UNIVERSITEIT VAN TILBURG

Public Sector Data Ethics

From principles to practice

dr. L.E.M. Taylor
Prof. dr. R.E. Leenes
S. van Schendel LLM

Tilburg University
TILT – Tilburg Institute for Law, Technology, and Society
P.O. Box 90153
5000 LE Tilburg
The Netherlands

April 2017

Colophon

Authors
Linnet Taylor
Ronald Leenes
Sascha van Schendel

Publisher
Tilburg University
TILT – Tilburg Institute for Law, Technology, and Society
P.O. Box 90153
5000 LE Tilburg
The Netherlands

Commissioned by
BZK, Ministry of the Interior and Kingdom Relations
Turfmarkt 147
2511 DP The Hague
The Netherlands

Date
April 2017

Table of Contents

Summary

Public Sector Data Ethics: From principles to practice

The problem

The new data technologies are a formidable and important tool for governance. However, the move from digitising government functions to government relying on the sharing and use of digital data for its functioning brings some entirely new challenges in terms of ethics. Public-sector ethics has traditionally focused on the behaviour of public servants rather than the behaviour of public institutions and departments towards the people. The latter is most relevant to the challenges of big data. The **example of the private sector** shows that despite legal checks and balances, it is still common for data to be used in ways that are invasive of privacy, discriminatory against protected groups, or are otherwise socially detrimental. The public sector shares many of the same risks, including lack of transparency.

The Dutch government's use of digital data is extensive, but for most departments the use of advanced data analytic techniques is still at a relatively early stage of evolution. The signals are, however, that governmental data science will increasingly involve external capacity, i.e. **collaboration** between government, academic and private-sector partners. For example, in areas such as predictive policing, tax fraud and fraud detection with regard to publicly funded services and subsidies, where large-scale data analytics are now being used as standard, the government frequently collaborates with private-sector technology firms who provide the expertise and operational skills to run the necessary systems. Furthermore, in the near future, much of the 'big data' which offers opportunities for real innovation in governance is likely to come from **outside the government's auspices**, for example from social media or citizen initiatives. These partnerships and diverse sources will make it hard to map and identify the different types of data analytics being used, and equally difficult to identify the red flags[1] that mark actual or potential problems with regard to the ethical use of data.

We identify a **gap in institutional oversight** in the public sector with regard to data ethics: the process of reasoning necessary to apply legal rules in different situations, to judge risk and to evaluate how to balance interests and rights. The current vision of oversight based on fiscal efficiency and legal compliance with data protection regulations encourages a short-term and compliance-based perception of responsibility which is insufficient to the longer-term social and democratic implications of data analytics as a tool of governance. What is currently missing is an ethical framework that can bridge the two sets of rules to provide meaningful and practical guidance as internal and external partners work together with different understandings of what is 'inside the box' in terms of experimentation with data**.**

Models, principles and practices

We argue for a model of data ethics oversight in government that puts accountability at the centre, and uses the idea of responsible data use to build a system of checks and balances. Such a framework also

[1] Scientific Council for Government Policy (WRR) (2011). iGovernment: synthesis of WRR report 86. The Hague: WRR. Available at https://www.wrr.nl/binaries/wrr/documenten/rapporten/2011/03/15/ioverheid/ioverheid.pdf

includes measures to create transparency, and will to some extent be public-facing. We identify different models for institutional oversight of data-driven governance:

1. **External (ad hoc) oversight**, through whistleblowing or civil society watchdog institutions;
2. **Internal oversight on the judicial level**, as with intelligence agencies which may have a judicial council to advise on internal rules;
3. **Independent government oversight**, such as the UK's National Statistician's Data Ethics Advisory Committee;[2]
4. **Internal oversight on the municipal/departmental level**, such as Amsterdam municipality's dedicated data protection body (Commissie Persoonsgegevens Amsterdam);
5. **External audits of the government's use of data**, for which a new public agency would have to be created, possibly on the model of EUROSAI (the European Organisation of Supreme Audit Institutions), which does public-sector audits on all sectors of government.
6. **Making projects public-facing**, for example through opening data or charting processes on public websites.

The international landscape provides different models: we focus on those of the UK and France which are each addressing similar issues to the Netherlands with quite different approaches. The UK has established six general principles for data use at the governmental level: 1) Start with a clear user need and public benefit; 2) Use the minimum level of data necessary to fulfill the public benefit; 3) Build robust data science models; 4) Be alert to public perceptions; 5) Be as open and accountable as possible, and 6) Keep data safe and secure. The strength - and weakness - of this approach is that it operates at a **high level of generalisation** and is not specific to any particular use of data or type of practice.

In contrast the French government has passed a law (the 'Loi Numerique') with three main elements: first, it aims to encourage the use of data 'in the public interest' by government and citizens, including as a tool for economic growth; second, it establishes protection for platforms, and third, it sets out an agenda for creating universal access to digital technology. The French approach offers various protections designed to promote equality in the digital sphere. Its advantage is that it creates an **enforceable set of measures** that both promote datafication and set out checks and balances, but it does not set out guidelines on how to **identify or address the societal risks** of data science, and will therefore be shaped by legal decisions rather than providing a vision of how governance should be shaped by new data sources.

These different examples raise the question, what kind of oversight creates both **responsibility** – where people think *ex ante* about their role, their effects and the future of the dataset in question – and **accountability** – the structuring of *ex-post* responses to problems. At the moment data ethics is being addressed in an ad hoc fashion when problems arise. The question is what kind of oversight can address it systematically, in a way relevant to the challenges of big data.

---

[2] Information about the National Statistician's Data Ethics Advisory Committee available at https://www.statisticsauthority.gov.uk/national-statistician/national-statisticians-data-ethics-advisory-committee/

Proposal: a model for the ethical use of big data in government

We devise a model that consists of three elements:

1. **Background principles**

First, the core values we want to protect and promote in government data science must be defined. Here we must choose whether to create concrete *rules* or general *principles*, and whether we want the values concerned to be *generic* or *domain-specific*.

2. **Accountability framework: the control loop**

For the core accountability scheme, we refer to the A4Cloud project's model, which goes beyond mere compliance to demand ethical reflection and respect for the information that is being used in an analysis. In this model, an accountable organisation 1) defines what it does; 2) performs what it has defined; 3) monitors how it acts; 4) remedies any discrepancies between the definition of what should occur (norms) and what is actually occurring (behaviour), and 6) explains and justifies its actions.

3. **Enforcement and oversight**

After defining norms (for instance, data protection and domain specific regulation), and principles (as defined in layer 1), along with other important issues such as preventing function creep by reviewing activities after the application stage, we move to the kinds of structures that form the core accountability layer. In this layer, choices need to be made: regarding **oversight and enforcement**, should the scope of reporting be *extensive* (e.g. impact assessment) or *marginal* (description of goals and methods)? Should reporting be *internal* or *external*? Should an enforcement body be *internal* or *external*; *ad hoc* (e.g. a regulator for a specific project) or *structural* (for example a Public Data Science Authority that oversees public sector data science projects). The *'distance'* to the project organisation is also an important point of choice: should a regulator be responsible for overseeing a sector, a ministry, a layer of government, or other units? This choice of level and position has implications for the skills and knowledge required of the regulator.

Recommendation: a sectoral approach

For the Netherlands, a model is necessary that can offer efficient oversight for both ministries and decentralised elements of government (such as the country's 388 municipalities), and that can grow with the government's use of data science. Given these criteria, the most appropriate solution at present seems to lie in **establishing principles centrally** that are then applied through **sectoral frameworks**. Through sectoral oversight, the control loop described here can function on different levels, potentially in relation to existing oversight structures such as the Central Statistical Bureau (CBS), which is already a meeting-point for operational information on activities with digital data across municipalities.

One challenge is to prevent the burden of oversight being placed solely on one organ, which would make addressing the full range of data ethics considerations on a governmental scale impossible. Instead, using **multiple structural options** as appropriate at the sectoral level makes it possible to achieve both 'horizontal' accountability (awareness of principles, ensuring departments report internally and to immediate management) and oversight from above (via the legal framework, higher-level oversight bodies, external bodies and/or routes to civil society accountability). There is an advantage to also incorporating a **public-facing element** of oversight, whether through open reporting requirements or the involvement of civil society institutions, because such a public-facing component adds to the

legitimacy of government's data projects within society, and would contribute to public awareness of, and debate about, the use of data in governance. The ideal outcome of the structural approach posited here would be to make data ethics, similarly to privacy by design, a mindset and a continuous process rather than a task of compliance with norms. Although this mindset would take time to become embedded, the process of negotiating, discussing and defining ethical principles is essential if accountability is to be legitimate and enforceable.

