

Data science for tax administration

Pijnenburg, M.G.F.

### Citation

Pijnenburg, M. G. F. (2020, June 24). *Data science for tax administration*. Retrieved from https://hdl.handle.net/1887/123049

Version:	Publisher's Version
License:	<u>Licence agreement concerning inclusion of doctoral thesis in the</u> <u>Institutional Repository of the University of Leiden</u>
Downloaded from:	https://hdl.handle.net/1887/123049

Note: To cite this publication please use the final published version (if applicable).

Cover Page



# Universiteit Leiden



The handle <u>http://hdl.handle.net/1887/123049</u> holds various files of this Leiden University dissertation.

Author: Pijnenburg, M.G.F. Title: Data science for tax administration Issue Date: 2020-06-24

# Various Topics

Besides the four main topics treated in the previous chapters, we have touched four others as well. We have included these contributions in a separate chapter. We hope that by seeing these various topics, the reader gets an impression of the broadness of the application field of data science to the administration of taxes.

We start with an application of analytics to Human Resources. This application area is also known as *HR Analytics*. In the section, a model will be constructed to explain differences in time allocation of employees. Then we will look at the connection between reinforcement learning and tax collection. Reinforcement learning forms a large domain of machine learning, next to *supervised learning* and *unsupervised learning*. Reinforcement learning has attracted a lot of attention lately, when the combination of deep learning and reinforcement learning proved to be able to beat the world champion of the board game Go. In the section we look at its potential to improve choices in the the collection of tax debts. In the third section, the focus is on a way to explain risk models to non-experts. Explaining complex algorithms from machine learning is often essential to introduce these techniques in a tax administration. Finally, we look at some ideas from Fuzzy Sets, and apply these on tax data, in line with the approach of Meza et al. [76]. The application on real data gives evidence of the predictive power of this approach.

Some parts of the chapter have been published in the following articles:

- C. Boon, F. D. Belschak, D. N. Den Hartog, and M. Pijnenburg. Perceived human resource management practices. *Journal of Personnel Psychology*, 2014
- M. Pijnenburg and K. Kuijpers. Explaining risk models to the business. *Tax Tribune*, 35:57–62, 2016

Besides, the idea about applying reinforcement learning has been worked out in the following Master Thesis, supervised by the author,

• M. Post. Tax data and reinforcement learning. Master's thesis, Leiden University, 2 2019. Under supervision of Mark Pijnenburg, Wojtek Kowalczyk, and Kaifeng Yang.



Figure 7.1: Model for the relation between (perceived) HR practices and time allocation

## 7.1 Perceived Human Resources Practices and Time Allocation

In a company, a difference can be made between 'core processes' and 'support processes'. Core processes deliver the products or services that are sold by the company. Support processes assist these core processes. Examples of support processes are financial processes and human resource (HR) processes.

Analytics can be applied to support processes as well. Applications of analytics to HR are frequently called *HR Analytics*. As an illustration of HR Analytics at the Netherlands Tax and Customs Administration, we present the use of confirmatory factor analysis to verify a hypothesized model. This model establishes a relationship between the time spent by employees and the (perceived) quality of HR practises. A full description of this research can be found in [18].

In short, based on available knowledge in the HR literature, a relation between perceived HR practices and time allocation of employees is hypothesized, as depicted in Figure 7.1. At the right side of the figure, three 'HR bundles' are shown, *people flow, appraisal and reward*, and *employment relation*. These three bundles, that relate to different HR aspects, can be seen as the components of a High-Performance Work System (HPWS). The three HR bundles are expected to influence the time allocation of employees. As can be seen from the Figure (left side), the time allocation is split into three components: time spend on task activities, time spend on contextual



Figure 7.2: Result after applying confirmatory factor analysis

activities, and absence (mainly due to illness). Task activities are tasks mentioned in the job description and related to an internal or external client. Contextual activities are other activities like networking and training, that are necessary to guarantee future production. The influence of the HR bundles can be directly, or indirectly by the concepts of *extra effort* and *job satisfaction*.

