller datasets work better with classical ML methods rather than DL.

Distributed ledger technology – is a technology that enables the creation of a de-centralised database. Instead of storing data in a single place, the data is divided and entrusted to multiple users and it creates a specific protocol. This generates a secure environment in which any change has to be replicated by every user in order to occur. In this way, fraud can be avoided and there is a high level of transparency surrounding the transactions/exchanges that occur in such a system.

Dynamic Time Warping (DTW) – is an algorithm technique used in time series analysis to compare similarities in different arrays that do not synch up perfectly. Dynamic Time Warping is applicable in a wide range of domains, such as speech recognition and financial markets data.

Environmental DNA (eDNA) - is a sampling approach that extracts the genetic material directly from an environmental sample (soil, sediment, water, etc.) without the need to interact directly with the biological source. This process is enabled by the analysis of the environmental traces left by the organisms interacting with the environment. This, paired with DNA

[19] Microsoft, 2018, *What is the cloud*, Available at https://azure.microsoft.com/en-us/overview/what-is-the-cloud/]
[20] Blackburn, P. and Bos, J., 2005. Representation and inference for natural language. *A first course in computational semantics. CSLI*.
[21] Bos (2011) "A Survey of Computational Semantics: Representation, Inference and Knowledge in Wide-Coverage Text Understanding" *Language and Linguistics Compass*, 5/6. Available at: https://www.let.rug.nl/bos/pubs/Bos2011Compass.pdf
[22] Frankenfield, Jane (2019) "Data Analytics", Investopedia. Available at: https://www.investopedia.com/terms/d/data-analytics.asp

sequencing technologies, allows for a non-invasive and effective approach to biodiversity monitoring.[23]

Geographic information system (GIS) – is a tool to understand geographical patterns and relationships. It involves bringing together location-based data and displaying it relative to Earth's surface. The GIS representation can be in the form of layers, representing divergent kinds of data in a single map, such as urban buildings and natural elements of the environment.[24]

Graphics Processing Unit (GPU) instance - is a single-chip processor that alters the devices' memory in order to accelerate the creation of images and thus boost the image performance. This feature is particularly relevant in applications that require a rapid succession and feedback of graphics (such as videos).

Intelligent Virtual Agent – are animated chatbots that are generally implemented to aid customer service on online platforms. These agents can be supplemented with AI-powered conversational tool to improve their service and efficiency[25]. **Internet of things (IoT)** – is a technology connecting ordinary smart devices or objects incorporating sensors to a computer network and to the internet. This allows devices to communicate with each other and to collect and exchange data in real time. By collecting data in this manner, large volumes of data can be extracted from objects in a short period of time and with little effort.

Machine learning (ML) is a subset of artificial intelligence. It refers to technology that enables machines to learn from the data that is inputted inside the machine and acquire skills to perform some tasks with partial human involvement.

Mainframe-based systems – refers to large computing systems, used primarily by big businesses and large organisations that allow for the rapid processing of substantial amounts of data.

Natural Language Processing (NLP) – is a technology that facilitates communication between computers and humans, working as a translator between machine code and human language. For example, NLP makes it possible for computers to extract information from natural texts, or support human interpretation by classifying text.

Natural Language Understanding (NLU) – is a subarea of NLP that relates to the comprehension and understanding part of natural language by the machine. Similarly to NLP, NLU accepts text in a natural language as input and extracts information that can be used in downstream tasks. NLU basically refers to the core NLP tasks: it allows the system to interpret the meaning of the natural language, extracting useful information from the text [26].

Neural network (neural net) – is a class of machine learning algorithms that can learn to produce target outputs from specific inputs by analysing training examples. Neural networks can be trained in both a supervised and unsupervised manner. Neural network models differ greatly with respect to their architecture, but all use a sequence of linear transformations

[23] Philip Francis Thomsen, Eske Willerslev, 2015, *Environmental DNA – An emerging tool in conservation for monitoring past and present biodiversity,* Biological Conservation, Volume 183, Pages 4-18, ISSN 0006-3207, Available at https://doi.org/10.1016/j.biocon.2014.11.019
[24] National Geographic, Resource library encyclopaedic entry. Available at: https://www.nationalgeographic.org/encyclopedia/geographic-information-system-gis/
[25] Chavez, David (2019) "A Buyer's Guide to Intelligent Virtual Agents" Avaya.com Available at: https://www.avaya.com/2019/10/buyers_guide_virtual_agents/
[26] https://expertsystem.com/natural-language-understanding-different-nlp/

interspersed with nonlinearities. There are several types of building blocks that are common: fully connected layers, attention layers, convolutional layers. These building blocks are combined using nonlinear transformations and are arranged hierarchically or in a recurrent way.