Public Sector Data Ethics: from principles to practice

The problem

The new data technologies are a formidable and important tool for governance. The ability to access detailed, large-scale, often real-time data on social and governmental dynamics, combined with the potential to enrich conventional governmental data by complementing, merging or linking it with other data sources, will provide an unprecedented power to visualise, to act and to evaluate in the policy sphere. It also, however, raises new risks that must be understood and addressed. However, this move from *e-government* (digitising administrative processes) to *i-government* (a government that relies for its functioning on the sharing and use of digital data)[3] brings some entirely new challenges in terms of ethics. Public-sector ethics has traditionally tended to focus on the behaviour of public servants, corruption and conflicts of interest, rather than on the behaviour of public institutions and departments towards the people. Although transparency, combating corruption and ensuring accountability and efficiency in public spending remain as important ethical goals as ever, as government becomes increasingly datafied[4] it is becoming clear that new forms of oversight and accountability are needed with regard to the ubiquitous use of digital data.

The use of big data by the private sector indicates both what is possible, and what the risks are where data is used without ethical frameworks and appropriate oversight. Firms' use of data has been shaped by the law in the shape of data protection regulation, by internal oversight, and by customer feedback. Despite these checks and balances, it is still common for data to be used in ways that are invasive of privacy, discriminatory against minorities or protected groups, or are otherwise socially detrimental. The public sector shares many of the same risk factors: the way it processes the data it collects from citizens is generally invisible to those citizens; it has areas that are effectively off-limits to scrutiny because of security concerns; it demands that data be used to improve efficiency and value for money, and the size and complexity of government bureaucracy makes it difficult to audit data flows effectively. Beyond these commonalities, however, governments may be even more at risk of creating unfair treatment through data because of their particular responsibilities: public welfare, and public security, are increasingly seen as demanding the linking and merging of data across domains, without the explicit responsibility of transparency.

---

[3] Scientific Council for Government Policy (WRR) (2011). iGovernment: synthesis of WRR report 86. The Hague: WRR. Available at https://www.wrr.nl/binaries/wrr/documenten/rapporten/2011/03/15/ioverheid/ioverheid.pdf

[4] Mayer-Schönberger, V. and Cukier K. (2013). Big Data: A Revolution that Will Transform How We Live, Work, and Think. London: John Murray.

Figure 1. China's Social Credit system

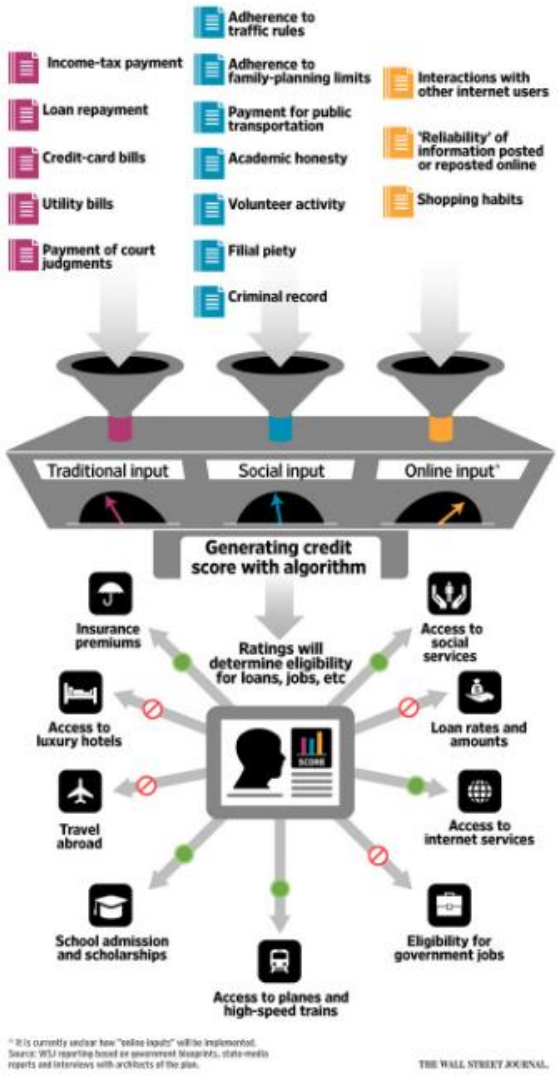

The power to collect and channel data from the public that is more detailed than any data used before by government, combined with access to emerging analytical technologies, has the potential to be problematic on many levels. One extreme instance that serves to illustrate this risk is the Chinese government's current project to create a social credit system (figure 1). In this system, flows of data about citizens from every area of life, public, private and economic, are merged into an algorithmic analysis that scores each person on their compliance with official state values, and correspondingly doles out freedoms or restrictions in almost every area of everyday life, from access to education and travel to the welfare system. The Chinese system demonstrates the power of today's data technologies to link and combine data from a huge range of sources, and to profile citizens in a highly dynamic and detailed way. This is one possible future for datafied governance; there are many others.

In the EU, we can see examples of data being used to change public behaviour in ways that are invisible to the public: one recent example is the UK's Behavioural Insights Team, formerly known as the Nudge Unit, which uses insights from behavioural economics and big data analytics to develop ways to 'nudge' people towards desired behaviours. These range from paying taxes on time to changing the way people give money to charity[5] and their engagement with preventive health. The Nudge Unit's work demonstrates the power of big data analytics to influence behaviour on the population level, but also the potential for misuse of that influence. The methods that can change people's health-seeking behaviour, beliefs about how to spend money, and everyday actions can also influence people in less desirable ways.

[5] The Behavioural Insights Team. Legacy giving and behavioural insights (2016, October 18). Retrieved from http://www.behaviouralinsights.co.uk/publications/legacy-giving-and-behavioural-insights/

In this report we will identify a gap in institutional oversight with regard to data ethics in the public sector. As datafication becomes increasingly embedded in government's functions, it becomes important to ensure that government establishes lasting criteria for good data governance, and sets a higher standard than the legal compliance that guides the private sector. Rather than data protection and anti-discrimination regulation, this report is concerned with data ethics: the process of reasoning necessary to apply legal rules in different situations, to judge risk and to evaluate how to balance interests and rights. Data ethics goes beyond the law, but can be encoded into specific rules or general principles, both of which possibilities will be discussed below. We also examine the aims implied by this encoding in different cases. Sometimes the aim is to provide sector-specific guidance to people dealing with a particular type of data (for example, the National Statistician's Data Ethics Advisory Committee in the UK). Alternatively, the aim may be to provide government-wide guidance that covers all uses of data (such as the UK Cabinet Office's code of data ethics).

We will outline the oversight gap and its implications for the Dutch government, using both current uses of data by the Dutch government and also drawing on international examples to show what is possible in terms of public-sector data use. We will then analyse different models of data ethics oversight to determine which elements or goals are appropriate for the Dutch case, and make recommendations as to how such oversight should be organised. Our aim is to explore and evaluate what kind of framework for data ethics adds up to a broader level of accountability where the government sets the tone and parameters for its use of data overall. Governments in many countries are experiencing a shift from data-informed to data-driven governance,[6] a shift that raises the question not only of what ethical oversight for data in government should do, but what its overall aim is. For data-informed governance, it may be sufficient to check that the data is accurate, compliant with the law, and securely stored. For data-driven governance, the ethical questions become of a larger scale. What values is the use of digital data supporting and promoting? How does it fit within a democratic context? The question moves from whether the numbers are correct to whether their use promotes values of good government and thus whether or not certain data science practices should be pursued.

