An audit of algorithms

Nine algorithms used by the Dutch government

2022

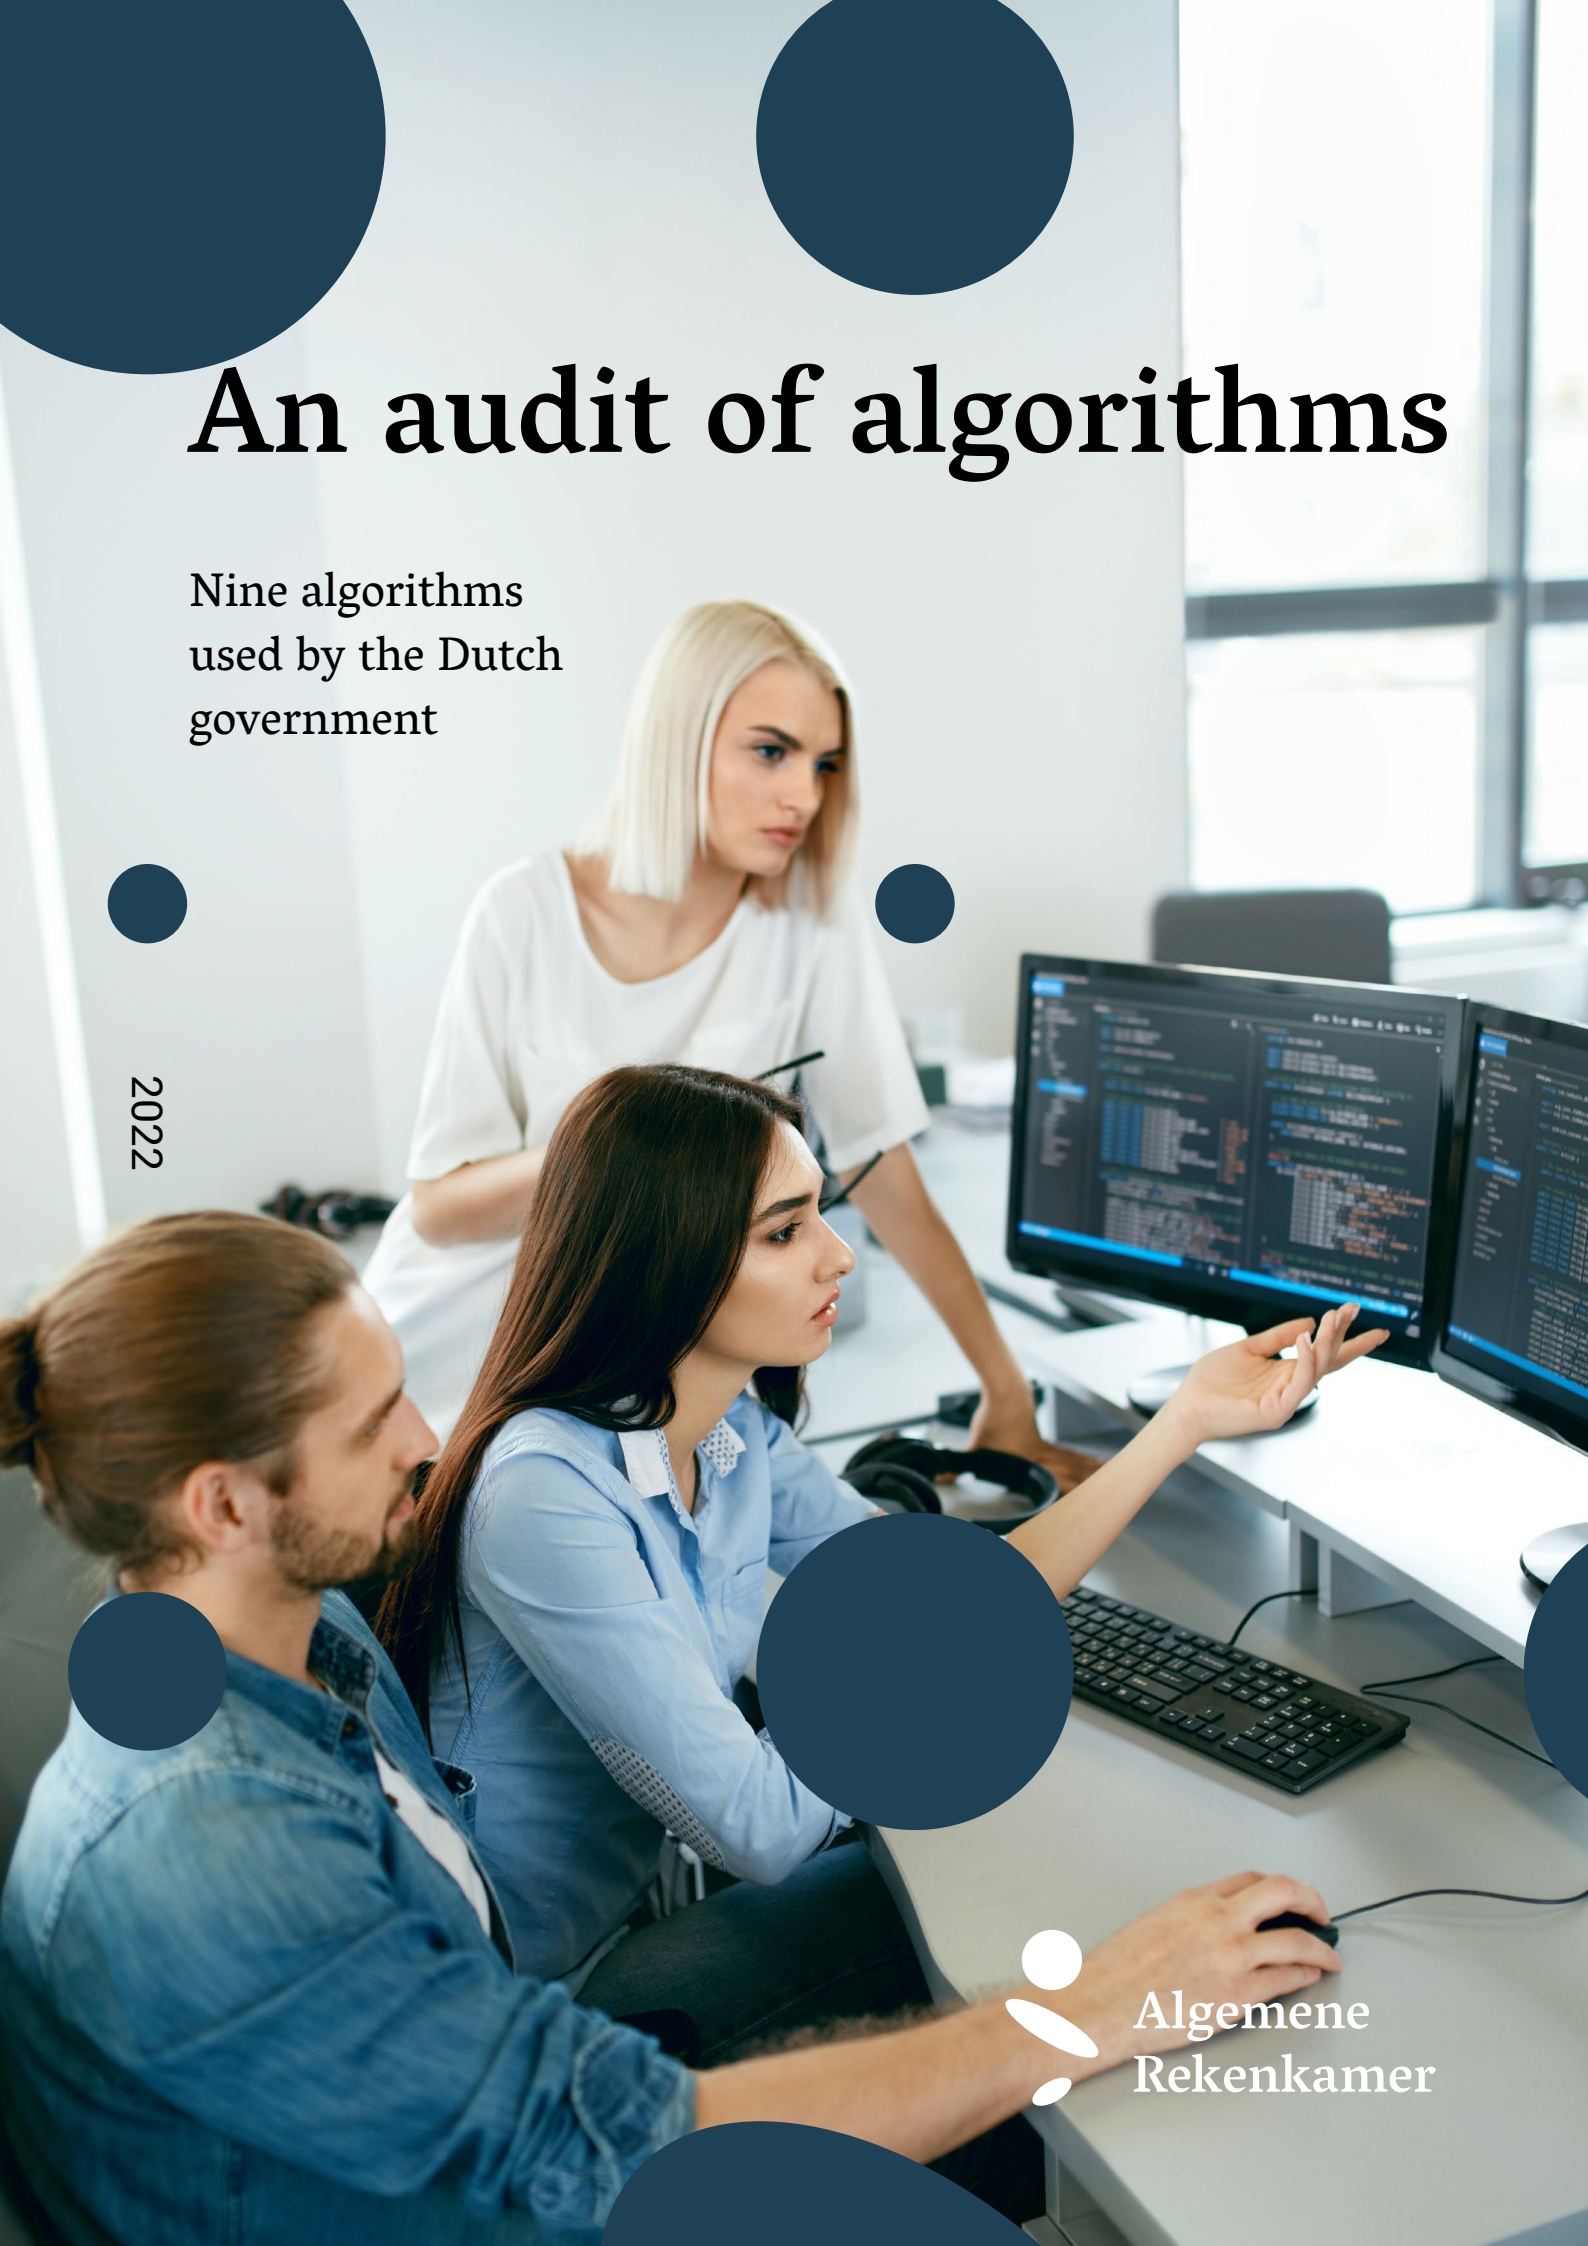

Algemene Rekenkamer

Inhoud

1. Summary

The government takes millions of decisions every month, for example on whether someone is or is not eligible for housing benefit or whether someone should be fined for speeding. In taking these decisions, the government uses algorithms. An algorithm is a set of instructions that a computer follows automatically to solve a problem or answer a question. This enables the government to perform some of its tasks either automatically or quicker. After all, government officials cannot perform all these tasks on their own. If this were the case, many more officials would be needed, and all of these decisions would take much longer. In other words, the government could not operate properly without using algorithms.

Using algorithms also creates certain risks. If they do not work properly, groups of people may be discriminated against. Private citizens have also expressed concerns that the use of algorithms may give the government too much information about their private lives. Others feel uneasy about the possibility of the government using algorithms without telling anyone, and of the government not being able to explain how certain decisions have been taken. This applies not only to technically complex algorithms, but also to simple ones. In fact, simple algorithms in particular are just as likely to have a big impact on private citizens as the more complex algorithms.

This is why we have audited the use of algorithms for a second time. Our first audit, entitled 'Understanding algorithms', looked at the purposes for which the central government uses algorithms, and at the types of algorithms it uses. We have already mentioned some of the risks associated with algorithms. It is possible to assess the severity of these in relation to any individual algorithm. In our first audit, we developed

a set of guidelines for assessing these risks and listed the requirements that algorithms must meet. Taken together, these requirements form what we have called an 'audit framework'. We can use this audit framework to check whether the government makes proper use of algorithms (Netherlands Court of Audit, 2021). In this second audit, we audited nine algorithms by seeing whether they satisfy the requirements of our audit framework.