The hypothesized relation can be tested by applying confirmatory factor analysis, a standard technique from statistics. Figure 7.2 shows the significance of the hypothesised relations after testing. Clearly, three hypothesised relations are not significant, while the others are. These insights may be used to optimize the HR practices at the NTCA or other organizations, in order to diminish unwanted absence or the increase employee satisfaction or time spent on task activities or contextual activities.

# 7.2 Reinforcement learning applied to tax debt collection

Reinforcement Learning is sometimes described as the third central pillar of learning from data, next to supervised learning and unsupervised learning [78]. *Supervised learning* is characterized by the presence of a sample of data for which we have additional information. Usually, the task of supervised learning is to learn the additional information for cases outside the sample. With *unsupervised learning* such additional information is lacking, and the task is usually to find structure in the data set, like the probability distribution of observations, or the location of clusters or anomalous observations.

*Reinforcement learning* is different from the previous two concepts, since there is no classical tabular data set with observations and features. Instead, in reinforcement learning, there is a player that is in a certain state (of an environment) who can take a limited number of actions. By taking an action, the state changes and there may occasionally follow a reward or punishment. By taking actions and observing the change of the states the player tries to understand the environment and to find a sequence of actions that will optimize his rewards.

Although the description of reinforcement learning may sound overwhelmingly at first, it is really close to human learning. Consider an example: when a child (the player) learns to walk, it is at a certain position (state) in a room (the environment). When it does some muscle movements (actions) it may change position and get some reward; a positive reward if it experiences that it has walked some distance, and a negative if it bumps against an object or falls. By reflecting on the muscle movements and the change in the state that followed these movements, the child slowly starts to understand how to interact with the environment it is in. Moreover, eventually it will learn what muscle movement will give good rewards: it has learned to walk! Another example is chess play; when a player has done an action (moved a piece), the state of the environment (the positions on the chessboard) has changed and at the end a positive (win) or negative (loss) reward occurs. By observing the effects of his actions, the player may learn how to play chess. A comprehensive introduction to reinforcement learning is the book of Sutton [105].

Reinforcement learning obtained some spectacular results recently. Supported by deep learning technologies, a reinforcement learning algorithm called AlphaGo was able to defeat the world champion 'Go' in March 2016. AlphaGo's successor AlphaZero is able to learn the games of chess, shogi, and go from scratch within a day and reach a superhuman level. In December 2017 AlphaZero defeated the best computer programs for these three games.

The debt collection process of a tax administration resembles the basic set-up of reinforcement learning: open tax debts may follow a sequence of actions taken by the tax administration (the player) starting with a reminder and usually ending with payment. Each action of the tax administration may lead to another action of the debtor, for instance a phone call of the tax administration may lead to a formal request for deferment from the debtor. As a result, a sequence of actions occurs, where each time the state of the debt may change. Rewards occur when the debt is (partially) paid, while a loss occurs (for the tax administration) when the debt is written off. Hence the idea of applying reinforcement learning to learn an optimal strategy for the tax administration to collect unsettled debts.

Actions that the tax administration can take include calling the taxpayer, offering a payment plan, initiating a pay claim, and allow a delay of payment. Currently, at the NTCA, the tax administration usually starts with mild measures that are getting harsher if it becomes clear that the taxpayer is unwilling to pay. If the taxpayer is unable to pay, the tax administration tries to take a more supportive position.

More on applying reinforcement learning on the collection process of a tax administration can be found in the master thesis of Martijn Post [95], written under the author's supervision. In summary, Post implemented several variants of a simple reinforcement algorithm to an anonymous data set of the NTCA (Dutch tax administration). The algorithm produces a value function for the current policy that gives sensible values for possible next actions. To improve on these initial results, more data of the taxpayers has to be added, like their ability to pay. Moreover, the algorithm must be able to explore the environment by taking some unconventional actions from time to time. Both these extensions were thought to be undesirable with a view on (the privacy of) the taxpayers and out of proportion. Especially since the current collection process can probably be improved with less far-reaching analyses. For these reasons we have not continued this line of research.