Governmental use of data analytics & the Dutch context

The Dutch government's use of digital data is currently extensive, but for most government departments the use of advanced data analytic techniques is still at a relatively early stage of evolution. Although most of the current applications that have been classified as Big Data within government would probably not meet the common definitions in literature, the use of data in innovative ways is clearly on the agenda in the Dutch public sector. Especially on the local level, municipalities are either experimenting with Big Data or are gathering and channeling data that could be analysed with data scientific. The Institute for Municipal Quality (Kwaliteits Instituut Nederlandse Gemeenten, or KING) has started a pilot together with several other actors, including a private actor specialised in data analytics, a research institution and a municipality, to analyse the possibilities of data collection and analytics in the

[6] Kitchin R. (2016). The ethics of smart cities and urban science. *Phil. Trans. R. Soc.* A 374: 20160115. http://dx.doi.org/10.1098/rsta.2016.0115

context of the 'omgevingswet' (environmental regulations).[7] Meanwhile, Amsterdam has established an office of Research, Information and Statistics that conducts big data analytics on city data.[8] Many of the current innovative uses of data focus on improving the living environment in the city, whether in terms of improving traffic flows and accessibility of areas and events,[9] helping protect the environment,[10] improving the level of safety and/or reducing criminal activity,[11] improving living conditions for elderly[12] or improving government service provision.[13] Several of these elements come together in city pulse projects, such as in Amsterdam, which aims to prevent shortages of energy, food, water and materials while simultaneously reducing CO2 output.[14]

Besides the development of data applications and data driven projects by decentralised government, collaboration with private parties on data collection and analysis is increasingly being explored by Dutch government institutions. For example, in areas such as predictive policing, tax fraud and fraud detection with regard to publicly funded services and subsidies, where large-scale data analytics are now being used as standard, the government frequently collaborates with technology firms who provide the expertise and operational skills to run the necessary systems. One recent example of this is the FinPro analysis undertaken by the Netherlands' Public Prosecutor (Openbaar Ministerie) in the period leading

---

[7] Kwaliteits Instituut Nederlandse Gemeenten. Radicaal anders werken door data? Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-9

[8] Onderzoek, Informatie en Statistiek: http://www.ois.amsterdam.nl/overois

[9] Kwaliteits Instituut Nederlandse Gemeenten. Sensor City zorgt voor slimmere wegen in Assen. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-2; Kwaliteits Instituut Nederlandse Gemeenten. Onbezorgd naar het strand van Scheveningen. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-4-copy; AMS. De mobiele stad: drukte in kaart (2015, September 6). Retrieved from http://www.ams-institute.org/events/event/de-mobiele-stad-drukte-in-kaart/

[10] Kwaliteits Instituut Nederlandse Gemeenten. Samenwerken voor schonere lucht in Nijmegen. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-6

[11] Kwaliteits Instituut Nederlandse Gemeenten. Tilburg ontwikkelt dashboard tegen diefstal Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-3; Atos. Atos zet Big Data in voor veiliger nachtleven Eindhoven (2015, July 21).Retrieved from https://atos.net/nl/2015/persberichten_2015_07_21/nl-pr-2015_07_21_01; Kwaliteits Instituut Nederlandse Gemeenten. Veiliger uitgaan met City Pulse in Eindhoven. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-4; Kwaliteits Instituut Nederlandse Gemeenten. Zaanstad zet big data in tegen huiselijk geweld. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-8; Kwaliteits Instituut Nederlandse Gemeenten. Straatkubus spoort leefbaarheidsproblemen vroegtijdig op in Almere. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-7; Kwaliteits Instituut Nederlandse Gemeenten. City Alerts verbetert veiligheid hulpverleners in Amsterdam. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-5

[12] Kwaliteits Instituut Nederlandse Gemeenten. Veldacademie geeft big data betekenis voor Rotterdamse buurten. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy

[13] Kwaliteits Instituut Nederlandse Gemeenten. Data voor digitale dienstverlening in Nederlandse gemeenten. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/1-amsterdam-copy-10

[14] Kwaliteits Instituut Nederlandse Gemeenten. Data als kloppend hart voor een duurzaam Amsterdam. Retrieved from http://magazine.kinggemeenten.nl/data-en-gemeenten/#!/gemeente-amsterdam-ams-institute

up to 2016,[15] where data from various public-sector departments and private-sector partners was brought together to research patterns that could identify fraud in welfare and subsidy claims. The research on this citizen data was conducted by a third party, Nyenrode University, so that data passed outside government systems for analysis, then conclusions were passed back in different form. A debate ensued about whether the program had gone far enough given the evident power of this merging of datasets to identify fraud,[16] but did not significantly address the questions of ethics that the process of shared data management and research brought up.

The mingling of public and private-sector data, and the potential for analysis by third parties in academia or quasi-governmental organisations, are increasingly features of governmental data science, and are likely only to grow with time. The box below shows an overview of the ways in which public and private data sources may mingle in processes of data-driven policymaking.

The potential for mixing public- and private-sector data sources (Kitchin 2016)[17]

- utility companies (use of electricity, gas and water);
- transport providers (location/movement, travel flow);
- mobile phone operators (location/movement, app use and behaviour);
- travel and accommodation websites (reviews, location/movement and consumption);
- social media sites (opinions, photos, personal information and location/movement);
- crowdsourcing and citizen science (maps, e.g. OpenStreetMap; local knowledge, e.g. Wikipedia; weather reports);
- government bodies and public administration (services, performance and surveys);
- financial institutions and retail chains (consumption and location);
- private surveillance and security firms (location and behaviour);
- emergency services (security, crime, policing and response);
- home appliances and entertainment systems (behaviour and consumption).

These examples suggest that while government's use of digital data is already highly diverse, it is likely to become more opportunistic and decentralised in the future, so that there is no easy way of knowing what is being done, by which departments, or how projects are evolving. The further decentralisation of many government data processing functions to the municipal level makes it even harder to map and

[15] Big data legt onzichtbare criminaliteit bloot. Het Financieele Dagblad (2016, May 9). Retrieved from https://fd.nl/economie-politiek/1150701/big-data-legt-onzichtbare-criminaliteit-bloot

[16] Ministry of Justice. Antwoorden Kamervragen Big data experiment legt onzichtbare criminaliteit bloot (2016, September 1). Retrieved from https://www.rijksoverheid.nl/documenten/kamerstukken/2016/09/01/antwoorden-kamervragen-big-data-experiment-legt-onzichtbare-criminaliteit-bloot

[17] Kitchin R. (2016). The ethics of smart cities and urban science. Phil. Trans. R. Soc. A 374: 20160115. http://dx.doi.org/10.1098/rsta.2016.0115

identify the different types of data analytics being used, and equally difficult to flag actual or potential problems with regard to the ethical use of data.

The WRR's 2011 *iGovernment* report[18] offers some 'red flags' that identify risky governmental information processes. The first is the networking of information – sharing use and management of data among a network of actors. The next, compiling and enhancing of information, occurs when new information and profiles are based on different sources from different contexts. The last, preventive and pro-active policy based on information, occurs when interventions are based on calculations of risk. Since 2011, uses of data that raise these red flags have multiplied in the Dutch governmental context, particularly around issues such as preventive policing, fraud detection and security. However, almost all the data government deals with, *including seemingly neutral forms of data such as information from environmental sensors or transport networks*, either stem from, describe or can be used to influence people's behaviour, activities and communications. This broadens the possible occurrences of the red flags identified by the WRR from these core security applications to almost every area of government policy where data may be merged, enriched and used to produce predictive guidance. It also argues for the development of an ethics oversight process that can identify data-related risks across a broad range of government functions, and that can grow with the government's use of digital data and analytic technologies.

Philosophers Floridi and Taddeo[19] define data ethics as having three components: first, *the ethics of data* (how data is generated, recorded and shared); second, *the ethics of algorithms* (how artificial intelligence, machine learning and robots interpret data), and third, *the ethics of practices* (devising responsible innovation and professional codes to guide this emerging science). So far, governmental data ethics have formally been defined by data protection instruments, which deals with the first point on Floridi and Taddeo's list – the data protection and fair processing aspect of data ethics. The key instrument has been the European Data Protection Directive (implemented in the Wet Bescherming Persoonsgegevens), which governs all flows of identifiable data with the exception of law enforcement-related data processing. This framework regulates the collection and processing of personal data founded on privacy and other fundamental rights in the EU. From 25 May 2018 onwards, the new General Data Protection Regulation (GDPR) (in Dutch 'Algemene verordening gegevensbescherming') will establish the new data protection framework covering Big Data practices outside of the law enforcement context. Data processing in the context of law enforcement[20] will be regulated in the Police Directive, which will need to be implemented into national law by 6 May 2018. The scope of application of the Police Directive is strictly defined, a lot of data processing by governmental actors will fall within the domain of the GDPR. The GDPR brings some significant changes compared to the Data Protection

[18] Scientific Council for Government Policy (WRR) (2011). iGovernment: synthesis of WRR report 86. The Hague: WRR. Available at https://www.wrr.nl/binaries/wrr/documenten/rapporten/2011/03/15/ioverheid/ioverheid.pdf

[19] Floridi, L., Taddeo, M. (2016) What is data ethics? Phil. Trans. R. Soc. A 374 (2083), 20160360. DOI: 10.1098/rsta.2016.0360.

[20] The exact scope is data processing by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, including the safeguarding against and the prevention of threats to public security.