1.1 Conclusions

Our audit revealed that three of the nine algorithms that we assessed satisfy the requirements of our audit framework. This shows that algorithms can be used both fairly and judiciously.

We also found that six of the nine algorithms tested do not comply with all the requirements set out in our audit framework. There is great scope for improvement here and we will discuss the three main improvements below.

We found that, in some cases, no clear agreements had been made about what the algorithm is required to do, and in other cases that no agreements had been made for checking whether the algorithm was operating in accordance with its intended purpose. This is especially important if the government asks another organisation to develop or manage an algorithm. The risk here is that government may be unable to monitor whether the algorithm is being used in a safe way.

Another area of improvement relates to IT controls. Algorithms use data obtained from people and businesses. Organisations using algorithms must make clear agreements about who is and is not allowed to work with the data concerned. We found that such agreements had not been made by all the organisations involved, which means that these organisations cannot be absolutely sure whether the data obtained from private citizens and businesses are safe.

Bias is a third area that is open to improvement. Bias means that an algorithm produces incorrect results in relation to certain groups of people because there is an error in the algorithm itself, or because the wrong information was fed into the algorithm when it was developed. Our audit revealed that organisations do not check whether the algorithms they use contain this type of error. This means that they do not know for certain whether their algorithms generate the same results for all groups of people.

1.2 Recommendations

We believe that the government should use algorithms in a responsible manner. Our recommendations for ministers as to how they may best achieve this are set out below.

We repeat the two recommendations that we made in our first audit 'Understanding algorithms' (Netherlands Court of Audit, 2021):

1. Particularly where algorithms are outsourced or bought from an external supplier, document agreements on their use and make effective arrangements for monitoring compliance on an ongoing basis.
2. Ensure that algorithms and the data required for their operation are protected by effective IT general controls.

These are supplemented by two new recommendations:

3. Check regularly – both during the design of algorithms and during their use – for the effect of bias, in order to prevent any undesirable systematic variations in relation to specific individuals or groups of individuals.
4. When supervising arm's-length institutions, take explicit account of how algorithms are used to perform public services.

2. About this audit

2.1 Why did we perform this audit?

Algorithms are used to automate activities, solve problems or make forecasts. An algorithm is a set of instructions that a computer follows automatically when performing calculations to solve a problem or answer a question. Organisations often use algorithms to help with their operational management and service delivery. Without algorithms, it would be impossible, for instance, to process large volumes of grant or benefit applications within the relevant time limits. Algorithms also allow organisations to use people and resources in a carefully targeted manner for checks and inspections. They can also make decision-making procedures more transparent and easier to audit. This is because the technical aspects of an algorithm, such as the data sources it uses and those aspects of the data on which it particularly relies, are defined in the form of instructions. This makes algorithmic decision-making in theory more transparent than human decision-making.

The use of algorithms creates opportunities for raising the efficiency and transparency of government. However, any careless use of algorithms also entails risks, not only for private citizens and businesses, but also for the government itself. Public authorities hold a growing volume of data obtained from private citizens, and the use of algorithms is becoming more and more commonplace. This means that private citizens are being confronted to a growing degree with the consequences of the government's use of data and algorithms (National Ombudsman, 2021). Concerns have been expressed, however, about the government's use of algorithms, for instance on the grounds that they may lead to discrimination. These concerns have emerged in news reports about

the government's use of algorithms, such as in an article on biased results in the detection of benefit fraud, which was prompted by an audit performed by the Rotterdam Court of Audit (Trouw, 2021, and Rotterdam Court of Audit, 2021). Another article in a national daily was on what became known as the 'grocery shopping case', in which algorithms were used to detect cases of fraud. (De Volkskrant, 2021). The Lower House of the Dutch Parliament has also raised concerns about the use of algorithms: a motion calling on the government to stop using discriminating algorithms was adopted in the Parliamentary debate on the childcare benefits scandal (Lower House of the Dutch Parliament, 2021).

Both the opportunities and the concerns surrounding the use of algorithms have prompted us to audit the use of algorithms by the central government. This is our second audit of algorithms. In our first audit report, entitled 'Understanding algorithms', we developed a framework for auditing the use of algorithms (Netherlands Court of Audit, 2021). This audit framework is a practical tool that public bodies can use to see whether their algorithms comply with certain predefined quality criteria, and also whether the associated risks have been properly identified, and/or whether action has been taken to mitigate these risks.

In this second audit, we assessed nine algorithms to see whether they met the requirements set out in our audit framework. In the case of one of the algorithms, we describe the context in which it operates, as an algorithm does not operate on its own, but is always part of a wider policy planning or policy implementation process.

2.2 What did we audit and how did we perform our audit?

We audited the use of algorithms by central government and its associated organisations. This audit consists of two parts:

1. The first part outlines the context of one specific algorithm.
2. The second part describes our assessment of the operation of nine algorithms, using our audit framework.

2.2.1 One algorithm in context

In this part of our audit, we examined one of the algorithms in detail, i.e. the risk model used by the Netherlands Enterprise Agency for assessing grant applications under the government's TVL scheme for the reimbursement of fixed costs. This scheme is designed to help business owners who suffer a loss of turnover due to the effects of government action to combat the COVID-19 pandemic. We explain how the algorithm

has been incorporated in the policy process and in the work performed by officials, and how it has affected both private citizens and businesses. Our audit centred on the following questions:

- How has the algorithm been incorporated in the policy process?
- Which role do humans play in working with the algorithm?
- How does the government arrive at a decision on the use of the algorithm?
- What is the impact on private citizens and businesses?

In order to answer these questions, we interviewed a number of staff working for the Netherlands Enterprise Agency, as well as officials at the Ministry of Economic Affairs and Climate Policy. We also analysed relevant process documents and other documents.

2.2.2 Audit of the use of algorithms

Which algorithms did we audit?

This audit builds on the stocktaking survey that we performed as part of our first audit (Netherlands Court of Audit, 2021). We supplemented this initial list of algorithms with information obtained from a number of sources, including a list of algorithms used by the Netherlands Organisation of Applied Scientific Research for providing public services, the National Ombudsman's views on the use of data and algorithms by the government, and various media reports (National Ombudsman, 2021 and Netherlands Organisation of Applied Scientific Research, 2021). We also quizzed the audit teams performing our annual accountability audits.

We then selected algorithms by applying the following criteria:

- *Impact on private citizens or businesses*
  The algorithms – or processes in which the algorithm plays a part – we selected have a big impact on individuals or businesses.
- *Risk-centred*
  We audited those algorithms which we believe have the greatest risk of not being used in the correct way.
- *Different domains*
  We selected algorithms from a number of different domains, e.g. the social domain and the security domain.
- *In operation*
  The selected algorithms also had to be currently in operation. We excluded from our audit any algorithms that had been discontinued or that were still in a pilot stage.

- *Different types of algorithms*
  In addition to the above, we selected different types of algorithms, ranging from technically simple algorithms, such as decision trees and electronic data interchange systems, to technically more complex algorithms, such as image recognition systems and self-learning applications.

In the end we selected nine algorithms for the purpose of this audit. These are listed in Table 1 below.

Table 1. *List of algorithms selected for the audit*

Ministry	Organisation	Status	Description of algorithm
Ministry of the Interior and Kingdom Relations	**National Office for Identity Data (RvIG)**	Executive agency	Helps staff to assess the quality of photographs for identity documents
Ministry of Economic Affairs and Climate Policy	**Netherlands Enterprise Agency (RVO)**	Executive agency	Risk classification model used for assessing grant applications under the TVL scheme for the reimbursement of fixed costs
Ministry of Finance	**Benefits department (Toeslagen)**	Ministry department	Helps staff to assess housing benefit applications as part of the system used for awarding benefits and tax allowances
Ministry of Infrastructure and Water Management	**Central Office for Motor Vehicle Driver Testing (CBR)**	Legal person with a statutory task/ autonomous administrative authority	Helps staff to assess whether people are medically fit to drive a motor vehicle
Ministry of Justice and Security	**Police force**	Legal person with a statutory task	The Criminality Anticipation System (CAS) forecasts where and when there is a high risk of incidents occurring
Ministry of Justice and Security	**Directorate-General for Migration (DGM)**	Ministry department	Performs an intelligent search of personal data on aliens to determine whether a subject was previously registered in the Netherlands
Ministry of Justice and Security	**Central Judicial Collection Agency (CJIB)**	Executive agency	Links data on traffic fines with traffic offences attributed to specific vehicle registration numbers

Ministry	Organisation	Status	Description of algorithm
Ministry of Social Affairs and Employment	**Benefits Intelligence Agency (IB)**	Legal person with a statutory task	Alerts local authorities to the need for performing regularity checks on social assistance benefits
Ministry of Social Affairs and Employment	**Social Insurance Bank (SVB)**	Legal person with a statutory task/ autonomous administrative authority	Helps staff in assessing applications for state pensions

The algorithms included in our audit are used both by central government itself and by organisations operating at arm's length from central government. There are different types of arm's-length organisations. In this audit, we audited a number of organisations known as 'legal persons with a statutory task' and 'autonomous administrative authorities'. These are independent organisations that operate at arm's length from central government and which perform certain government tasks. A special 'institutional act' governs each organisation's powers and responsibilities and sets out the nature of its relationship with the minister. As a result, the minister is responsible for supervising these organisations and is required to report to the House of Representatives on their operation and performance.

Our selection is not a cross-section of all the algorithms used by central government as it is based on the degree of risk that they pose. This means that it is not possible to use our audit findings as a basis for drawing general conclusions about processes, organisations or central government as a whole.

Differentiating between different types of algorithms

The algorithms included in our audit differ both in complexity and in terms of the role that they play in implementing policies.

Complexity of the algorithm

Algorithms come in all shapes and sizes. They range from technically simple algorithms, such as decision trees and electronic data interchange systems, and technically complex algorithms, such as image recognition systems and self-learning algorithms (Ministry of Justice and Security, 2019).

The algorithms included in our audit break down as follows:

- Technically simple algorithms:
  - Decision trees: Benefit department, Central Office for Motor Vehicle Driver Testing and Social Insurance Bank
  - Electronic data interchange systems: Central Judicial Collection Agency and Benefits Intelligence Agency
  - Search function: Directorate-General for Migration at the Ministry of Justice and Security
  - Risk classification model: Netherlands Enterprise Agency
- Technically complex algorithms:
  - Image recognition system: National Office for Identity Data
  - Self-learning algorithm: police force

We examined different types of algorithms

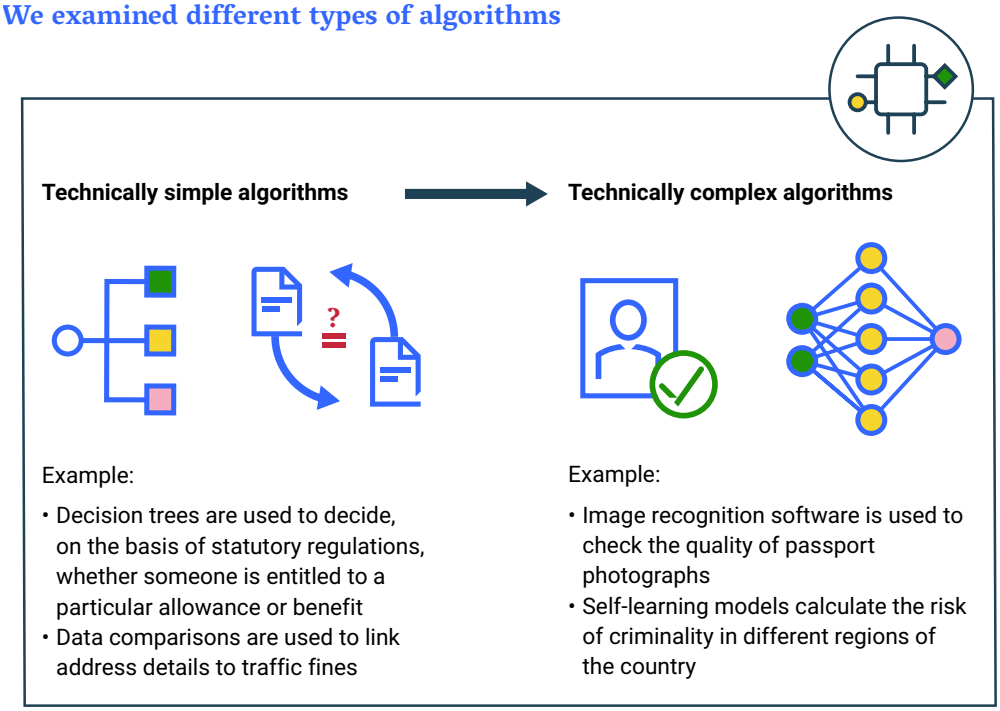

The algorithm's role in implementing government policy

In addition to the above, the algorithms that we audited may be broken down into algorithms supporting policy implementation and algorithms supporting automated (or partially automated) decision-making. Algorithms may for instance be used to support policy implementation by supplying missing data making forecasts, or comparing data for the purpose of risk management. In the case of the second category, applications received from private citizens are sent directly to the algorithm, which then takes automated decisions that have a direct impact on private citizens or businesses. If an application satisfies all the requirements for the award of the allowance or benefit in question, it then follows a completely automated process in which the algorithm

decides. If the application does not satisfy all the requirements, a government official performs a manual check.