Open Source – is a type of licence that is used for disseminating software code. This licence allows for the free and open distribution of the source code behind a software. Software that is derived from the original open source software must then also be freely distributed.

Polygon vector data – are a subset of vector data that are comprised of the types of geospatial data used in spatial representations. Polygon vector data (divergent from point vector data and line vector data) are constituted by three or more connected vertices, thus forming closed polygon shapes. They are used to describe areas.[27] Vector data allow the representation of discrete geographical locations into defined spatial objects.

Predictive analytics (or predictive modelling) – is a process that uses data analysis and modelling to predict future events or behaviours based on past performance or patterns. Generally, this technique involves assigning a probability or risk score to selected events depending on their likelihood of occurrence.

Random forest – is a type of supervised machine learning algorithm. In technical language, processes are being described as ‘decision trees’. The "forest" is a group of decision trees. Random forests are multiple decision trees that are merged together to draw a more accurate and stable prediction from the data it analyses.

R-CNN algorithms – are neural networks algorithms for object detection, designed to create and classify feature vectors for image recognition. The output of this process allow machine to recognise the objects contained in an image.

Robotic Process Automation (RPA) – is a technology that allows the user to configure a software (also commonly referred to as “bots” or “robots”) in order for it to carry a specific process in an automated manner. The “bot” can be used to mimic the actions that a human would take to complete a specific task. For example, RPA can be used to search for information in a specific system and then input it into a form. This replaces the manual inputting normally done by a human. The technology can be built on various processes and can also overlay on other software applications.

Scraping – is a process of data extraction through which a computer gathers and computes data from human-readable interfaces. The best-known application is web scraping, whereby the content of websites is translated into large amounts of data for computational use.

Source code – is a text, or code, written in programming language that provides a set of commands to be executed by a machine or programme.

Species Distribution Modelling – is a quantitative spatial technique related to quantitative biogeography. It uses tools that model and predict the environmental suitability for species, combining data on species occurrence and environmental conditions.

VGG 19 neural network – is a trained convolutional neural network used for image classification. The name derives from the team of developers – Visual Geometry Group (VGG) – and the number of layers in the network.

---

[27] Dempsey (2017) available https://www.gislounge.com/geodatabases-explored-vector-and-raster-data/

Word/Sentence Vectors – are numerical vectors that represent the meaning of a word or sentence. They are vectorised representations of words and/or sentences that capture certain lexical, syntactic and semantic properties of a word/sentence. This is more advanced than traditional NLP, as it allows machines to capture the meaning and association of a word or sentence rather the mere dialectical absence/presence of a word or sentence. Typically, these representations are learnt by a neural network from text data in a supervised or an unsupervised manner.

Annex 3 – Topic Guide

Introduction

Thank you for agreeing to take part in this interview. We are interested in finding out more about your use of emerging technologies and will use this information to help develop case studies and best practice lessons for other regulators. This interview forms part of a wider study on this topic that ICF is carrying out on behalf of the Better Regulation Executive.

The interview is expected to last between 45 and 60 minutes, depending on your responses.

Your participation is voluntary and you do not have to answer any questions that you do not wish to. Your responses will not be attributed to you personally, but they will be reported on in relation to [emerging technology case] and attributed to the organisation you represent.

With your permission I will now start recording the interview. The recording will not be shared with BRE and is only used to ensure we have an accurate record of your responses to support the analysis.