Directive. One of the new elements that is interesting from a data ethics point of view is the approach of privacy by design that is laid down in some form as a legal requirement in the article of data protection by design and by default. This article requires actors who process personal data to take technical and organizational measures to implement data protection principles into the design of systems and to make sure that by default not more personal data is processed than necessary.

Although requiring data protection by design and by default is not that big a change, the GDPR gives legal import to the philosophy of privacy by design. Privacy by design focuses on more than mere legal compliance by integrating considerations of privacy throughout the process of system design, organisational structure and staff training. The new regulation will also require data protection impact assessments to describe the processing, risks, necessity and proportionality of any high-risk data processing, and the measures to be taken to mitigate that risk. The GDPR also requires the installation of data protection officers (DPO's), for example in the case of a public authority. In the case of a public authority a DPO may be assigned for several bodies or authorities together. The DPO will monitor compliance with data protection requirements within an entity, advise the entity and its employees on data protection matters and communicate with the national Data Protection Authority. Under both the GDPR and its predecessor, governments are held to the same standard in terms of data protection as private actors. However, it can be argued that government should be held to a higher standard with regard to processing data because it has a high and legally enforced level of access to very private data including education, health and crime records, and that this should go beyond law and regulation into the territory of ethics. Governmental ombudsmen could theoretically be part of addressing this gap, but have not played an identifiable role so far in doing so.

The second and third components of ethical use of data, however – concerning the use of algorithms and establishing principles for responsible innovation – pose new challenges to Dutch governmental oversight. First, because ethical codes for governance were largely developed in the era before big data and do not address the positive or negative ways in which large-scale data, and particularly the use of algorithms, may influence the operation of government. Second, because innovation related to data analysis is not restricted to one sector or operation of government, but is becoming something almost every department is expected to engage with, often on an exploratory level where identifying risks is not prioritised.[21] The provision of ethical guidelines has not kept up with these developments, nor have new institutional structures been established which can keep track of new practices, collaborations and experimentation with citizens' data. Instead the tendency is to rely on the general guidelines provided by data protection law or overarching ethical codes, without creating institutions that can provide day-to-day oversight of projects or audit the way data is being used, particularly at the public-private interface.[22]

---

[21] The UK's Nudge Unit is an example of this experimental dynamic, information about this Unit is available at http://www.behaviouralinsights.co.uk/

[22] One example of this is the lack of an audit trail with regard to the data being shared by the UK National Health Service with Google: one member of a patient rights organisation has said that [an audit procedure] 'will eventually demonstrate to a patient that data was copied to Google "for direct care" when they were nowhere

Although the ethical principles for the Dutch public sector mainly address efficiency and the use of public funds, 'principles of appropriate governance', *de beginselen van behoorlijk bestuur*, have been developed in administrative law and are applicable to interactions between an administrative organ and a citizen or organisation.[23] Many of the principles developed there are strongly focused on administrative procedures, however some can also be applied to a broader range of government activities; The prohibition of *détournement de pouvoir* together with the principle of specialisation prevent governmental actors from using a competence for another purpose than the one for which it was developed, and oblige actors to guard the specific interest that is the focus of given legislation; the principle of proportionality requires that adverse effects for citizens do not weigh heavier than the importance of the general interest that is served; the prohibition of *détournement de procédure* prohibits governmental actors from following a procedure with fewer safeguards to come to a decision, if there is also a procedure open which offers more safeguards; and finally the principles of legal certainty and motivation require governmental actors to make well motivated or reasoned decisions, to give citizens insight into why a decision was made and to create predictability in government decision-making.

While these principles form a good starting point for ethical governance, there have also been attempts to update them for the digital age. One such is the set of principles developed by Franken, *beginselen van behoorlijk IT-gebruik*.[24] Franken argues that ICTs create power and therefore put strain on traditional principles of democracy, and that the siting of ICT-related competency does not align with the power that the government has in terms of information.[25] Therefore Franken proposes six main principles to protect norms and values that might otherwise be put at risk by this structural change in the relation between government and citizens. These are: availability (of information and knowledge to make sure citizens are well informed and able to form their opinions); confidentiality (making systems secure and only open to authorised employees); integrity (of both IT systems themselves and of the information they contain); authenticity (the ability to check whether information really originates as indicated); flexibility (being able to adapt the system to new demands from users) and transparency (of the working of the system and of the origin of the decisions it makes).[26] Franken's principles predate Big Data, but nevertheless agree in many respects with frameworks evolving to guide today's data

---

near the hospital at the time. It should irrevocably log that Google got data that they were not entitled to access, and now refuse to answer questions about.' Retrieved from https://techcrunch.com/2017/03/09/deepmind-says-no-quick-fix-for-verifying-health-data-access/?imm_mid=0eeb7e&cmp=em-data-na-na-newsltr_20170315

[23] Some of these principles have been codified in the Algemene Wet Bestuursrecht.

[24] Franken, H. (1993) 'Kanttekeningen bij het automatiseren van beschikkingen' in Beschikken en Automatiseren. Preadvies voor de Vereniging voor Bestuursrecht, var-reeks nr. 110, Den Haag.

[25] Idem, p. 18.

[26] Idem, p. 18 - 22.

engineering, for example the Fairness, Accountability and Transparency in Machine Learning (FAT/ML) principles that are becoming influential as a tool for the ethical governance of algorithms.[27]

The Netherlands has various bodies dedicated to enforcing public sector principles, including Chief Information Officers, and the Bureau ICT Toetsing (BIT) within BZK, which oversees large projects. None of these, however, oversee the collection and use of data. The Algemene Rekenkamer has released several statements and reports[28] that are critical of ICT policies, but its critique focuses more on possibilities, feasibility, and on what the government should do in terms of a practical strategy. It is less concerned with the values that are reflected or should be reflected in governmental policy or accountability of policies.

Oversight on data protection is provided by the Dutch Data Protection Authority (Autoriteit Persoonsgegevens) which oversees compliance with data protection legislation such as the Wet Bescherming Persoonsgegevens (Dutch Data Protection Act) and data protection acts in the domain of policing, and which advises the government on matters of data protection legislation. Its focus is on legality, rather than ethics. There are also lower-level processes in place: examples include the requirement that departmental-level officials consider requests from outside researchers to use government-collected data, and data protection bodies such as the Commissie Persoonsgegevens Amsterdam, which regulates the city's use of personal data in planning and governance. Both of these, however, are focused strongly on the legality of data collection and use, and need to be complemented by other bodies focusing on data ethics in order to promote a vision that goes beyond legal compliance in data processing.

The current vision of oversight based on fiscal efficiency and legal compliance with data protection regulations leaves a gap with regard to the ethical dimension of data use, namely because it encourages a short-term and compliance-based perception of responsibility which is insufficient to the longer-term social and democratic implications of data analytics as a tool of governance. What is currently missing is an ethical framework with regard to the public sector's use of data science, something that can bridge the two sets of rules**.** Ethical oversight for data use goes beyond either data protection rules or conventional public sector ethics because it encompasses both systems and practices. Such oversight requires consultation and consideration at the stage of planning and procurement, and second, ongoing application review to prevent function creep (the repurposing of data or systems, for example when a camera system designed to monitor traffic flows becomes useful for monitoring the movements of suspected criminals via their license plate numbers). One example of this two-step approach comes from the EU's 2015 report on oversight of data use by military intelligence,[29] which provides a dual focus

[27] FAT/ML. Principles for Accountable Algorithms and a Social Impact Statement for Algorithms. Retrieved from http://www.fatml.org/resources/principles-for-accountable-algorithms

[28] For example: Algemene Rekenkamer. Trendrapport open data (2015). Retrieved from http://www.rekenkamer.nl/Publicaties/Onderzoeksrapporten/Introducties/2015/03/Trendrapport_open_data_2015

[29] European Commission for Democracy through Law (Venice Commission). Report on the democratic oversight of signals intelligence agencies (2015). Retrieved from http://www.venice.coe.int/webforms/documents/default.aspx?pdffile=CDL-AD(2015)011-e

on rules and their translation into practices. The report notes that 'institutionalising respect for human rights comes in both when the internal rules are being devised, and, because of the importance of automated minimisation systems, at the point these rules are "translated" into software.'

Similarly, a dual focus is necessary on insiders and outsiders in the public sector. Government departments challenged to innovate with data largely know what risks their own data sources entail. In contrast, it is difficult for outsiders to understand the mission of the agency. When the two collaborate, it becomes easy for the outsiders to stray outside the box in terms of the kinds of experimentation they do with data (as could be seen from the Snowden revelations of consultants to the US government being comfortable using data gathered under the rubric of national security to view intimate photos sent by citizens via Yahoo mail). Outsiders will tend to have an interest in data rather than the domain, and therefore have an interest in discovery rather than user needs and public benefits.