The algorithms we audited break down as follows:

- Algorithms that support policy implementation: the Benefits Intelligence Agency, the Central Judicial Collection Agency, the Directorate-General for Migration at the Ministry of Justice and Security, the National Office for Identity Data, and the police force.
- Algorithms that take automated (or partially automated) decisions that have a direct impact on private citizens: the Netherlands Enterprise Agency, the Central Office for Motor Vehicle Driver Testing, the Social Insurance Bank, and the Benefits department of the Tax and Customs Administration.

The following table shows that none of the algorithms we audited is both technically complex and also responsible for taking automated decisions.

Table 2. *Breakdown of the audited algorithms*

	Automatic decision-making	Supporting policy implementation
Technically simple	• Housing benefit (Benefits department) • Assessment of medical fitness to drive (Central Office for Motor Vehicle Driver Testing) • Applications for state pensions (Social Insurance Bank) • Applications for grants under the TVL scheme for the reimbursement of fixed costs (Netherlands Enterprise Agency)	• Traffic fines (Central Judicial Collection Agency) • Social assistance benefits (Benefits Intelligence Agency) • Aliens (Directorate-General for Migration at the Ministry of Justice and Security)
Technically complex		• Assessing the quality of passport photographs (National Office for Identity Data) • Assessing the risk of criminality (Police force)

Audit framework for algorithms

We audited the use of algorithms based on our audit framework for algorithms (Netherlands Court of Audit, 2021). This is a practical tool that the central government can use to manage the main risks associated with the use of algorithms.
The audit framework assesses algorithms from five different perspectives:

1. governance and accountability;
2. model and data;
3. privacy;
4. IT general controls;
5. ethics.

We identified the main risks relating to each of the above perspectives, and linked the aspects we wanted to examine and the audit questions we wanted to answer, to these risks. The risks and audit questions included in the audit framework are listed in Appendix 3.

We used our audit framework as the basis for auditing the algorithms

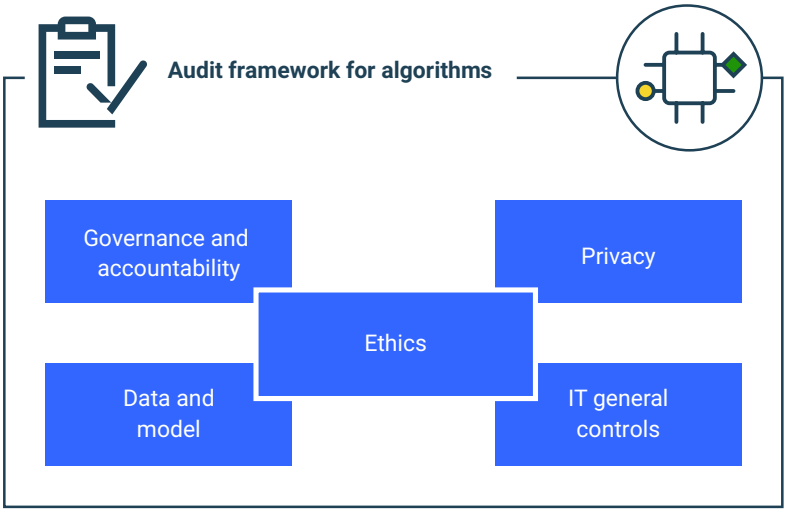

In examining the algorithms from an ethical perspective, we applied the following four ethical principles:

1. respect for human autonomy;
2. the prevention of damage;
3. fairness;
4. explicability and transparency.

Rather than constituting a genuinely separate aspect, ethics form an integral part of the four other aspects of our audit framework. This means that the ethical principles are linked to the risks relating to the other four aspects, i.e. governance and accountability; model and data; privacy; and IT general controls.

Assessing algorithms

By answering all the questions included in the audit framework and allotting scores, we formed a picture of the extent to which the risks associated with a particular algorithm have been mitigated. The degree of risk associated with a specific algorithm depends in part on its impact on private citizens.

Our approach to this audit is explained in detail in Appendix 1.

2.3 Format of this report

This audit report consists of two parts. Chapter 3 describes the context of one specific algorithm and analyses how this algorithm is used as part of a policy process. Chapter 4 contains our main observations and criticisms arising from our audit of the use of algorithms by central government. Chapter 5 presents our conclusions and recommendations.

3. One specific algorithm in context

In order to be able to analyse and explain how an algorithm works, we need to be aware of the context in which it is used. This chapter looks at the way in which one particular algorithm helps the Netherlands Enterprise Agency (RVO) to pay out grants under the government's TVL scheme for the reimbursement of fixed costs. This scheme is designed to help business owners confronted by a sharp decline in turnover due to the effects of government action to combat the COVID-19 pandemic. As we indicated in the previous chapter, this is a technically simply algorithm in which part of the decision-making process has been automated. This chapter explains how it has been incorporated in the policy process and in the work performed by government officials, how the government arrived at a decision on the use of the algorithm, and what its impact has been on both private citizens and businesses.

3.1 How the algorithm works

The Netherlands Enterprise Agency uses an algorithm to help it assess grant applications from business owners. The algorithm checks whether each application meets the criteria laid down by the minister. This is done with the aid of a risk classification model: when an application arrives at the Netherlands Enterprise Agency, the first step is to make a rough appraisal of whether the applicant is indeed entitled to a grant. This involves, for example, comparing the information provided by the applicant with other information that he or she has already supplied to the government, for example in the form of a tax return. The application is automatically approved if it is classified by the algorithm as a low-risk application, for example because the applicant is only applying for a small amount of money and there are no indications that any misuse or

improper use has taken place in the past. In that case, no government official intervenes in the process. The money is transferred to the applicant's bank account within five days.

However, if an application is classified as being high-risk, for example because a large amount of money is involved or because there are indications of previous misuse or improper use, a manual check is made by a government official. If evidence of misuse or improper use has been found in relation to a previous grant application, all future grant applications made by the same applicant will be subjected to manual checks. Similarly, if an application is so complex as not to be capable of automatic processing, it is then subjected to a manual check. This means that the application is assessed by an official, who may ask the applicant to supply certain additional information. Once this information has been received, it is then up to the official to decide whether to approve or reject the application.

Algorithms can be used in support of policy processes

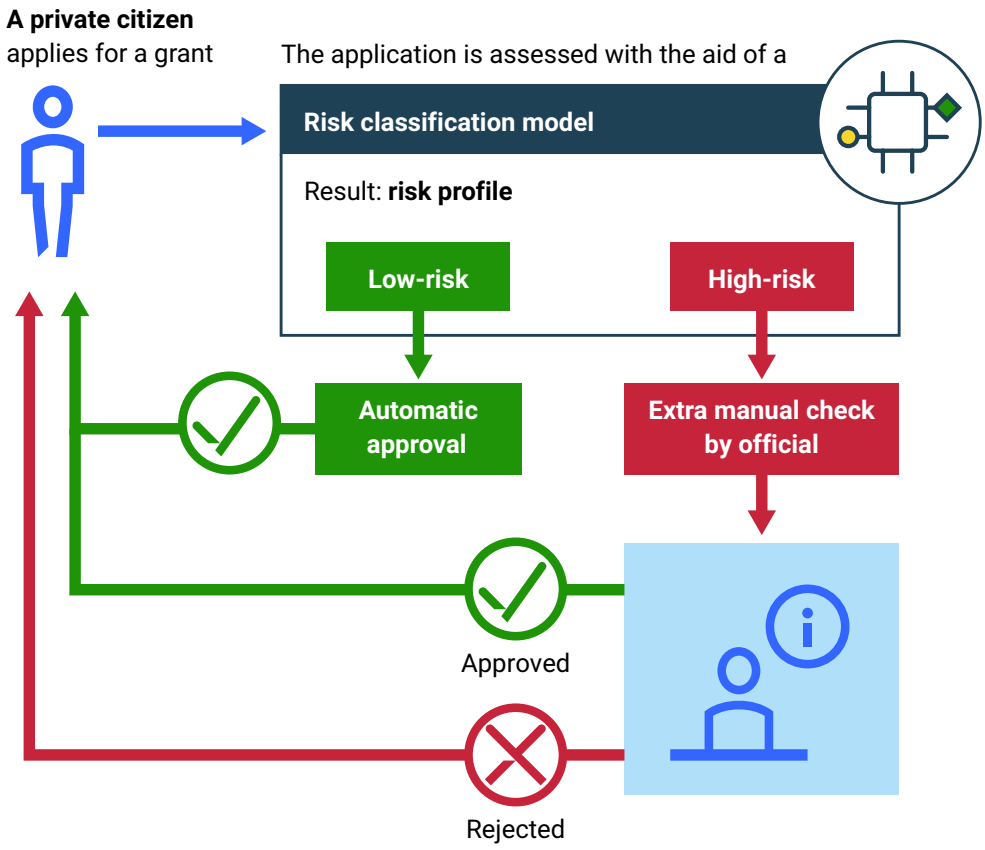

3.2 The decision-making process on the use of an algorithm

Executive agencies such as the Netherlands Enterprise Agency, the Social Insurance Bank or the Tax and Customs Administration are responsible for performing certain government tasks in accordance with laws made by the government and enacted by parliament. They make regular use of algorithms for this purpose. Both the ministry and the executive agency in question should be involved from the outset in the decision-making process on the use of an algorithm. This enables the executive agency to advise the ministry on whether the policy is workable and for both parties to reach agreement on the purpose of the algorithm, the conditions that need to be put in place in order for the algorithm to use used in practice, and on the acceptability of certain risks.

In the case of the TVL scheme (for the reimbursement of fixed costs in connection with the COVID-19 pandemic), the Ministry of Economic Affairs and Climate Policy worked together with the Netherlands Enterprise Agency in planning the scheme and deciding how to strike the right balance between the need to disburse the grants as quickly as possible while minimising the risk of public funds being misused or improperly used, and ensuring that the scheme remained workable in practice. Where the Netherlands Enterprise Agency found that certain adjustments needed to be made to the algorithm, these were submitted to the ministry for approval.

3.3 A limited role for officials

We have already seen that grant applications are not automatically rejected, i.e. without any human intervention, by an algorithm. An official is always involved in any decision to reject a grant application. The question is, though: what does this actually mean in practice? How important is the role played by officials?

In the case of the TVL scheme, an official enters the equation as soon as the algorithm classifies a particular application as high-risk. This may be, for example, because a large amount of money is involved, because certain information is missing, or because there are indications of misuse or improper use having taken place in the past. If one of these criteria is met, the application is forwarded for manual assessment by a human. The algorithm informs the relevant official about the specific aspects on which they should focus, such as information that is missing from the application or the presence of certain inconsistencies. This may be useful, as it helps the official to quickly build up a clear picture of the situation. The official then assesses the application as a whole,

and decides either to approve or reject it, where appropriate after requesting the applicant to provide some further details.

3.4 The algorithm's impact on private citizens and businesses

What impact does the use of an algorithm have on business owners? The use of an algorithm for assessing grant applications under the TVL scheme enables such applications to be dealt with quickly and efficiently. The vast majority of applicants quickly receive the money in question.

If a particular application is selected for a manual check, this means that it must now be assessed by an official. As a consequence, it takes longer to decide whether the application can be approved and it therefore also takes longer for the money in question to be transferred to the applicant's bank account.

Moreover, if evidence of misuse or improper use has been found in relation to a previous grant application, all future grant applications made by the same applicant will be subjected to manual checks.