Introduction

1. Could you briefly introduce yourself and your role in relation to [emerging technology case]?

Description of the case study

1. Could you start off by telling me a little bit about [emerging technology case]? Prompt for:
    a. What has it been designed to help with? How was it determined that this technology could be useful? What regulatory function was it serving?
    b. What technologies does it incorporate?
    c. What triggered the decision to develop the use of this tech by your organisation?
    d. When did development begin? Where in the process are you now?

Developing the technology

2. Were there any legal barriers or perceived legal uncertainty (e.g. lack of an appropriate framework) when developing [emerging technology case]? Prompt for:
    a. How was this addressed?
    b. Did these legal barriers of uncertainties change or limit what you set out to do/achieve? How?
3. What sorts of resources did you need to develop [emerging technology case]? (examples may include financial resources, human resources, infrastructure, computational resources) Prompt for:
    a. Did you encounter any challenges when trying to find these resources? How have you addressed these?
    b. Have you worked with external partners? Who were they/what was their role?
        i. If yes, how would you describe the process of working with them? Did you encounter any specific challenges? Were you able to address these?

   c. Did you encounter any challenges in finding the skills needed within your organisation? If so, how have you addressed this?
      i. Do you anticipate the need for any upskilling in the future?
   d. Did you need to collaborate with other teams within your organisation? Were there any challenges arising from this?
4. To what extent was UX/the user experience taken into account when developing [emerging technology case]? Prompt for:
   a. Was this something you were able to discuss with the contractor?
   b. Was this something you were able to carry out in house? If so, who took care of this?
   c. What have the outcomes of this been (e.g. has the tool been easy for staff to use and incorporate)?

Timeline

5. Could you describe in more detail how long it took to develop the use of the technology? Could you describe the length of time each stage of development took? E.g. securing internal clearances, implementation?
6. How long has the technology been in use?

Accessing/using data [for cases that make use of data]

7. Was there a need to access data from external sources (or internal sources) as part of the development of [emerging technology case]? (examples might be data from other government departments, non-government data such as social media etc) If so, prompt for:
   a. What types of data did you need to access? Where did these come from?
   b. What was this process like?
   c. Were there any challenges in accessing this data? How did you address these?
8. Did you encounter any challenges in making use of the data you had available? How did you address these? (examples might be issues with data quality, issues with data confidentiality)
9. Was there a need to collect any new data to support [emerging technology case]? If so, prompt for:
   a. What types of data were collected?
   b. Were there any challenges in collecting this data? How did you address these?

Maintaining the technology

10. Does [emerging technology case] require any on-going maintenance or monitoring? Prompt for:
    a. How have you addressed this?
    b. Are your own staff involved? What role does the contractor play in this?

Impact

11. What impact has the implementation of [emerging technology case] had so far? (or, if not yet implemented: What impact would you expect [emerging technology case] to have?) Prompt for:
    a. Does this differ from the impact you first anticipated? Why?
    b. What impact has it had compared to if the tech had not been introduced? (counterfactual if they are able to shed some light- even qualitative or anecdotal evidence would be useful)

   c. Are you looking to measure impact in any specific way (e.g. is any sort of impact evaluation planned)? Or if already implemented - Has your organisation done any assessment/measurement of the gains from deploying this tech? If so, what does it show?
   d. Do the gains relate to efficiency, or of new capability, or both? Please elaborate.

Links to other technologies

12. Has the implementation of [emerging technology case] built on any previous work you have completed?
13. Do you anticipate that future projects might build on what has been started by [emerging technology case]?

Other

14. Are there any other challenges you have encountered related to the development and implementation of [emerging technology case]?
15. Based on your experience, would you have gone about developing your use of this technology/ies differently?

Helpful resources and closing

16. What types of resources would you recommend for other regulators looking to pursue something similar? Prompt for:
    a. Are there any specific individuals it would be helpful to speak with? (if necessary, clarify also that we are looking for diverse backgrounds: policy as well as technical)
    b. Are there any publications or guidance documents you would recommend?
17. What issues would you say an organisation should carefully consider before using this technology?
18. Would you be happy to have your name passed onto BRE? This would not be included in any report, but we will include text in the report stating that BRE have relevant contacts that they could put a regulator in contact with if they would like to discuss the case further.

This publication is available from: www.gov.uk/beis