Models, principles and practices

There are two possible models for data ethics oversight in government. One puts responsible data science and innovation at the centre, and works from there to a set of standards that can be used by any governmental office using digital data. The other puts accountability at the centre, and incorporates various dimensions of responsibility in a framework that also includes checks and balances. Such an accountability-based framework must include establishing responsible data science principles, but it also sets out a series of clear steps: first, define the rules you want to abide by; second, monitor what you do; third, correct deviations, and finally be prepared to take responsibility for the whole circle.[30] Such a framework inevitably also includes measures to create transparency, and will to some extent be public-facing.

Here, the specifics of public-private sector collaboration complicate matters; while workers in a particular sector will understand the kind of data they are dealing with (and the (ethical/social/legal) concerns these data may raise), external collaborators may not. Now that governments are largely collaborating with the private sector for any large-scale data analytics work, there will inevitably be different understandings of how to keep data in context and ensure experimentation does not overstep the bounds of the original purpose of data collection, both important principles in data ethics.[31] For example, municipal employees usually use data from the monitoring of public space responsibly, but the ability to merge, link and add new data sources may tempt external consultants to experiment in ways that provide predictive power, extra analytical scope or simply to see what can be done.[32] One way to address this is to have all those involved in a private-public collaboration around data agree and write

[30] Jaatun, M. G., Pearson, S. Gittler, F., Leenes, R.E., Niezen, M.. Enhancing accountability in the cloud. International Journal of Information Management (2016), http://dx.doi.org/10.1016/j.ijinfomgt.2016.03.; PiLab (2012), iOverheid, burger in beeld (final report). Accessible at: https://pilab.nl/onewebmedia/PI.lab-2012-R11216-Eindrapport-iOverheid-burger-in-beeld1.pdf

[31] Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

[32] 'Een biertje met Big Brother erbij op Stratumseind'. Volkskrant (2015, May 9). Retrieved from http://www.volkskrant.nl/binnenland/een-biertje-met-big-brother-erbij-op-stratumseind~a4192665/

down the aims of the project and envisioned 'issues' relating to data analytics on the data at hand at the start, which will have the effect of sharing standards with the private-sector collaborators about what is ethical to do with data in a particular domain, but also of exposing public-sector workers to the possibilities of the data and challenging them to define what kinds of new exploration would be useful.

The private-public interface also functionally poses two problems for the establishment of any independent oversight: visibility and scale. First, it becomes difficult to know what is happening outside of government offices in private firms conducting analytics on behalf of government. Second, collaboration increases the scale at which government can engage in data science exploration, but also makes it harder oversee the use of data across institutions and across the public-private divide.

Other principles for data ethics oversight in government would include proactivity - that oversight should be empowered and proactive; specificity to the use of digital data, rather than to ethical behaviour in general; that activities within government should create a 'data trail' to make it possible to trace who has handled data and made decisions about it.

Institutional models for ethical oversight:

On the international scale, we identify several models for institutional oversight of data-driven governance. These operate at different levels and with different degrees of independence, and together offer a picture of the prevailing collective wisdom. An overview of these follows. These represent a variety of dimensions that add up to oversight, from which we can generate the models we currently have. Each has advantages and disadvantages.

1. External (ad hoc) oversight

   This can be conducted through whistleblowing (for example, the Netherlands' new law on whistleblowers), but in the case of problems with the use of data by security or intelligence units of government, exposes the whistleblower to significant political pushback (for example, Edward Snowden in the US or David Kelly in the UK). An alternative is the activities of civil society institutions, which can use legal remedies against government misuse of data. The ACLU and EFF in the US, Privacy International in the UK, Bits of Freedom in the Netherlands and the organisation Droits des Lycéens in France have all had success with legal challenges to government practices with data.

2. Internal oversight on the judicial level

   This is usually weaker the more highly funded and powerful the agency[33], but has worked in some countries where, for example, intelligence oversight bodies are empowered in ways that take into consideration the closed loop of security services oversight: for example the Swedish Signals

[33] Demmke, C. and Moilanen, T. (2012). Effectiveness of Ethics and Good Governance in Central Administration of EU-27: Evaluating Reform Outcomes in the Context of the Financial Crisis. Available at http://www.integriteitoverheid.nl/fileadmin/BIOS/data/Publicaties/Downloads/Effectiveness_of_Public-service_Ethics_and_Good_Governance_in_the_Central_Administrations_of_the_EU-27.pdf

Intelligence Agency has an "integrity council" consisting of three judges who have an advisory role when the agency is devising internal rules.

3. Independent government oversight

This is the new mode of establishing and enforcing data ethics principles, but is in its infancy. Examples include the National Statistician's Data Ethics Advisory Committee[34] in the UK. A council of data ethics has been proposed and agreed to by the UK government, but has not yet been created.

4. Internal oversight on the municipal/departmental level

Examples exist of ethics codes tailored to specific governmental tasks and functions. For example, Michigan's Dept. of Transport has an ethics guide,[35] that focuses on identifiable personal data, but also includes broader ethical positions such as the importance of preserving the possibility of anonymity in public space. Amsterdam has its own data protection body (Commissie Persoonsgegevens Amsterdam) that focuses on the city's use of data.

5. External audits of the government's use of data

The Rathenau Institute has recommended[36] establishing a public agency to audit the code used in governance and determine whether it is fair. Such an external audit function could also be conducted by an international body, for example by adapting the existing external governmental audit function of a body such as EUROSAI (the European Organisation of Supreme Audit Institutions), which does public-sector audits on all sectors of government. Its work has included audits of ethics and integrity, but the institution has not yet conducted auditing on data handling. EUROSAI is part of INTOSAI, which works under the UN's ECOSOC body.

6. Making projects public-facing

This model outsources oversight more generally to civil society by making data science projects public-facing. One current example is the French case where the Min of Education has been forced by public pressure to make its college placement algorithm (partly) public. The city government of Chicago is doing this purposefully, making the data from 'Array of Things' sensor project[37] open as a strategy to substitute for ethical oversight.

---

[34] Information about the National Statistician's Data Ethics Advisory Committee available at https://www.statisticsauthority.gov.uk/national-statistician/national-statisticians-data-ethics-advisory-committee/

[35] Michigan Department of Transportation & Center for Automotive Research (2014). ITS Data Ethics in the Public Sector. Available at https://www.michigan.gov/documents/mdot/06-14-2014_ITS_Data_Ethics_in_the_Public_Sector_464226_7.pdf

[36] Kool, L, J. Timmer & R. van Est (2015). De datagedreven samenleving - Achtergrondstudie. Den Haag: Rathenau Instituut. Retrieved from https://www.rathenau.nl/nl/files/rapportdatagedrevensamenlevingrathenau-instituutpdf

[37] Urban Center for Computation and Data. Chicago becomes first city to launch array of things (2016, August 29). Retrieved from http://www.urbanccd.org/news/2016/8/29/chicago-becomes-first-city-to-launch-array-of-things

What does the international landscape tell us about government data ethics?

Looking at the international scene we see different models, each of which raises questions and points out possibilities in how to define and enforce ethical behaviour with data. Governmental efforts to understand how to best use the new data sources and technologies are evolving along different paths. We will compare here different approaches from the EU and US, with a special focus on those of the UK and France which are each addressing similar issues to the Netherlands with quite different approaches.

Establishing general ethical guidelines for data use at the governmental level has the benefit of setting out clear boundaries for all government departments. The UK has done so, in order to provide guidance as it attempts to stimulate innovation in the use of data in governance. These include internal learning, supplementing national statistics and behaviour change.[38] The ethics code (see box below) was established by the Cabinet Office in 2016 (a department with an oversight function for the whole of government, and access to all government functions).

UK Cabinet 2016 data ethics guidelines

The UK Cabinet has established a set of guidelines accompanied by a list of questions for government departments to answer, creating a checklist. This evaluation of the risks of data use does not focus on one specific stage of the project, the guidelines focus on several stages of government programs. The proposed principles are:

1. Start with a clear user need and public benefit
2. Use the minimum level of data necessary to fulfill the public benefit
3. Build robust data science models
4. Be alert to public perceptions
5. Be as open and accountable as possible
6. Keep data safe and secure

The riskier the project in terms of privacy and data protection, the more caution is required from these actors. Projects that might pose more severe risks can be subjected to a further assessment. For these projects the guidelines offer more detailed questions which can function simultaneously as a Privacy Impact Assessment.