The use of algorithms can save time and staff capacity

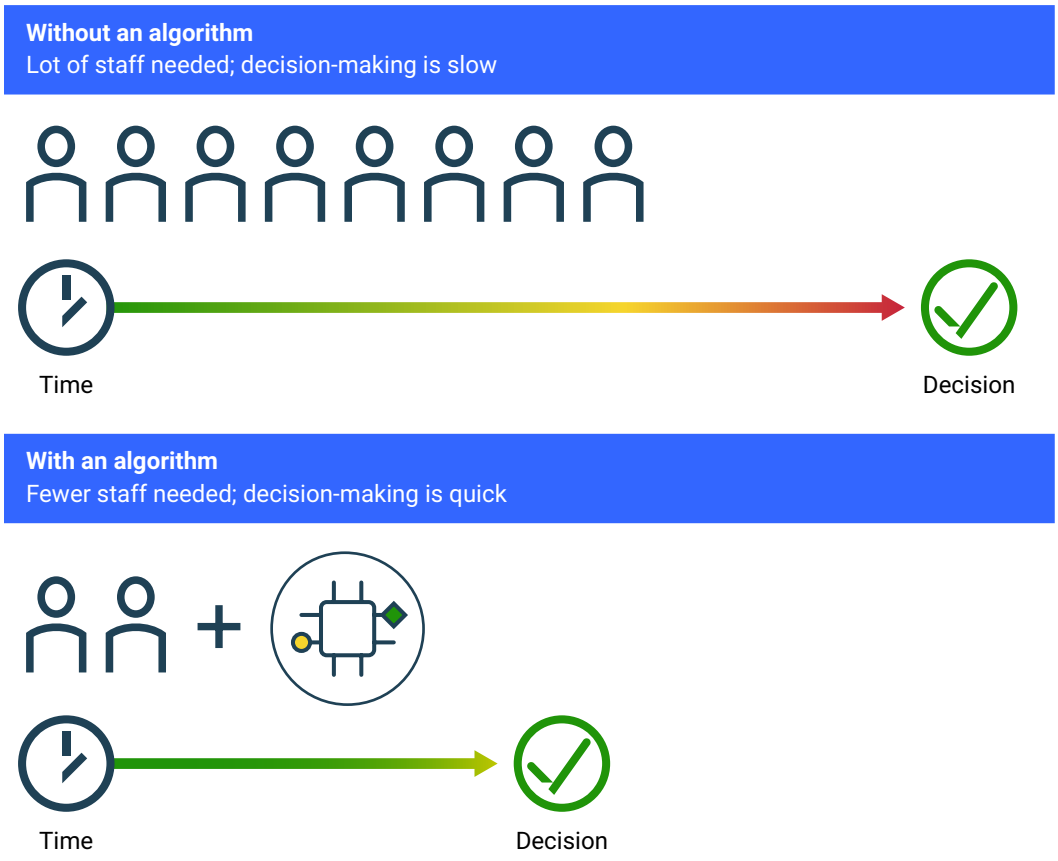

4. An assessment of the use of nine algorithms

We assessed nine algorithms with the aid of our audit framework for algorithms, which looks at algorithms from four different perspectives, viz. governance and accountability; model and data; privacy; and IT general controls.

The risks and audit questions included in the audit framework are listed in Appendix 3. By answering all the questions in the audit framework and then allotting a score, we can show how the risks associated with the algorithm in question has been mitigated. The degree of risk associated with a particular algorithm depends in part on its impact on private citizens, among other factors. Appendix 1 explains in more detail how we assessed the algorithms.

This chapter describes the findings of our assessments of the nine algorithms. The main findings are summarised in the following figure, which shows the extent of the residual risk for a number of key aspects of each perspective. The lower row at the bottom of the figure shows whether the algorithms meet the criteria set out in the audit framework. An algorithm is deemed to be compliant if measures have been taken for all key aspects to mitigate the risks affecting the algorithm. An algorithm is not fully compliant if there is a high residual risk in respect of one or more key aspects.

Our findings in relation to each perspective are described in greater detail later on in this chapter (see sections 4.3 to 4.7).

Six of the nine algorithms do not meet the requirements set out in the audit framework

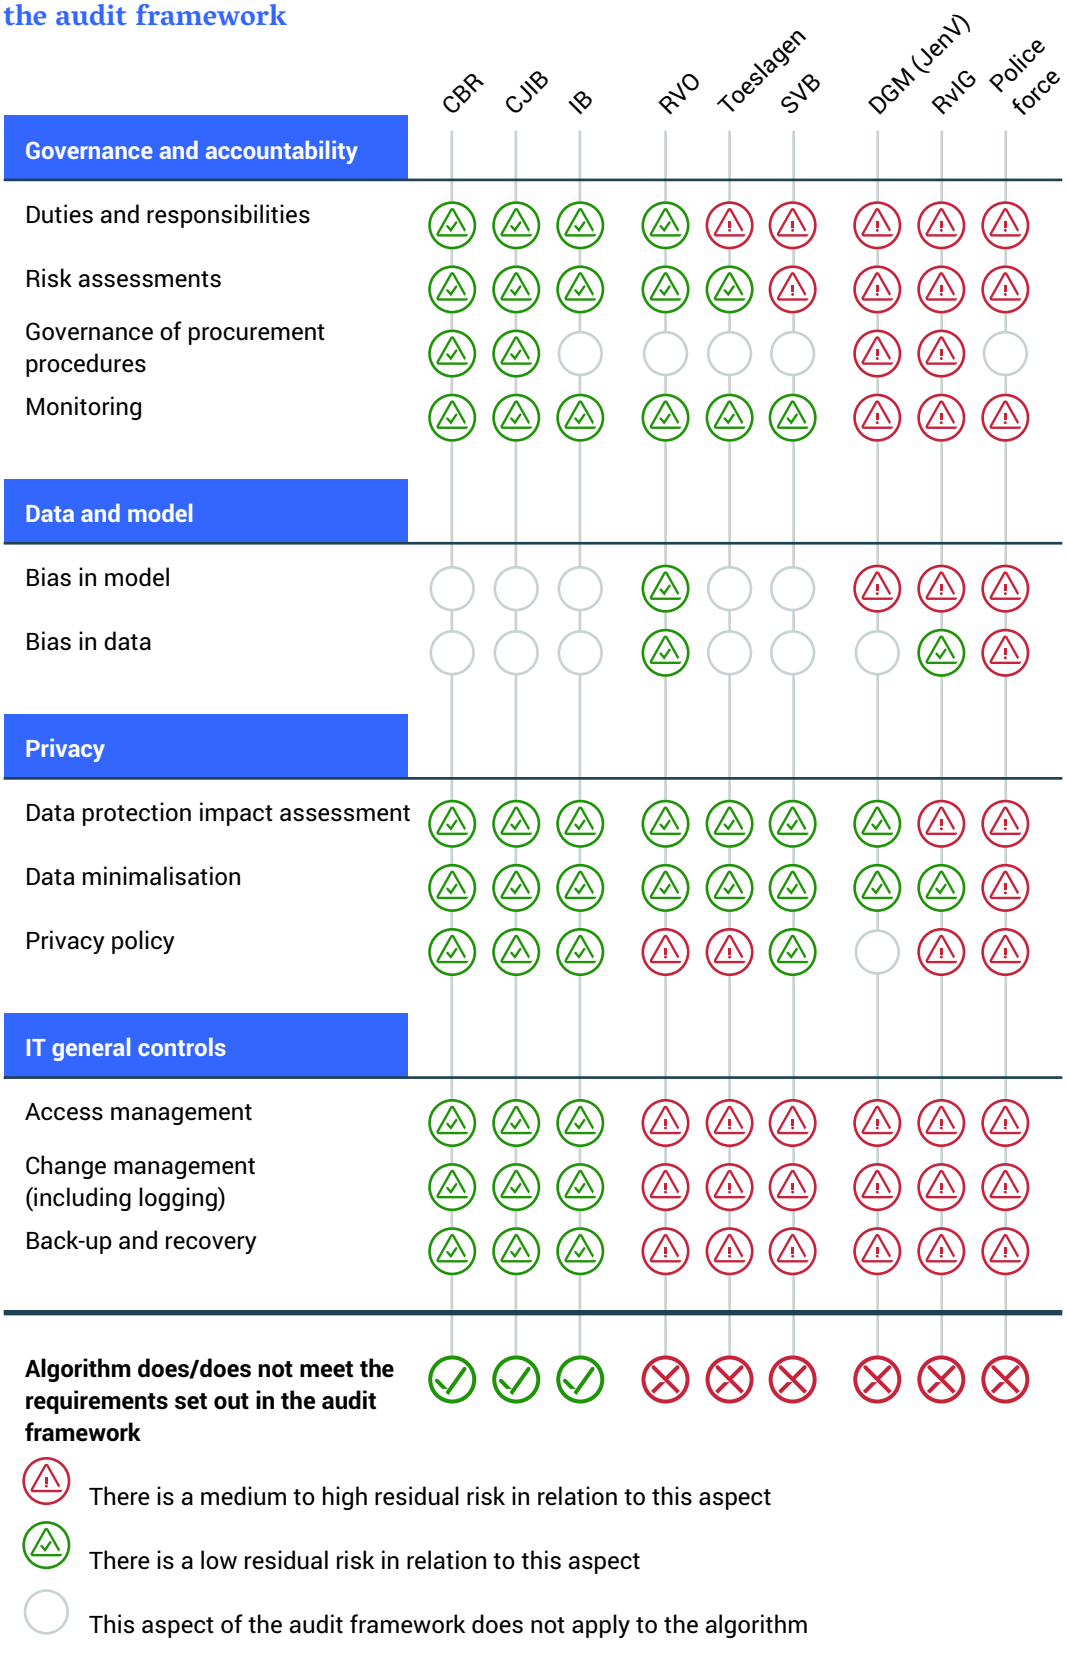

4.1 Three of the nine algorithms we assessed comply with the audit framework

Three of the algorithms we assessed complied with the audit framework. This means that the organisations in question have taken action to mitigate the risks associated with the algorithm as specified in the audit framework. This proves that it is indeed possible to use algorithms in a responsible manner.

These three algorithms are currently in use by the Central Office for Motor Vehicle Driver Testing (CBR), the Central Judicial Collection Agency (CJIB) and the Benefits Intelligence Agency (IB). The algorithms used by the Central Judicial Collection Agency and the Benefits Intelligence Agency that we assessed are electronic data interchange systems that are used to support a particular process. The algorithm used by the Central Office for Motor Vehicle Driver Testing that we assessed is a decision tree involving a form of partially automated decision-making that directly affects private citizens.

The Central Judicial Collection Agency, for example, uses an algorithm to ensure that traffic fines get to the right people. The algorithm links contact details (i.e. name and address) to traffic offences that have been linked to specific vehicle registration numbers, using data provided by the National Vehicle and Driving Licence Registration Authority (RDW). Our audit team found that the Central Judicial Collection Agency had made certain arrangements with the latter Authority about the quality of its services, and that these had been recorded in the form of a service level agreement. We also found that regular checks were made of the availability and supply of data, as well as of any interruptions in the service. This enables the Central Judicial Collection Agency to ensure that traffic fines reach the right people and do so in an efficient manner.

4.2 Six of the nine algorithms we assessed do not comply with the audit framework

Six of the algorithms we assessed are not fully compliant with the audit framework. These are used by the Netherlands Enterprise Agency, the Benefits department of the Tax and Customs Administration, the police force, the Directorate-General for Migration at the Ministry of Justice and Security (DGM), the Social Insurance Bank (SVB) and the National Office for Identity Data (RvIG). These algorithms have a high residual risk in relation to one or more key aspects of the audit framework.

In the case of three of the above organisations (viz. the Netherlands Enterprise Agency, the Benefits department of the Tax and Customs Administration, and the Social Insurance Bank), the risks in question relate primarily to IT general controls. These organisations still need to adopt demonstrably effective IT controls for the algorithms in question (see also section 4.6). They do not have a clear picture, for example, of whether only authorised members of staff actually have access to the data and the algorithm, or are authorised to make changes to the algorithm. This means that there is a relatively high risk of other members of staff gaining access to the algorithm and the data used by it. If people are able to gain unauthorised access to the systems, this may for example result in data being altered, damaged or lost – with potentially dire consequences for private citizens and businesses.

We found that the three other organisations (viz. the police force, the Directorate-General for Migration at the Ministry of Justice and Security and the National Office for Identity Data) are insufficiently able to manage the risks pertaining to a number of aspects of the audit framework. We identified the following residual risks, among others: inadequate arrangements and monitoring procedure in relation to outsourcing, no checks of bias, and inadequate IT controls. The main risks are discussed in more detail in section 4.3 to 4.6 of this chapter.

There are also disadvantages in not using algorithms

Two of the algorithms that we had originally selected for the purpose of this audit were no longer in use by the time that the audit got underway. The algorithms in question were a risk classification model used by the Benefits department of the Tax and Customs Administration and a risk scan used by the Social Insurance Bank in connection with personal care budgets.