The strength – and weakness – of the UK's approach is that it operates at a high level of generalisation and is not specific to any particular use of data. There is also a National Statistician's Data Ethics Advisory Committee dedicated to ensuring ethical use of data in government statistics, which has its own sector-specific focus. This combination of providing general principles and subjecting certain offices to specific scrutinising has the advantage of focusing more attention on priority areas, but also has the

[38] For more information: https://gdsdata.blog.gov.uk/category/data-science/

potential to lead to fragmentation where some departments have explicit and enforceable rules and others do not.

Moreover, the UK's general principles approach uses 'data' and 'data science' as general terminology for what is actually a wide range of practices with very different types and levels of public risk attached. This has the benefit of accessibility - everyone from a local town council member to a member of the intelligence services can, in theory, apply its principles to their work. However, its status as a code rather than a rule or a directive means that it serves only as a guide for self-regulation, and there is no oversight or enforceable penalty for contravening it. This approach is useful for low-risk uses of data, but arguably less so for the direct use of data about citizens to conduct behavioural experiments on them. For instance, using data to visualise levels of poverty across the London region is not the same as the practice of gaining 'behavioural insights' from data which can then be used to alter behaviour on a mass scale (also known as 'nudging'). The UK's Nudge Unit[39] was formerly part of the Cabinet Office and gained a lot of publicity, both positive and negative, which may have been one factor leading to the office's publication of data ethics guidelines. In 2016, however, the unit was privatised, which will insulate the office from oversight by any future public-sector ethics body set up to regulate the use of big data to influence behaviour. It also places the team under the rubric of innovation (through a quasi-autonomous innovation institute) rather than policy making, framing it as experimental and exploratory rather than as a policy body (which it effectively is). Thus it has access to government data but no government oversight. These problems may be resolved if the promised 'council of data ethics' is set up in the UK, as agreed in 2016, but so far the government has not followed through on its promise to do so.[40]

The French government has taken an alternative approach by passing a law that packages together the government's overall aims with regard to the information society. The 'Loi Numérique',[41] passed in 2016, is aimed to regulate both the private and the public sectors. It raises sanctions for commercial misuse of data to €3m, provides some checks on government misuse of citizens' data such as 'white hat' hacker protection in public sector cases including national security, and sets open data goals. The law acknowledges both positive and negative rights on the part of the population with regard to digital data, and deals with the functions of data across the public sphere, rather than just the public sector as in the UK. The law has three main elements: first, it aims to encourage the use of data 'in the public interest' by government and citizens, including as a tool for economic growth; second, it establishes protection for platforms (through net neutrality and 'digital service platform fairness' and 'digital rights' for individuals. Third, it sets out an agenda for creating universal access to digital technology, including the disabled and disadvantaged. The French approach constitutes a broad ethical perspective on data that

[39] For more information: https://www.gov.uk/government/organisations/behavioural-insights-team

[40] Government agree to set up 'Council of Data Ethics' (2016, April 26). Retrieved from https://www.parliament.uk/business/committees/committees-a-z/commons-select/science-and-technology-committee/news-parliament-2015/big-data-dilemma-government-response-15-16/

[41] The explanatory memorandum for this legislation is available at http://www.republique-numerique.fr/pages/digital-republic-bill-rationale

encompasses virtue ethics (which ask how data should be used for good) and various protections designed to promote equality in the digital sphere.

This broad legal approach has had the advantage of entailing a long and detailed public debate about what a 'loi numerique' should entail, what kinds of rights and freedoms people should have with regard to digital data, and how digital technologies should shape the future. The data protection elements of the law broadly echo the forthcoming GDPR, including its provision that people should not be subjected to algorithmic decision making. This provision has already been used by the high-school students' association, which in October 2016 succeeded in forcing the Ministry of Education to release some of the source code from its college allocation algorithm (ABP).[42] The release was analysed *pro bono* by hackers, who found it seemed to be prioritising students from French territories over those from within France. However, since the whole algorithm was not released, and it is not clear whether the ministry must do so, a definitive answer is still pending.

These two approaches – the British and the French – aim in very different directions. The British model prioritises a broad statement about recognising and considering the broader societal implications of a data science project, but does not make the code enforceable. The French approach creates an enforceable set of measures that both promote datafication and set out checks and balances. The disadvantage of the British model is its fundamental vagueness - the definition of ethical behaviour is ultimately left up to the individual conducting the project, and there are no explicit sanctions for disregarding the code. The French model, however, does not aim to specify what ethical behaviour with data is: it includes a section on protecting personal data, but does not set out guidelines on how to identify or address the societal risks of data science[S7] [SvS8] . Instead, the Loi Numerique will be tested and shaped by juridical decisions, and the body of legal precedent that emerges will map where the thin ice exists rather than provide a vision of how governmental policymaking should be shaped by new data sources.

A third possible model that can suggest how data science in governance may develop can be drawn from the US government research unit DARPA. The unit has transitioned over the last decade from a model where it seeded projects directly to one where it aims to 're-architect social networks among researchers so as to influence new technology directions'. Although DARPA is only one unit within government, this model has importance for predicting the way in which governments may increasingly engage with big data. The reason is capacity: the Netherlands, like most governments, still has relatively low specialised capacity in terms of data science and is likely, as many municipalities are already, to become more reliant on partnerships with firms and academia to conduct projects as time goes on. Outsourcing and network-shaping thus inevitably become ways through which the government influences the development of data technologies, but also influences the way data is handled and the purposes to which it is put. This has implications for ethical oversight because it suggests that at least

[42] APB : ce qu'une première analyse du code source nous révèle. Le Monde (2016, October 19). Retrieved from http://ingenuingenieur.blog.lemonde.fr/2016/10/19/apb-ce-quune-premiere-analyse-du-code-source-nous-revele/

one component should be external to government in order to look at how government is exercising network governance rather than developing data applications itself.

These different examples raise the question, what kind of oversight creates both *responsibility* – where people think ex ante about their role, their effects and the future of the dataset in question – and *accountability* – the structuring of ex-post responses to problems. At the moment data ethics is being addressed in an ad hoc fashion when problems arise, but not systematically. This process is likely to be subject to demands from other areas of oversight such as efficiency, economy and internal transparency, but in this case they will not be adapted to the particular concerns relevant to big data and important issues may be overlooked.

Proposal: a model for the ethical use of big data in government

Providing guidance on what specific actors must do in order to conduct data science ethically must incorporate legal obligations and parameters. However, it also extends beyond the law: not everything that is permissible should be pursued. The models and problems identified above suggest that to be effective, data ethics oversight for government must be multi-level and composed of overlapping checks and balances. Ethics is rooted in principles that can be stated and can serve as a backdrop to such a sector-specific system of checks and balances. But beyond that, what can be deemed ethical needs to be determined within the context of each particular application, and within a given function of government, and therefore there is a need to create guidelines in a context-specific process. In this context-specific process, it is up to the stakeholders concerned to define what they deem appropriate and act accordingly, and therefore some of the challenge rests in convening the kind of process that can identify the relevant stakeholders and take their views into account. In order to make sure that practice as defined by the stakeholders stays within the boundaries defined, both oversight and enforcement are then required.

On this basis we can devise a model that consists of three elements:

1. Background principles
2. Accountability framework/control loop
3. Enforcement/oversight

We will elaborate this model below and outline choices that need to be made within the three layers.

Dimensions of the core model

1. Background principles

The background consists of a set of principles that serve as a baseline for ethical practice. The table below provides an example of core values grouped into three perspectives: individual, relationship between two individuals, society/social system. Values such as these may help outline the core values at play in a particular data science application area.

Table 1. Core values at stake in DS projects source: KNAW 2016[43]

Primarily applicable to a single individual	Primarily applicable to the relationship between two individuals	Primarily applicable to a social system
health	responsibility	respect
wellbeing	accountability	dignity
physical integrity	justice	non-discrimination
happiness	equity	transparency
privacy	solidarity	trust
security	autonomy	democracy
safety	confidentiality	freedom
knowledge	access	utility

In the context of public sector data science, a more concrete and tailored list of core values needs to be developed. On the basis of the list of values more specific guidance can be formulated. With respect to this decisions will need to be taken on two dimensions:

i. Rules or principles

Guidance can be provided in the form of concrete rules with clear and specific conditions and normative conclusions, or in the form of more abstract principles. An example of the former is: 'Information about ethnic background may not be used in analytics pertaining to X'. An example of principle-based guidance is 'Care should be taken not to incorporate variables in the analysis that are proxies for ethnicity'.

ii. General or specific

The principles or values that serve as a background in the model can be either generic (as in table 1), or be made more specific for the public sector or even specific domains within the public sector. Within the domain of healthcare, for instance, values such as privacy, autonomy, empowerment could be defined (preferably elaborated in more detail, such as 'autonomy to make decisions about conditions and treatment within the home sphere').