The risk classification model used by the Benefits department of the Tax and Customs Administration

The risk classification model was adopted by the Benefits department in April 2013 as a means of selecting benefit applications for manual processing. It was decided in July 2020 that the department should stop using the model, as it did not meet the requirements of the General Data Protection Regulation (Ministry of Finance, 2020). No data protection impact assessment had been carried out that properly described the privacy risks and the associated measures for reducing or eliminating the adverse effects (Ministry of Finance, 2021). The Dutch Data Protection Authority found that the department had acted contrary to the terms of the Regulation in using nationality as a characteristic

in the algorithm, and decided that such action was discriminatory (Dutch Data Protection Authority, 2020).

However, the decision to stop using the model does have certain adverse consequences for the performance of regularity checks. When the risk model was still in use, those applications with the highest risk of errors were selected for manual processing. The number of applications needing checking was based on the available staff capacity. In other words, if not many staff were available, the Benefits department passed on only the highest-risk applications for manual checking by assessors (Ministry of Finance, 2021). Now that the algorithm is no longer used, more staff are needed to check the applications.

The risk scan used by the Social Insurance Bank in connection with personal care budgets

In 2019, the Social Insurance Bank launched a 'risk scan' as a pilot project for reducing the incidence of improper use and fraud in relation to personal care budgets. The risk scan provided the staff of local authorities with information about the potential risks of misuse and improper use in relation to personal care budgets. Despite the fact that the results of the pilot project were described as extremely promising, the Minister of Health, Welfare and Sport nonetheless decided to curtail the project. The problem was that the law needed to be changed in order to accommodate a national roll-out of the risk scan – and
the performance of risk analyses on a structural basis (Ministry of Health, Welfare and Sport, 2020). The Minister of Health, Welfare and Sport said at the time that he would probably decide by the end of 2020 whether or not to initiate a legislative procedure for creating a legal basis for the national roll-out of the risk scan (Ministry of Health, Welfare and Sport, 2020). The situation at the beginning of 2022 is that no decision has yet been taken.

The fact that the risk scan cannot be used has certain operational consequences. Those responsible for distributing personal care budgets do not currently have access to information and warning signs such as would enable them to respond adequately to evidence that the regulations on personal care budgets have been misused. The current mode of operation is based largely on manual processing, which absorbs a great deal of staff capacity.

4.3 Governance and accountability

The ‘governance and accountability’ aspect refers to the presence of written records of roles, responsibilities and expertise, the performance of risk assessments in relation to the use of the algorithm, and the making of arrangements with external parties about issues such as liability.

Purpose, risk assessment and monitoring

While many of the organisations in question have defined a clear purpose for the algorithm used by them, they have been less diligent in terms of regularly reassessing the risks associated with the use of the algorithm. In order to use an algorithm in a responsible manner, the owner must weigh up the advantages created by the algorithm against the risks posed by its use. And the results of this analysis may well change over time. Greater priority also needs to be given to regular assessments of whether the algorithm still satisfies the objectives or quality criteria that have been set for it.

Governance in relation to outsourcing

Ministers or executive agencies are not relieved of their responsibility for algorithms by dint of outsourcing them to external parties. They are obliged to keep an eye on the algorithms used by their organisations and on the risks associated with their use. Particularly where outsourcing is used or where the government buys an algorithm from an external supplier, it is important to make clear arrangements about who is liable for what. The development and management of two of the algorithms we assessed had been outsourced. The algorithms in question were those used by the Directorate-General for Migration at the Ministry of Justice and Security and the National Office for Identity Data. These two organisations had not adopted any controls in relation to governance and accountability. This was the case, for example, with the algorithm used by the National Office for Identity Data for assessing the quality of photographs used for identity documents, based on certain characteristics. Due to the fact that the National Office has outsourced the development and maintenance of the algorithm to an external supplier, the algorithm has become a ‘black box’ for the National Office. This is because the contract with the supplier does not contain any clear agreements about the algorithm, thus making it impossible for the National Office to check whether the algorithm is still doing what it is supposed to do.

4.4 Model and data

'Model and data' refers to aspects such as the development and maintenance of an algorithm. Checks to ensure that the algorithm is working properly form a key aspect of this perspective.

Checking the completeness of data processing

Completeness is a prerequisite for all forms of data processing. Private citizens (or other entities) should not be unintentionally excluded from data processing, as this may result in incorrect outcomes. An organisation can check whether this is the case by comparing the input for a particular algorithm with the output. We found that these checks are not always performed, however. In many cases involving the processing of large volumes of data, the owner relies on errors and omissions being flagged up in the form of error messages. The assumption is then that the absence of any error messages is in itself a guarantee that the algorithm is working properly. This is not always the case, though.

Checking the accuracy of data

Checks of the accuracy of the data used by the model are the next logical step after completeness checks. Certain data may not be available, or there may not be adequate guarantees that the data required by the model will be supplied on time. If these checks are not performed, there is no guarantee that the algorithm is working properly.

Checking for bias

In order to be sure that the algorithm is working properly, the owner needs to check whether it contains an undesirable bias in relation to certain individuals or groups of people. This requires a study of the causes of potential bias, and of how to deal with it should it arise, to be incorporated as a standard feature of the risk management process. We found that the organisations included in our audit do not all, as a matter of routine, employ staff with the relevant expertise in this respect. We also found that, in most cases, no checks are made to identify the presence and consequences of bias and hence the inaccurate overrepresentation or underrepresentation of certain groups of people. As a result, we are unable to rule out the possibility that certain algorithms may contain undesirable, systematic variations in relation to certain individuals or groups of people (see section 4.7).

What is bias?

Bias is taken to mean the presence of undesirable, systematic variations in the results of an algorithm in relation either to a group as a whole or to specific individuals or clusters of individuals.

How does bias arise?

The audit framework distinguishes between bias in the model and bias in the data. The model may contain a bias if, for example, it uses ethnicity or other characteristics that may be pointers to ethnicity (such as semi-literacy or attendance of a mosque). An example of bias in data is an algorithm that uses historic data about the application of a specific policy to specific groups of people. Let's say that the government mounted a big campaign in the past for detecting cases of benefit fraud stemming from people living together who had been classified as single for tax purposes ('cohabitation fraud'), and that the data generated by this campaign was then used to design an algorithm for detecting benefit fraud. In that case, the algorithm will be better at detecting cohabitation fraud simply because this type of fraud is relatively prevalent in the data. And if cohabitation fraud is committed predominantly by women, the algorithm will be biased towards women. This type of bias will then lead to more intensive checks being performed on female benefit claimants, and may hence result in women accounting for a higher proportion of offenders. This is a form of discrimination.

Is it possible to prevent bias?

Technically, it is not usually possible to completely rule out bias. What is possible, though, is to uncover instances of bias by analysing the results of an algorithm. And the way in which an algorithm is used also creates opportunities for making adjustments to account for bias. This can be done by incorporating certain checks in the algorithm or by checking the results to see whether they contain any undesirable systematic deviations.

Evaluating an algorithm

During the process of developing an algorithm, the quality of the model used needs to be evaluated by testing the predictions made by the algorithm, for example by comparing them with previous results. We found that a number of the algorithms we audited are not evaluated along these lines, even though there is scope for doing so. In certain cases, an evaluation is performed, but no proper record is kept of the results. This applies, for example, to the procedure for assessing grant applications under the TVL

scheme. Although data from previous grant award procedures are incorporated in new versions of the algorithm, no analysis is made of the differences between the data on grant applications and the data on grant awards. This means that no answer is provided to the question of whether the applications in question were rightly rejected or rightly approved.

4.5 Privacy

All the algorithms included in our audit make use of personal data. This imposes certain responsibilities on the organisation processing the data. After all, both private citizens and businesses must feel confident that the government handles their data in a responsible manner. Our audit team found that most of the organisations are aware of the privacy risks: they compile processing registers, for example, describing the type of personal data that they collect and process.

Data protection impact assessments

We found that the protection of personal data has become an integral part of the decision-making procedure on policies and regulations. Under the GDPR, government bodies are in certain instances obliged to undertake a data protection impact assessment (DPIA) in relation to new forms of data processing (Ministry of the Interior and Kingdom Relations, 2017). This type of assessment is mandatory, for example, for existing processes associated with substantial privacy risks. The government-wide template for DPIAs helps organisations to keep a record of their decision-making processes in this connection. We found that a DPIA had been made for seven of the nine algorithms we audited (or, as the case may be, for the processes of which the algorithms form part). Neither the National Office for Identity Data nor the police force have performed a DPIA. This means that they do not have a way of checking whether all risks have been mitigated.

Data minimalisation

Data minimalisation means that an organisation is not allowed to collect more data than are needed in order to achieve the intended purpose. We found that most organisations have taken action to prevent the disproportionate use or collection of data. The police force was the only organisation that collected more data than were needed to achieve the intended purpose. The police collected data on nationality, for example, without using the data for performing calculations.

Transparency

The extent to which information on privacy and data processing is shared varies from one organisation to another. While most organisations have published a privacy policy, this does not generally contain any information on specific algorithms. The Central Office for Motor Vehicle Driver Testing, the Benefits Intelligence Agency and the Social Insurance Bank have all published privacy policies with information on the way in which data is processed in relation to a number of specific algorithms. The Benefits Intelligence Agency's website, for example, includes a page on which its general privacy policy is set out (https://www.inlichtingenbureau.nl/Privacy-en-Veiligheid/Privacy-en-burgers). Information on the algorithms it uses may be found in a number of processing registers. These contain, for example, information on the purpose of processing, the sources used and the relevant retention periods. Benefit claimants and other interested parties can use this information to find out how the Agency uses personal data and what their rights are in this connection. The Netherlands Enterprise Agency and the Benefits department of the Tax and Customs Administration have both published general privacy policies, as well as a privacy policy designed specifically with algorithms in mind, which is incorporated in the application procedure.

4.6 IT general controls

One of the key preconditions for a responsible use of algorithms is that both the data used and the algorithm itself should be protected. Any unauthorised access to systems, for example, may result in data being altered, damaged or lost. This could potentially have drastic consequences for both private citizens and businesses.

IT general controls: ensuring that the basics are in order

Effective IT general controls – in relation to algorithms as well as in other contexts – are crucial for safeguarding the integrity, confidentiality and continuity of data systems. This means, for example, that only authorised persons should have access to the systems, that members of staff should not be given wider access rights than is necessary, and that all changes should be tested before being implemented in production.

What exactly are IT general controls?

IT general controls (ITGC) are basic measures taken by organisations in relation to IT systems. The audit framework for algorithms specifies certain

requirements that the following traditional IT maintenance processes are in any event expected to meet: access management (i.e. password management, authorisation management and network component security), change management, and back-up and recovery. An assessment of these controls traditionally forms part of our annual accountability audits.

IT controls for algorithms: some examples

- Where algorithms are used for taking decisions that affect private citizens (such as assessing applications for state pensions), the organisations concerned are required to ensure that only a limited number of staff are able to change the code used for the algorithm (and hence alter the decision-making procedure).
- Changes should be made with the aid of a predefined procedure that applies throughout the organisation, thus preventing any major changes in the algorithm from being made on the spur of the moment. It must be possible to reverse a major change in the algorithm if, for example, it becomes clear that it is detrimental to private citizens.
- Precisely because the majority of the algorithms that we assessed are capable of having a big impact on private citizens, it is important to keep a record of all aspects of the algorithm and its computer code. This allows the owner at all times to identify who carried out what type of work on the algorithm. As a result, the owner can trace the person responsible for making an error, as well as anyone who has misused his or her authority to access the algorithm. The organisation and the member of staff concerned can then rectify the error or learn from it.

How can an organisation make clear that it has a clear picture of the risks associated with the algorithms it uses?

Our aim is not just to understand exactly what the algorithms do, but also whether organisations are aware of the risks pertaining to the algorithms and are in control of them. This means that it is important for organisations to be able to prove that they have taken action to mitigate the IT risks or otherwise that they have very consciously decided to accept the risks associated with the algorithms. Showing that IT controls apply to algorithms (or to the environment in which algorithms operate) is one way of satisfying this requirement. IT audits can help in this respect.

IT controls of algorithms are inadequate

The results of our audit make clear that these general IT controls are inadequate in many respects in relation to six of the nine algorithms we audited. We wish to point out that six of the organisations included in our audit do not know whether access to the data and the algorithm is restricted to authorised members of staff, or whether only authorised members of staff are entitled to make changes to the algorithm. These IT governance problems are not unique to algorithms. For many years now, our audits have identified problems affecting IT management in central government (Netherlands Court of Audit, 2020). The picture emerging from this audit only serves to confirm this.

On the other hand, we also found that a number of organisations are in effective control of the risks associated with access management, change management and back-up and recovery in relation to the algorithm. The organisations in question are the Central Office for Motor Vehicle Driver Testing, the Central Judicial Collection Agency and the Benefits Intelligence Agency.