2. The control loop

The second layer consists of the core accountability scheme. We base our accountability approach on the model developed within the A4Cloud project. The A4Cloud project centers around '*accountability under data protection laws for personal data processed in cloud service provision ecosystems*'.[44] However, the methods and conceptual framework developed there are also relevant and easy to apply outside of the cloud computing context. Accountability, in the A4 Cloud project, goes beyond mere

[43] KNAW (2016) *Ethische en Juridische Aspecten van Informaticaonderzoek*. Available at https://www.knaw.nl/nl/actueel/publicaties/ethische-en-juridische-aspecten-van-informaticaonderzoek
[44] More information available at http://www.a4cloud.eu/scope

compliance with the law as it also requires demands ethical reflection and respect for the (personal) information that is being used.[45] The A4 Cloud project developers note that 'Accountability... is considered not to be a state, but rather a learning process in which organisations mature with respect to good governance. In our view the notion of accountability is not only descriptive ('accountability of some agent to some other agent for some state of affairs'), but also has a strong normative claim ('the promise of fair and equitable governance') that requires interaction between the two agents to establish the norms and reflect upon what responsible behaviour is, in line with scholarship and practice in public administration.'[46]

In this model, an accountable organisation:

- defines what it does,
- performs what it has defined,
- monitors how it acts,
- remedies any discrepancies between the definition of what should occur (norms) and what is actually occurring (behaviour),
- explains and justifies its actions.[47]

The definition of the norms is based on legal requirements (for instance, the data protection and domain specific regulation), but also the principles as defined in layer 1 (background). In the governmental sector it is also of importance to prevent function creep, which can be prevented best by also reviewing activities after the application stage. With the control loop actors have to be able to explain and justify actions, meaning that detection of actions that are strictly outside of a competency are easier to detect. It is important that for any data science project the relevant stakeholders convene to define the norms the project should adhere to. The relevant actors should also discuss up front what can be made public and what not.

In the United Kingdom this development is visible, for example in the Cabinet Office guidelines described above. The underlying idea is to start with a clear idea of the project up front and to get external actors that are relevant to the project in line with the project goals. The guidelines also focus on the idea of a 'loop' or circle, meaning that follow up on the project goals is also required and that placing checks on the data use is a continuous process.

In terms of accountability in data protection, accountability quickly turns into a form of keeping track of what data is being processed and why and keeping a paper trail of all relevant activities. Instrumental in this layer is the notion of impact assessment. Impact assessment models are based on the idea that goals and risks are thought through prior to actual action. Various models for doing this have been and are being developed. Relevant to DS in the public sector is the Data Protection Impact Assessment as

[45] More information available at http://www.a4cloud.eu/content/aspects-accountability

[46] Jaatun, M. G., Pearson, S. Gittler, F., Leenes, R.E., Niezen, M.. Enhancing accountability in the cloud. International Journal of Information Management (2016), http://dx.doi.org/10.1016/j.ijinfomgt.2016.03.004

[47] Idem.

mandated by art. 35 of the GDPR. This could provide a foundation for developing a more extensive and elaborate model for data science projects.

In this core accountability layer, choices need to be made on two dimensions:

i. Extensive or marginal reporting - impact assessment vs. description of goals and methods?

The first dimension relates to the scope and depth of the reporting related to the accountability loop. An extensive variant imposes a full blown impact assessment model on each project meaning that goals, risks, scope etc of the project needs to be formulated in clear norms, that during the project lifecycle will be reported in view of the norms, and that any deviation are being mitigated and also reported. In its minimal form, a summary of relevant goals, aims, risks, etc are being formulated that serve as guiding lights for the project.

ii. Internal or external reporting?

Secondly, a choice needs to be made with respect to the audience and purpose of the reporting/assessment. The record-keeping and reporting process could on the one hand serve as guidance for the project team and its managers only, but on the other hand it could serve as public accountability for the project, providing transparency and allowing external scrutiny (and thus full accountability).

3. Enforcement and oversight

Although the first two layers of oversight could already contribute to responsible data science, oversight and enforcement (of both internal and external norms) further strengthen the model.

In this enforcement layer, choices need to be made around the following issues:

i. Internal versus external oversight

As with reporting, the oversight of a given project can be organised either within the project organisation (or one level up), or can be fully external and thus performed by an independent authority. If oversight is external, then other choices also need to be made with respect to who this authority should be: for example, it could be a domain-specific board within government, an existing governmental authority whose function is extended to this responsibility, or a body from civil society.

ii. Ad hoc versus structural oversight

Oversight can be done on an ad hoc basis, with a regulator (either internal or external) formed specifically for a particular project, or be more structural, for example a Public Data Science Authority that oversees all (or large) public sector data science projects.

iii. Level/distance to the oversight organisation

The ‘distance’ to the project organisation is also an important point of choice. A regulator can be responsible for overseeing a sector, a ministry, a layer of government, or other units. This choice of level and position has implications for the skills and knowledge required of the regulator: one who is responsible as well as other factors.

For example, the UK system outlined in the previous section would fall under the categories of *principle-based* guidance that is *specific* to the public sector; with currently minimal and *internal* record-keeping and reporting (this is implied though not specified in the existing guidelines); and finally no oversight is implied, though the involvement of the Cabinet Office in producing guidelines and in the promise to set up a Council of Data Science Ethics implies that external oversight a the government-wide level will be chosen. In contrast the French approach consists of *rules* (expressed in the form of law) that are *generic* and thus designed to be used across domains and sectors. In theory this implies extensive record-keeping though more for audit purposes than for explicit reporting to a particular body; and finally, *external* oversight that is also *structural*, as it is performed through the courts.

For the Dutch case, one relevant factor in making these decisions is the extent to which the practice and effects of data science will be decentralised, given that municipalities hold population records and other relevant data, and can channel these to the central government where they can be combined with data collected by ministries. Municipalities are already cooperating with the Dutch Central Bureau for Statistics (CBS) by using its data for policymaking. Several cities are also setting up Urban Data Centers together with the CBS, taking advantage of both its expertise and the possibility to connect and compare data and approaches among municipalities.[48] It could be that similar initiatives will be set up with other governmental institutions on the central level, who also have control over large datasets. However, the question is also to which extent there is an oversight of data science practices on the central government level: different departments might have varying practices of data science, choices need to be made on which scale oversight will take place.

Beyond the categories of tax and security data, the Dutch governmental landscape presents many different and highly distributed domains where problematic uses of data may emerge. Examples include the country's 388 municipalities, the employment and income domain (SUWI), the welfare and social security infrastructures, and the administrative infrastructures in the health domain. This suggests that sector-specific legislation may be the most effective option for oversight and accountability. Pointing the discussion of principles and guidelines in a sectoral direction helps to avoid replicating the oversight conundrum of the UK's 'nudge unit', whose work reaches across sectors including the tax authorities, public health and everyday behaviour, and which aims to invisibly influence behaviour in ways that are often openly experimental. Departments set up to conduct experiments based on large-scale data collection and analytics but that do not distinguish between domains and types of experiment are difficult to make accountable. However, the WRR's recent report on the security sector's use of data scientific approaches[49] demonstrates that it is also important to establish a cross-sectoral dialogue

[48] Mudde L. (2017) CBS gaat lokaal met Urban Data Centers: Decentraal Bureau voor de Statistiek. VNG Magazine nummer 2. Retrieved from https://vng.nl/cbs-gaat-lokaal-met-urban-data-centers-decentraal-bureau-voor-de-statistiek

[49] Scientific Council for Government Policy (WRR) (2016). Big Data in een vrije en veilige samenleving. Amsterdam: Amsterdam University Press. Available at https://www.wrr.nl/publicaties/rapporten/2016/04/28/big-data-in-een-vrije-en-veilige-samenleving

about the development of guidelines and principles in order not to end up with vague overarching norms that do not hold departments to account effectively.