IT outsourcing

In those cases in which the development or management of an algorithm has been outsourced to an external supplier, the minister or executive agency in question remains responsible for ensuring that the algorithm is working properly. This was the situation at the National Office for Identity Data and the Directorate-General for Migration at the Ministry of Justice and Security. Our auditors found that the organisations in question do not know whether the risks associated with the algorithms are properly managed. For example, the organisations did not supply us with any information about the IT general controls, or were unable to show that the IT controls outsourced to external suppliers were effective.

4.7 Ethics

Ethics form an integral part of the four aspects we have already described, i.e. governance and accountability, model and data, privacy and IT general controls. The ethical perspective may be subdivided into the following four sub-topics:

1. respect for human autonomy;
2. the prevention of damage;
3. fairness;
4. explicability and transparency.

Respect for human autonomy

In the case of four of the algorithms included in the audit, requests received from private citizens or businesses are sent directly to the algorithm. This applies to the algorithms used by the Netherlands Enterprise Agency, the Central Office for Motor Vehicle Driver Testing, the Social Insurance Bank and the Benefits department of the Tax and Customs Administration. All four organisations use an algorithm for performing an automatic decision-making process that has a direct impact on private citizens or businesses. If an application meets all the criteria for the award of the grant or document in question, the application then follows a fully automated process in which the final decision is taken by an algorithm. If the application does not meet all the criteria, it is then subjected to a manual check by a member of staff.

We found that, where the decision-making algorithms included in our audit are used, the applicant is told about the reasons underlying the decision, for example in the form of an explanatory letter. Where the applicant is a private citizen, he or she is then entitled to lodge a formal objection to the decision.

The prevention of damage

As far as the prevention of damage is concerned, there is a need to safeguard people's privacy and to ensure that personal data are handled with care. We found that organisations spend a great deal of time and energy on the protection of personal data (see section 4.4). However, in order to prevent any damage, it is also important to protect both the data used and the algorithm itself. In other words, the general IT controls need to be in order. The risks are considerable: any unauthorised access to systems may result in data being altered, damaged or lost. This could potentially have drastic consequences for businesses and private citizens. The results of our audit make clear that general IT controls are inadequate (see section 4.6).

Fairness

Fairness means that an algorithm takes account of diversity in the population and does not discriminate. One should be aware at the same time that people themselves may also discriminate. For example, in an audit of the method of fraud detection used by the Tax and Customs Administration, PwC studied communications among members of staff and found instances in which the risk of fraud was based on nationality and outward appearance (Ministry of Finance, 2022, and PwC, 2021). A study by the Netherlands Institute for Human Rights showed that the use of algorithms may not only raise the risk of discrimination, but also lessen it (Netherlands Institute for Human Rights, 2020).

The following section looks at two ways in which the use of an algorithm may raise

the risk of discrimination, viz. the presence of bias in the algorithm and the presence of bias in the data.

Bias in the algorithm

Bias may be built into the model if, for example, it makes use of a variable such as nationality. The best example of this is the use of nationality in the risk classification model used by the Benefits department of the Tax and Customs Administration during the period up to July 2020 (Ministry of Finance, 2020). The Dutch Data Protection Authority came to the conclusion that the algorithm's use of nationality constituted an infringement of the regulation and was discriminatory (Dutch Data Protection Authority, 2020).

However, this does not imply that variables such as nationality or characteristics based on nationality may not be used in any circumstances. In certain instances, this information is needed in order to check whether a person is entitled to an allowance or benefit. In the case of two organisations included in our audit, this information was needed in order to check whether a person was entitled to an allowance or benefit. The Social Insurance Bank bases the value of a state pension on the number of years that the recipient has spent living or working in the Netherlands. For its part, when assessing applications for housing benefits, the Benefits department of the Tax and Customs Administration looks at whether the applicant lives in the Netherlands and either has Dutch nationality or is in possession of a residence permit.

Six of the organisations included in our audit (viz. the Central Office for Motor Vehicle Driver Testing, the Central Judicial Collection Agency, the Benefits Intelligence Agency, the Directorate-General for Migration at the Ministry of Justice and Security, the Netherlands Enterprise Agency and the police force) have not incorporated any variables in their algorithms that may lead to discrimination, such as nationality or variables based on nationality. In the case of one organisation, we do know which specific variables are used in its algorithm. The National Office for Identity Data has outsourced the algorithm to an external supplier, and is not aware of the variables that are used for checking passport photographs. In other words, the algorithm in question has become a black box for the National Office. This means that there is a risk of bias or discrimination if, for example, the algorithm is more likely to reject passport photographs of people of a certain nationality than those of people of other nationalities.

Bias in data

An undesirable systematic variation in relation to specific individuals or groups of individuals may also arise if, for example, use is made of historic data containing a given bias (see section 4.4). As regards the organisations included in our audit, the police force produces forecasts of future crime levels based on probability calculations using historic data. If this historic data contains a bias (for example, because more intensive checks were carried out in particular neighbourhoods in the past), there is a considerable risk of undesirable systematic variations appearing in the forecasts.

This was also identified in a study among Dutch local authorities that was commissioned by the Netherlands Institute for Human Rights. All the local authorities included in the study said that a risk was posed by the practice of consistently working with historic data and that there could also be a bias in the data set (Hooghiemstra & Partners, 2021).

In order to ensure that an algorithm works properly, the owner has to check it for any undesirable bias in relation to specific individuals or groups of individuals. We cannot rule out the possibility that some algorithms nonetheless contain an undesirable systematic variation, i.e. bias, in relation to certain individuals or groups. This is because we found that many organisations do not check for the effects of bias (see section 4.4), even though this would be a useful means of countering the risk of an undesirable bias.

Explicability and transparency

There are two types of transparency in relation to algorithms: technical transparency and procedural transparency. Technical transparency means that the owner of an algorithm must be able to explain how it works. Procedural transparency means that the owner of an algorithm must be able to explain how it was designed and how it arrives at its results. The nature of the target group is a key factor in deciding on the degree of transparency that is required. A study performed by a research institute (PON & Telos) and commissioned by a consortium operating under the name of 'Publieke controle op algoritmes' ('Public scrutiny of algorithms') revealed that the general public are not particularly interested in information about the technical details. Rather, people want to be informed about privacy, the extent to which the algorithm is subject to human scrutiny, and the reason for using the algorithm (Het PON & Telos, 2021). On the other hand, technical transparency is an important consideration for those responsible for exercising scrutiny, as they need to work out how the algorithm works.

Our auditors found that, generally speaking, technical transparency was not a problem. Most algorithms are not black boxes for audit institutions such as the Netherlands Court of Audit. For the purpose of our audit, we were given access to the models used for the algorithms, and the owners were able to explain broadly speaking how the models worked.

We did find, however, that there is scope for improving certain aspects of the procedural transparency. Private citizens and businesses do not always have a clear picture of what data are used by which algorithms, what their basic mode of operation is, and what the impact their results have. While many organisations have published privacy policies, these do not always contain information on specific algorithms (see section 4.5). The creation of a register of algorithms could be a means of increasing the degree of transparency and ensuring that the general public is better informed. The city council in Amsterdam has already developed an algorithm register along these lines (see https://algoritmeregister.amsterdam.nl/). This lists the algorithms used by the local authority for delivering its public services.

5. Conclusions and recommendations

We examined nine algorithms for the purpose of this audit and found that three of them satisfy the criteria set out in the audit framework for algorithms. This shows that it is indeed possible to use algorithms in a responsible manner. The remaining six of the nine algorithms do not comply with all the criteria set out in our audit framework. Our audit team found that there is still plenty of room for improvement. The main aspects in need of improvement are outsourcing (i.e. keeping a record of the arrangements made and monitoring the status), IT general controls and bias. We have formulated a number of recommendations for making the necessary improvements.

We repeat here two of the recommendations made in our 2021 audit report, entitled 'Understanding algorithms' (Netherlands Court of Audit, 2021):

1. Particularly where algorithms are outsourced or bought from an external supplier, document agreements on their use and make effective arrangements for monitoring compliance on an ongoing basis.
2. Ensure that algorithms and the data required for their operation are protected by effective IT general controls.

These are supplemented by two new recommendations:

3. Check regularly – both during the design of algorithms and during their use – for the effect of bias, in order to prevent any undesirable systematic variations in relation to specific individuals or groups of individuals.
4. When supervising arm's-length institutions, take explicit account of how algorithms are used to perform public services.

5.1 Arrangements and monitoring in relation to outsourcing

Further to the findings of our initial audit, we would like to stress the need for making clear arrangements when outsourcing algorithms or buying them from an external supplier. We found that outsourced algorithms become black boxes, and that it is unclear which aspects of the data are taken into account in arriving at a decision.

First recommendation: Particularly where algorithms are outsourced or bought from an external supplier, document agreements on their use and make effective arrangements for monitoring compliance on an ongoing basis.
We urge the government to reach an agreement at an early stage, when buying or outsourcing algorithms, about the purpose of the algorithms and the assessment of the risks involved, to document these agreements and to adopt mechanisms for monitoring compliance with these agreements on an ongoing basis. This will give the government a clear picture of whether the algorithms are being used in accordance with their intended purpose.

5.2 IT general controls

One of the key preconditions for a responsible use of algorithms is that both the data used and the algorithm itself should be protected. In other words, the basic IT controls should be in order and should be geared specifically towards algorithms. The algorithms included in our audit make use of sensitive data, including personal data. Private citizens and businesses must feel confident that their data are safe. Our auditors found, however, that these basic IT controls are not sufficiently effective in many respects. They also found that the IT controls were too generic and not specifically geared towards the algorithm in question. On too many occasions, the owners referred us to a set of generic controls, despite the fact that an algorithm requires specific protection against changes or the manipulation of data. For many years now, our audits have identified problems affecting IT management in central government (Netherlands Court of Audit, 2020). The picture emerging from this audit serves to confirm this.

Second recommendation: Ensure that algorithms and the data required for their operation are protected by effective IT general controls.
We recommend the government to ensure that algorithms and the data required by them are protected by IT controls that are designed specifically with the algorithms in mind.

5.3 Bias

A well-functioning algorithm takes account of diversity in the population and does not discriminate. If no effective action is taken, an undesirable systematic variation (i.e. bias) may arise in relation to specific groups of individuals. Checks need to be performed in order to ensure that an algorithm is working properly, by comparing the results generated by the algorithm for different groups of people, so that any systematic variations can be identified. We found that no such checks are performed, even though they would form a useful means of countering the risk of an undesirable bias.

Third recommendation: Check regularly – both during the design of algorithms and during their use – for the effect of bias, thus preventing any undesirable systematic variations in relation to specific individuals or groups of individuals.
We recommend the government to ensure that checks are made for the effect of bias and the inaccurate overrepresentation or underrepresentation of certain groups of people, so that it becomes clear whether there is any undesirable bias in relation to specific individuals or groups of individuals. In order to ensure that algorithms are used in a responsible manner, such checks for bias should form a standard part of routine risk management processes.

5.4 Supervision of algorithms used by arm's-length institutions

Algorithms are used not only by central government itself, but also by organisations operating at arm's length from central government. The above recommendations apply not only to the use of algorithms by ministries, but also to their use by arm's-length institutions, i.e. 'legal persons with a statutory task' and 'autonomous administrative authorities'. In supervising such institutions, ministers should take explicit account of how they use algorithms for performing public services.

Fourth recommendation: When supervising arm's-length institutions, take explicit account of how algorithms are used to perform public services.
We urge the government, in supervising arm's-length institutions, to take explicit account of how algorithms are used for performing public services. The audit framework may prove a useful aid in this connection.

6. Minister's response and Court afterword

On 11 May 2022, we received a response to our draft report from the Minister for Digitalisation, who is responsible for coordinating all central government action on digitalisation. Her response is summarised in section $6.1 below. The full Dutch text of her response, including a number of appendices containing responses from the relevant organisations, has been posted on our website (www.courtofaudit.nl). The chapter concludes in section $6.2 with our own afterword.

6.1 Response of the Minister for Digitalisation

In her response, the Minister for Digitalisation ('the Minister') thanks the Court of Audit for undertaking the audit and for the added value it brings to enhancing the quality of government action. The Minister says that she will be using the report's findings and recommendations to continue her work in improving the standard of service delivery and policy implementation.

The Minister writes that our audit framework for algorithms can be used to build up a picture of the extent to which the risks associated with the use of an algorithm have been mitigated. The use of the framework also helps those responsible for developing algorithms and those planning how to supervise their use. The framework forms part of the further developments outlined by the Minister in the rest of the letter.

The Minister also writes that the Court of Audit reaches a number of tough conclusions in its audit report and that, for this reason, she will be working together with the ministries concerned to remedy the shortcomings. Because this is a time-consuming

exercise, the Minister has asked the ministries to instruct their specialist staff, i.e. the Chief Information Officer (CIO) and the Data Protection Officer (FG) or Chief Privacy Officer (CPO), to examine whether it would be wise to continue using the six algorithms categorised as either medium- or high-risk until the problems have been resolved.

On the basis of the assessments received from the staff in question, the Minister is sufficiently confident that, for the time being, the risks identified can be adequately controlled by the measures that the organisations have put in place to guarantee the continued safe use of the algorithms and the data required by them. The organisations' responses are enclosed as appendices to the Minister's letter (also posted – in Dutch – on www.courtofaudit.nl).

The Minister regards these responses as the start of a process of, firstly, raising awareness of the risks associated with the use of algorithms and, secondly, engaging in a dialogue on the way in which the use of guidelines and frameworks can supplement the controls already adopted by the organisations in question.

The Minister also states that the organisations have all now set to work with the report's findings and recommendations, with the initial aim of dealing with (and hence reducing) the residual risks. They will be seeking to learn from those organisations whose algorithms were found by the auditors to comply with the audit framework. The next step will be to examine the Court of Audit's findings in greater detail. The Minister will be asking for information, in relation both to each individual organisation and to the specific algorithm included in the Court's audit, on the audit framework and the way in which the organisation arrived at a decision about the algorithm, assessed the risks and checked for bias. This information will form the basis of a letter containing further information, which will be sent to both the Lower House of Parliament and the Court of Audit before the summer recess.

6.2 Court of Audit's afterword

The Minister's actions show not just that she takes our audit findings seriously, but also that she takes her responsibility for coordinating government action on digitalisation equally seriously.

We welcome the fact that the Minister writes that the organisations will be seeking to remedy the shortcomings. The Minister says that she will is planning to inform both us and the Lower House about possible further action to be taken in response to our audit findings.

It is clear from the responses submitted by a number of the organisations included in our audit that they do not agree with some or all of our findings. They would appear to diverge from the Minister's views in this respect. For this reason, we would like to stress that careful thought has gone into all the various aspects of our audit framework. They are all based on existing standards, criteria, guidelines and legislation. The general IT controls included in the framework, for example, are based on the international ISO/IEC 27002 standard and on the Dutch government's information security baseline (BIO).

We should also point out that we consider the responses from both the Directorate-General for Migration at the Ministry of Justice and Security and the police force as a matter of concern. In our opinion, the Minister of Justice and Security should ask these organisations to address the risks identified in our audit report, as they can have a very real impact on private citizens.

The use of algorithms can improve the government's operational efficiency. This, in turn, can have a beneficial effect on both businesses and private citizens. At the same time, both private citizens and businesses must feel confident that the government uses algorithms in a responsible manner. We hope that this audit report will contribute to the responsible use of algorithms. We will be following the progress made in relation to all our recommendations and will continue to pay attention, in our audit work in the coming years, to the responsible use of algorithms.

Appendices

Appendix 1 How we performed the audit

Audit questions

Our audit questions were as follows:

1. Does the central government make responsible use of the algorithms that we selected?
   a. Have sufficiently effective controls been put in place to mitigate the risks?
   b. Do the algorithms that we selected meet the criteria set out in our audit framework for algorithms?
2. How do the selected algorithms operate in practice? How does each algorithm fit in with the policy process as a whole?
   a. How does the government arrive at a decision on the use of the algorithm?
   b. What do officials do with the algorithm's output? On which basis are decisions taken?
   c. What impact does this have on private citizens?

Selection of cases

This audit builds on the stocktaking survey that we carried out for our first audit of algorithms. We supplemented our list of sources with sources including:

- a media analysis;
- answers to questions supplied by project managers of our accountability audit;
- the results of the evaluation presented in *Quickscan AI in publieke dienstverlening II* (Netherlands Organisation of Applied Scientific Research, 2021);

- the National Ombudsman;
- the Netherlands Institute for Human Rights.

Which algorithms did we audit?

We assessed the use of nine algorithms on the basis of the criteria set out in our audit framework for algorithms. These algorithms were selected by applying the following criteria:

- *Impact on private citizens or businesses*
  The algorithms – or processes in which the algorithm plays a part – we selected have a big impact on individuals or businesses.
- *Risk-centred*
  We audited those algorithms which we believe have the greatest risk of not being used in the correct way.
- *Different domains*
  We selected algorithms from a number of different domains, e.g. the social domain and the security domain.
- *In operation*
  The selected algorithms also had to be currently in operation. We excluded from our audit any algorithms that had been discontinued or that were still in a pilot stage.
- *Different types of algorithms*
  In addition to the above, we selected different types of algorithms, ranging from technically simple algorithms, such as decision trees and electronic data interchange systems, to technically more complex algorithms, such as image recognition systems and self-learning applications.

Audit framework for algorithms

We audited the use of algorithms based on the audit framework for algorithms that we drew up as part of our 2021 audit, entitled 'Understanding Algorithms' (Netherlands Court of Audit, 2021). The audit framework is available to the public at: https://english.rekenkamer.nl/publications/publications/2021/01/26/audit-framework-for-algorithms. The audit framework assesses algorithms from four different perspectives:

1. governance and accountability;
2. model and data;
3. privacy;
4. IT general controls.

The risks and audit questions included in the audit framework are listed in Appendix 3.

Opinion

We formed our opinion of the use of algorithms as follows:

1. We audited the effectiveness of all the controls included in our audit framework (audit question 1a), based on the documentation submitted to us and also on our interviews. A control may be classified as 'effective', 'partly effective' and 'ineffective'.
2. We classified the residual risk as low, medium, or high. The residual risk is always high if the controls are ineffective. The risk classification may be lowered to either medium or low on the basis of the context and/or other supplementary measures.
3. We then formed our final opinion, i.e. the use of the algorithm does or does not comply with the requirements set out in our audit framework. This is an over-arching opinion.

Audit of the effectiveness of controls

We audited the effectiveness of controls based on documentation, interviews and observations during interviews. We asked the audited organisations to explain and show (by presenting the relevant documentation) how they mitigated each risk listed in the audit framework. We then assessed this information and asked the audited organisations to confirm the outcome of our assessment. Based on the responses and any supplementary documentation provided, we then re-assessed the effectiveness of the controls and completed the details in the audit framework. The initial assessments were always made by at least two auditors. These were followed by an overall assessment of all the algorithms performed in conjunction with the whole of our audit team. We also verified the internal and external consistency of our assessment.

Appendix 2 References

Dutch Data Protection Authority (2020). *Belastingdienst/Toeslagen. De verwerking van de nationaliteit van aanvragers van kinderopvangtoeslag*. The Hague.

European Commission (2019). *Ethics guidelines for trustworthy AI*.

European Commission (2020). *White Paper on Artificial Intelligence – A European approach to excellence and trust*.

Het PON & Telos (2021). *Informatiebehoeften van burgers over de inzet van algoritmes door overheden*. Tilburg.

Hooghiemstra & Partners (2021). *Hoe gemeenten besluiten over algoritmen & mensenrechten. Onderzoek voor het College voor de Rechten van de Mens*. The Hague.

Lower House of the Dutch Parliament (2021). *Motie van het lid Klaver c.s. Parlementaire ondervraging kinderopvangtoeslag.* Lower House of the Dutch Parliament, 2020–2021 session, 35 510, no. 16.

Ministry of Finance (2020). Letter from the State Secretaries for Finance. *Nadere informatie over de Fraude Signalering Voorziening (FSV) en het gebruik van FSV binnen de Belastingdienst*. Lower House of the Dutch Parliament, 2019-2010 session, 31 066, no. 681.

Ministry of Finance (2021). Letter from the State Secretary for Finance – Benefits and Customs. *Openbaarmaking Risicoclassificatiemodel Toeslagen*. Lower House of the Dutch Parliament, 2021-2022 session, 31 066, no. 923.

Ministry of Finance (2022). Letter from the State Secretary for Finance – Taxation and Tax and Customs Administration. *Rapporten PwC over FSV – Particulieren en externe gegevensdeling*. Lower House of the Dutch Parliament, 2021-2022 session, 31 066, no. 957.

Ministry of Health, Welfare and Sport (2020). Letter from the Minister of Health, Welfare and Sport. *Persoonsgebonden Budgetten*. Lower House of the Dutch Parliament, 2019-2020 session, 25 657, no. 332.

Ministry of the Interior and Kingdom Relations (2017). *Model gegevensbeschermings-effectbeoordeling Rijksdienst (PIA). Report no. 105172*. The Hague.

Ministry of Justice and Security (2019). *Bijlage bij Brief over waarborgen tegen risico's van data-analyses door de overheid*. Lower House of the Dutch Parliament, 2019-2022 session, 26 643, no. 641.

National Ombudsman (2021). *Een burger is geen dataset. Ombudsvisie op behoorlijk gebruik van data en algoritmen door de overheid*. Report no. 2021/021. The Hague.

Netherlands Court of Audit (2020). *State of Central Government Accounts 2019*. The Hague: own publication.

Netherlands Court of Audit (2021). *Understanding algorithms*. The Hague: own publication.

Netherlands Institute for Human Rights (2020). *Als computers je CV beoordelen, wie beoordeelt dan de computers? Algoritmes en discriminatie bij werving en selectie*. Utrecht.

Netherlands Organisation for Application Scientific Research (2021). *Quickscan AI in publieke dienstverlening II*. Commissioned by the Ministry of the Interior and Kingdom Relations. The Hague.

PwC (2021). *Onderzoek effecten FSV Particulieren*. Amsterdam.

Rotterdam Court of Audit (2021). *Gekleurde technologie. Verkenning ethisch gebruik algoritme*s. Rotterdam.

Trouw (2021, 15 April). *De Rekenkamer waarschuwt Rotterdam: groot risico op discriminatie en profilering door algoritmes*. Retrieved on 28 February 2022 from https://www.trouw.nl/binnenland/de-rekenkamer-waarschuwt-rotterdam-groot-risico-op-discriminatie-en-profilering-door-algoritmes~b58eb5b9/.

De Volkskrant (2021, 23 April). *Staat licht in jacht op uitkeringsfraude burgers volledig door tot verbazing van privacy experts*. Retrieved on 28 February 2022 from https://www.volkskrant.nl/nieuws-achtergrond/staat-licht-in-jacht-op-uitkerings-fraude-burgers-volledig-door-tot-verbazing-van-privacy-experts~b41a35c6/.

Appendix 3 Audit framework for algorithms

We developed our audit framework for algorithms for the purpose of our 2021 audit, entitled 'Understanding algorithms' (Netherlands Court of Audit, 2021). It is a practical tool that the central government can use to manage the main risks associated with the use of algorithms. In developing our audit framework, we made use of existing standards, criteria, guidelines and legislation.

General questions

Before using the audit framework, the user must first answer a number of general questions. The answers to these questions paint a general picture of the algorithm and its context. It is this context and general impression that determine which questions should be selected from the audit framework for the purpose of assessing the algorithm in question.

1. What is the name of the algorithm or the system of which the algorithm forms part?
2. What is the operating process or the product or service in which the algorithm is used?
3. Does the algorithm make use of personal data (GDPR)?
4. Is it a self-learning algorithm, i.e. an algorithm that evolves and improves over time by making use of data and/or experiences?
5. Does the algorithm support or provide advice in relation to human activities or decisions, or does it act independently or automatically without any human intervention?
6. What technology and/or what applications or software does the algorithm use?
7. What data and data sources does the algorithm use?

Five perspectives

The audit framework consists of five perspectives:

1. governance and accountability;
2. model and data;
3. privacy;
4. IT general controls;
5. ethics.

We identified the main risks relating to each of the above perspectives, and linked the elements of the assessment and our audit questions to these risks. Once all the various questions have been answered and scores allotted, a picture emerges of the extent

to which the risks relating to a particular algorithm have been mitigated. The degree of risk associated with a specific algorithm depends (among other factors) on its impact on private citizens.

Ethics

The questions in our assessment framework are based in part on ethical principles (European Commission, 2019; European Commission, 2020). There are four key issues in connection with the ethical perspective:

1. respect for human autonomy;
2. the prevention of damage;
3. fairness;
4. explicability and transparency.

The numbers refer to an ethical principle. The ethical principles are linked to the risks associated with the other four perspectives, i.e. governance and accountability, model and data, privacy, and IT general controls.

No.	Ethical framework	Ethical principle
1.1	Respect for human autonomy	The decisions made by the algorithm are open to human checks.
2.1	The prevention of damage	The algorithm is safe and always does what it is supposed to do.
2.2		Privacy is safeguarded and data protected.
3.1	Fairness	The algorithm takes account of diversity in the population and does not discriminate.
3.2		The algorithm's impact on society and the environment was taken into account during its development.
4.1	Explicability and transparency	It is possible to explain the procedures that have been followed.
4.2		It is possible to explain how the algorithm works.