For the Netherlands, a model is necessary that can both offer efficient oversight on this decentralised scale, and can grow with the government's use of data science over the coming decade. The uneven use of data science across different domains within the public sector is like to change over time, as new data sources become available and different departments and groups are challenged to use them to create efficiency and insights. The current state of affairs - extensive use of data science approaches in certain key areas such as fraud detection and security, combined with much less use across the bureaucracy in general - suggests that it is necessary to establish a model for responsible data science that can grow with the government's use of data. Such a model will both guide and oversee the process of development from primarily data-informed to more data-driven policy over time, as testing out the benefits and opportunities of datafication moves higher on the list of governmental priorities.

Conclusion: a sectoral approach

Putting together the different elements of the model we have presented above in different combinations will produce different results, but the elements chosen will also be determined by place and politics. It is unlikely that replicating or scaling up models from other countries will work as well as determining how guidelines fit best with the operations of each sector. What is best for the Netherlands, therefore, will need to be determined through discussion and debate so that it reflects the Dutch government's aims and values.

Given these elements and criteria, the most appropriate solution at present seems to lie in sectoral frameworks within which actors using data scientific methods must follow the accountability process to make clear what they are doing in terms of data projects. The sectoral approach makes it possible to use a combination of overarching *principles* (of the kind set out by the UK Cabinet Office), and a set of *rules* that give clearer guidance on sector-specific issues. With a sectoral approach it becomes possible to indicate particular red flags (for example, 'in anti-fraud operations, when sharing data with law enforcement, x and y criteria must be observed'; or 'when merging public with private-sector health data, precautions must be taken to ensure x does not happen'). Such a combination of strategies both raises awareness amongst government actors that merging, linking or enriching digital databases may raise risks so that attention must be paid to the possible consequences, but also provides specific measures which help to avoid those negative consequences. This combination of awareness and concrete rules is likely to be more productive in shaping behaviour than either on its own, since principles on their own are not enforceable, but without such principles rules are likely to be experienced as restrictive rather than helpful.

A simplified version of the model can be found in the table below:

Table 2. The core model for Dutch governmental actors for data ethics

Principles		
Central	Starting point and foundation	Applicable to: - individuals - relations between individuals - social system
Accountability		
All levels	Control loop	Public-facing elements combined with internal accountability
Oversight & Enforcement		
Sectoral	Identify gaps and create new oversight where necessary	Possibility of sanctions

The priorities for the Dutch model, then, begin with establishing principles centrally. Such principles, once activated, should play a role across different levels from the individual to the departmental and governmental. As a common basis, they can draw together the government's policy on ethical data practices across decentralised elements such as municipalities, and can form the framing for other checks and balances as they are developed. The establishment of principles can then be followed by building sectoral oversight so that the control loop described here can function on different levels, potentially in relation to existing structures. For example, the Algemene Rekenkamer and the AP (or its successor under the GDPR) both provide potential models for structural oversight. One possible way to address the decentralised municipality system is to centre oversight in a body that already coordinates between municipalities. One candidate would be the CBS, already a meeting-point for operational information on activities with digital data across municipalities.

One challenge is to prevent the burden of oversight being placed solely on one organ, which would make addressing the full range of data ethics considerations on a governmental scale impossible. To take one example, the AP - one likely candidate for such an approach - is already working to capacity implementing the GDPR and the Police Directive into Dutch law and, despite its existing work on data protection, would be challenged to also develop data ethics oversight capacity on top of its current portfolio. Instead, using multiple structural options as appropriate at the sectoral level may prove the most efficient solution, with the goal of achieving both 'horizontal' accountability (awareness of principles, ensuring departments report internally and to immediate management) and oversight from above (via the legal framework, higher-level oversight bodies, external bodies and/or routes to civil society accountability). There is an advantage to also incorporating a public-facing element of oversight, whether through open reporting requirements or the involvement of civil society institutions. Such a public-facing component adds to the legitimacy of government's data projects within society, and would contribute to public awareness of, and debate about, the use of data in governance. This would be more

of a challenge for those sectors that cannot make documents publicly available, but even there it may be possible to combine some public-facing elements with internal accountability.

The challenge of creating structures for oversight in this domain is substantial. This is because the use of data by government is also substantial, and is likely to grow exponentially as capacity and opportunities for collaboration increase. Resolving such issues as establishing overall principles, the best fit among different bodies for sectoral oversight, and the role of civil society in the accountability process all constitute, however, a valuable process for negotiating the standards for the use of data science by the Dutch government. The ideal outcome of the structural approach posited here would be to make data ethics, similarly to privacy by design, a mindset and a continuous process rather than a task of compliance with norms. Although this mindset would take time to become embedded, the process of negotiating, discussing and defining ethical principles is essential if accountability is to be legitimate and enforceable.

Literature overview*

Algemene Rekenkamer. Trendrapport open data (2015). Retrieved from http://www.rekenkamer.nl/Publicaties/Onderzoeksrapporten/Introducties/2015/03/Trendrapport_open_data_2015

Demmke, C. and Moilanen, T. (2012). Effectiveness of Ethics and Good Governance in Central Administration of EU-27: Evaluating Reform Outcomes in the Context of the Financial Crisis. Available at http://www.integriteitoverheid.nl/fileadmin/BIOS/data/Publicaties/Downloads/Effectiveness_of_Public-service_Ethics_and_Good_Governance_in_the_Central_Administrations_of_the_EU-27.pdf

European Commission for Democracy through Law (Venice Commission). Report on the democratic oversight of signals intelligence agencies (2015). Retrieved from http://www.venice.coe.int/webforms/documents/default.aspx?pdffile=CDL-AD(2015)011-e

Floridi, L., Taddeo, M. (2016) What is data ethics? Phil. Trans. R. Soc. A 374 (2083), 20160360. DOI: 10.1098/rsta.2016.0360.

Franken, H. (1993) 'Kanttekeningen bij het automatiseren van beschikkingen' in Beschikken en Automatiseren. Preadvies voor de Vereniging voor Bestuursrecht, var-reeks nr. 110, Den Haag.

Jaatun, M. G., Pearson, S. Gittler, F., Leenes, R.E., Niezen, M.. Enhancing accountability in the cloud. International Journal of Information Management (2016), http://dx.doi.org/10.1016/j.ijinfomgt.2016.03

Kitchin R. (2016). The ethics of smart cities and urban science. *Phil. Trans. R. Soc.* A 374: 20160115. http://dx.doi.org/10.1098/rsta.2016.0115

KNAW (2016) *Ethische en Juridische Aspecten van Informaticaonderzoek*. Available at https://www.knaw.nl/nl/actueel/publicaties/ethische-en-juridische-aspecten-van-informaticaonderzoek

Kool, L, J. Timmer & R. van Est (2015). De datagedreven samenleving - Achtergrondstudie. Den Haag: Rathenau Instituut. Retrieved from https://www.rathenau.nl/nl/files/rapportdatagedrevensamenlevingrathenau-instituutpdf

Mayer-Schönberger, V. and Cukier K. (2013). Big Data: A Revolution that Will Transform How We Live, Work, and Think. London: John Murray.

Michigan Department of Transportation & Center for Automotive Research (2014). ITS Data Ethics in the Public Sector. Available at https://www.michigan.gov/documents/mdot/06-14-2014_ITS_Data_Ethics_in_the_Public_Sector_464226_7.pdf

Ministry of Justice. Antwoorden Kamervragen Big data experiment legt onzichtbare criminaliteit bloot (2016, September 1). Retrieved from https://www.rijksoverheid.nl/documenten/kamerstukken/2016/09/01/antwoorden-kamervragen-big-data-experiment-legt-onzichtbare-criminaliteit-bloot

Mudde L. (2017) CBS gaat lokaal met Urban Data Centers: Decentraal Bureau voor de Statistiek. VNG Magazine nummer 2. Retrieved from https://vng.nl/cbs-gaat-lokaal-met-urban-data-centers-decentraal-bureau-voor-de-statistiek

Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

PiLab (2012), iOverheid, burger in beeld (final report). Accessible at: https://pilab.nl/onewebmedia/PI.lab-2012-R11216-Eindrapport-iOverheid-burger-in-beeld1.pdf

Scientific Council for Government Policy (WRR) (2011). iGovernment: synthesis of WRR report 86. The Hague: WRR. Available at https://www.wrr.nl/binaries/wrr/documenten/rapporten/2011/03/15/ioverheid/ioverheid.pdf

Scientific Council for Government Policy (WRR) (2016). Big Data in een vrije en veilige samenleving. Amsterdam: Amsterdam University Press. Available at https://www.wrr.nl/publicaties/rapporten/2016/04/28/big-data-in-een-vrije-en-veilige-samenleving

* For an overview of all source materials used for this report, please refer to the footnotes.