The audit framework for algorithms

The following is a list of the risks and audit questions included in our audit framework.

No.	Risk	Audit question	
1	Governance and accountability		
1.01	Management or accountability is not possible without clarity about the purpose of an algorithm	Has the purpose of the algorithm been clearly defined?	4.2
1.02	Without an up-to-date analysis of the risks, it is impossible to reach an informed decision as to whether the benefits of using the algorithm outweigh the drawbacks	Are regular assessments made at pre-determined intervals of the risks associated with the use of the algorithm?	4.1
1.03	Without adequate expertise in both qualitative and quantitative terms there is a greater risk of error	Does the organisation have access to sufficient expertise in both qualitative and quantitative terms?	
1.04	An incomplete picture of the algorithm's life cycle makes it more difficult to operate and manage	Has the entire life cycle management process for the algorithm been documented?	
1.05	Lack of clarity about roles, tasks, responsibilities and powers creates risks	Have roles, tasks, responsibilities and powers (including ownership) been defined and have these been applied in practice?	4.1
1.06	Performance and quality targets cannot be measured or discussed if no policy has been adopted in this respect	Is there an agreed and documented policy on quality and performance targets for algorithms?	4.2
1.07	A dependency on external experts who leave after developing the algorithm, taking their knowledge and experience with them, means that continuity and management are no longer safeguarded	Where certain elements or activities relating to the algorithm have been outsourced, have arrangements been made with external suppliers and have these been documented?	4.1
1.08	The algorithm cannot be managed without any monitoring	Is the algorithm monitored at regular intervals? Aspects that could be monitored include continuity, performance/quality, and compliance with current legislation.	

No.	Risk	Audit question	
2	**Model and data**		
2.01	The algorithm does not operate in accordance with the objectives set	Does the algorithm have a clearly formulated purpose and has this been translated into practical features with respect to the model and data used? Which particular task or aspect of operational management is the algorithm intended to support?	4.2
2.02	Without a shared picture of the objectives, there is a greater risk of errors and/or differences of interpretation	Does the algorithm have an agreed purpose, and is this clear and explicable to the owner, developer and user?	4.2
2.03	It is either impossible or difficult to explain how the algorithm is used	Is it possible to explain how the algorithm is used, and has an effort been made to strike a balance between the model's explicability and its performance?	4.2
2.04	It is no longer possible to trace the reasons underlying the choices made in designing and implementing the algorithm	Has a record been made of the reasons underlying the choices made in designing and implementing the algorithm?	4.1, 2.1
2.05	No continuity in the process or the performance of activities, due to the absence of documentation	Is there any documentation describing the design and implementation of the algorithm?	4.1
2.06	Hyper-parameters were selected at random, and the wrong choices were made in doing so. A hyper-parameter is a parameter or variable that can be used to manage a training or learning process.	Was the selection of hyper-parameters supported by arguments and evidence?	
2.07	A lack of transparency for private citizens, businesses and stakeholders; non-compliance with legislation on transparency	Has the model (i.e. the code and mode of operation) been published and is it available to interested parties? Have the data, or has a description of the data, been published and made available to interested parties?	
2.08	The algorithm uses automated decision-making even though this is not permitted, or there are no opportunities for human intervention	If the algorithm leads to automated decision-making, does it comply with the relevant legislation?	1.1, 2.1
2.09	Very limited sources of input mean a higher risk of error and a failure to meet the objectives set and to comply with the relevant legislation	Were the various stakeholders and end users of the algorithm involved in the development process?	3.1

No.	Risk	Audit question	
2.10	The algorithm does not operate according to plan	What input-output checks are performed to safeguard the accuracy and completeness of data processing?	2.1
2.11	The model was based on the legislation applying in year t-1, and is now being used in year t. The legislation (e.g. on margins of tolerance and limits) may have changed in the meantime, or certain legal provisions may no longer apply	Is the model updated at regular intervals to bring it into line with current legislation?	
2.12	Incorrect training or testing may lead to overfitting or underfitting, or bias (undesirable systematic variation).	Have safeguards been put in place to guarantee the quality of the choices made in relation to training and test data?	4.1
2.13	The model leads to undesirable systematic variation in relation to certain individuals, groups of people or other entities (bias)	Have guarantees been put in place to avoid any bias as a result of the choices made in relation to the model?	3.1, 3.2
2.14	There is an undesirable systematic variation (i.e. bias) in the data	Is there any undesirable bias in the data?	3.1, 3.2
2.15	If training, test and validation data are not processed separately, this leads to overfitting, which means that the model cannot be used for new observations	Are training, test and validation data processed separately?	
2.16	The data are not representative	Are the data representative for the application for which the algorithm is used?	2.1, 3.1, 4.1
2.17	Dependency on third parties with respect to the data used	Does the organisation or government body have full control (ownership) of the data used for the model?	
2.18	Violation of basic premises and rules on data minimalisation and proportionality	Is there evidence of data minimalisation? Have proportionality and subsidiarity been taken into account?	2.1
2.19	The performance metrics are not consistent with the algorithm's objectives	Has the quality of the model been documented?	4.2
2.20	The data on which the model is based are available only after the outcome has been identified	Is there evidence of target leakage? That is to say, do the model's features of the model include the forecasts that it is designed to make?	
2.21	The quality of forecasts is not up to standard	Are performance indicators or performance metrics used?	2.1, 4.2

No.	Risk	Audit question	
2.22	The model does not always work in practice, or no longer works in practice	Is the model's output monitored?	2.1
2.23	People do not know that they are dealing with an algorithm, and do not know about the consequences or the limitations of the algorithm. This may result in incidents or errors and hence claims for damages	Has the operation of the model or algorithm, including its limitations (i.e. what it can and cannot do), been communicated to external parties?	4.2
2.24	There is a risk that all efforts are concentrated on developing and producing the algorithm, and that no account is taken of the transfer of information and knowledge to the officials responsible for managing the algorithm, or of the need for maintenance to take account of business considerations	Have arrangements been made for the maintenance and management of the algorithm?	
3	**Privacy**		
3.01	Not compliant with statutory regulations under the GDPR for keeping the register up to date	Is the use of personal data recorded in a register?	2.2
3.02	The design of the algorithm does not take sufficient account of the need to protect personal privacy	Is there evidence of data protection by design?	2.2
3.03	Not compliant with statutory regulations under the GDPR for performing a data protection impact assessment (DPIA)	Has a DPIA been performed (if applicable)?	2.2
3.04	The algorithm uses automated decision-making even though this is not permitted under the GDPR .	Is there evidence of automated decision-making, and if so, is this permitted?	2.2
3.05	Not compliant with statutory regulations under the GDPR for serving mankind	Can those involved opt out of automated decision-making (if applicable)?	2.2
3.06	Disproportionate use or collection of personal data	Is there evidence of data minimalisation?	2.2
3.07	Unlawful action in relation to data processing.	Is data processing based on a statutory duty?	2.2
3.08	Not compliant with GDPR in relation to purpose limitation	Is the use of the algorithm to process (special-category) personal data consistent with its original purpose?	2.2

No.	Risk	Audit question	
3.09	Not compliant with statutory regulations under the GDPR in relation to recording of responsibilities	Have the data controller and the processor of personal data for the algorithm and the data used been identified?	2.2
3.10	Violation of Article 1 of the Constitution or Article 14 of the European Convention on Human Rights (ECHR)	Do the data used and the model lead to discrimination?	2.2
3.11	Profiling as defined in Article 4 (4) of the GDPR; risk of contravening the GDPR	Has an assessment been made of whether there is evidence of profiling and whether this is permitted?	2.2
3.12	Not compliant with statutory regulations under the GDPR in relation to informing data subjects	Have arrangements been made for informing, either pro-actively or on request, individuals whose data are processed or used (in relation to both data and the algorithm)?	2.2
3.13	Not compliant with statutory regulations under the GDPR and the general principles of good governance in relation to logic and accessibility	Is the logic behind the algorithm and the data used sufficiently clear to data subjects?	2.2
3.14	Not compliant with statutory regulations under the GDPR in relation to impact on data subjects	Is the impact of the algorithm clear to data subjects?	2.2
3.15	Data subjects have not been informed of their rights or of the algorithms and data used	Is there a publicly available privacy policy describing the data and algorithms?	2.2
4	**IT general controls**		
4.01	Without any logging information, there is no audit trail for tracing when adjustments were made	Is logging information about the operation of the algorithm recorded and stored in an accessible manner?	
4.02	Access rights are no longer up-to-date	Are any checks made of whether access rights are up-to-date with respect to the algorithm's operating environment?	2.2
4.03	Unlawful access to the algorithm	Are access rights updated when a member of staff leaves or moves to a different post?	2.2
4.04	Access rights are issued by unauthorised staff	Are access rights issued by staff who are authorised to do so?	2.2

No.	Risk	Audit question	
4.05	Risk of the algorithm being manipulated in cases where access rights are incompatible	Is there a mechanism for preventing individuals who are entitled to access the algorithm from playing a number of different roles at the same time?	2.2
4.06	The more users are granted special powers, the greater the risk of manipulation	Are management accounts generic? Is there a logical relationship between the number of management accounts and the number of managers?	2.2
4.07	User groups are difficult to identify	Are naming conventions used when granting access rights to different user groups or roles? Is this done on a systematic basis?	2.2
4.08	Managers and users are difficult to identify.	Are naming conventions used for users and managers, so that they can be identified?	2.2
4.09	It is not clear who made changes to or worked on the algorithm.	Do managers perform management and ordinary user activities under two different user names?	2.2
4.10	The database is open to manipulation if holders of user accounts have access to underlying components	Do user accounts have direct access to underlying components?	2.2
4.11	The database is open to manipulation in respect of the separation of duties if holders of user accounts have access to underlying components	Is there a separation of duties in relation to applying for, authorising and processing changes in user accounts and access rights?	2.2
4.12	The database is open to manipulation in relation to password management if holders of user accounts have access to underlying components	Are passwords managed interactively, and are they of adequate quality?	2.2
4.13	Unauthorised access, changes, damage to and/or loss of data. Non- compliance with the law	Are changes made to the code of the algorithm verifiable? Are changes tested, approved and authorised, for example?	2.2
4.14	Unauthorised access, thus creating a risk of the algorithm being manipulated (i.e. data being changed, damaged or lost)	Is the algorithm protected, so that there is no risk of unauthorised access, or of data being changed, damaged or lost?	2.2
4.15	Back-ups are not consistent with the back-up policy. There is no recovery option, and hence a risk of data loss, if the algorithm stops working	Are back-ups made of the algorithm, and can the algorithm and the data be restored?	
4.16	The lack of security by design creates certain risks	Is there evidence of security by design?	2.1

Appendix 4 Abbreviations and key terms

	Description
Algorithm	A set of rules and instructions that a computer follows automatically to solve a problem or answer a question
Bias	An undesirable systematic variation in relation to specific individuals or groups of people
CAS	Criminality Anticipation System
CBR	Central Office for Motor Vehicle Driver Testing
CJIB	Central Judicial Collection Agency
DPIA	Data protection impact assessment
IT general controls	Basic measures taken by organisations to protect and control IT systems: access security, change management, and back-up and recovery.
Model and data	The development and maintenance of an algorithm
Procedural transparency	The owner of an algorithm must be able to explain how it was designed and how it arrives at its results
RvIG	National Office for Identity Data
RVO	Netherlands Enterprise Agency
Governance and accountability	Written records of roles, responsibilities and expertise, the performance of risk assessments in relation to the use of the algorithm, and the making of arrangements with external parties about issues such as liability
SVB	Social Insurance Bank
Technically complex algorithms	Algorithms such as image recognition systems and self-learning algorithms
Technically simple algorithms	Algorithms such as decision trees, search engines and electronic data interchange systems
Technical transparency	The owner of an algorithm must be able to explain how it works
TVL	A government scheme for the reimbursement of fixed costs. The scheme is designed to help business owners confronted by a sharp decline in turnover due to the effects of government action to combat the COVID-19 pandemic

Appendix 5 Endnote

The draft version of this report stated that the National Office for Identity Data does not have a published privacy policy. This passage has been removed from the report because the National Office for Identity Data made clear, when we sent them the draft version, that they do indeed have a published privacy policy.

Netherlands Court of Audit

Department of Communication

PO Box 20015

2500 EA Den Haag

The Netherlands

Phone +31 70 342 44 00

voorlichting@rekenkamer.nl

www.courtofaudit.nl

Translator: Tony Parr

Cover photo: Shutterstock.

The Hague, May 2022