#### 7.3 Explaining risk models to non-experts

Risk models are becoming more popular in tax administrations for audit selection. Although risk models are not new to tax administrations, only recently they are becoming mainstream and are embedded in central business processes. The fact that an, often complex, computer algorithm is at the heart of a risk model makes it harder to explain to auditors and managers, who use or are responsible for the model. A good explanation is an integral part of getting these new, often efficient, techniques accepted.

Tax administrations are most familiar with risk rules, like "if this year's amount differs more than 20% from previous year's amount, then select for audit". Domain experts construct these rules. In contrast, risk models are created by applying a computer algorithm to historical data. Note that in practice, often, a combination of risk rules and risk models is applied.

An analogy can help explain a risk model. Below we present an analogy that has been tested by the Netherlands Tax and Customs Administration. The analogy explains as well some essential concepts such as 'Analytical Base Table' and 'target'.

To begin with, auditors and managers are asked to imagine that they are teachers in a class of students. As a teacher they are interested in knowing what students failed to complete their homework. Of course, they can check every student, but this will consume most of the lesson, leaving no time to teach. Just like auditors who only have limited capacity to control a large number of companies. The analogy is continued by assuming that the teacher decides to choose three students for each lesson to be checked for their homework. The teacher wants to check the three students with the largest chance of not having completed their homework. To select these students, the teacher uses different *features* of his students, see Figure 7.3. These features can be any characteristic of a student, such as sitting in the back of the classroom, gender, results of previous homework checks, and perceived kindness. In the context of taxation, features are usually characteristics of taxpayers or tax returns and we want to check the correctness of tax returns.



Figure 7.3: Explaining a risk model by analogy of a classroom.

Then it is explained that not all features are suitable for the selection of students. Some features might be in conflict with the way we want to treat students; it is usually not desired to discriminate between students (or taxpayers) based on features like race or gender. Moreover, some features cannot be determined objectively; a feature like 'annoying' is not objective since a student might be annoying for one teacher, but not for another. Such a feature can thus only be used for a risk model for that teacher, not for all teachers. A third class of features that are not suitable for selection are features whose value is unknown at the time we want to apply the model. This may seem obvious, but a risk model is usually build on historical data. We are thus in the future with respect to the date the data was created. Therefore, we may have information that was unknown at the time the data was created (e.g. the taxpayer has raised a dispute). In our student example, the fact that a student is expelled from the class is an example of such a feature.

After selecting suitable features, the teacher can start building a data set, called *Analytical Base Table* (ABT), see Figure 7.4. This means that for a while, the teacher



Figure 7.4: Building an Analytical Base Table for the student analogy of a risk model.

registers features of each student who has been checked. Each student receives a row in the Analytical Base Table and each column contains a feature of the students. The last column indicates whether or not the student has completed the homework. This column, often called the *target*, contains the information that we want to learn to predict. The Analytical Base Table is the starting point for creating a risk model.

The next step is to construct a model. Although in a previous step we have selected a set of suitable features, not all features will have predictive power. So the model building process is started with *feature selection*, i.e. selecting the most promising features. In the class example, a teacher may start looking at the values of the individual features in the Analytical Base Table and compare them with the target. The teacher can note that the feature of low average grades often goes together with the fact that the homework has not been completed. There might be no such relation between the target and a feature like 'wearing a hat'. This way, the teacher comes to a shortlist of features that are thought to be promising for selecting risky students. Now, selected features have to be combined to come to a final risk score for each student. Although many ways exist to do this, in the analogy we take the approach of assigning individual risk points to each feature and then compute the final score by adding up these risk points. In the student's example we can, for instance, assign two risk points to the student if the homework was not completed the last time and one risk point when the average grade is low. The total risk score is then 3. The determination of the right risk points for each feature is an essential part for obtaining a good model. We then demonstrate a brute force way that tries a large collection of different risk points and finally selects the one that gives the best result on the Analytical Base Table.

Finally the scoring of new students with the help of this risk model is demonstrated, visually. Here the model is represented as a machine that gets as input the features of a new student and returns a risk score This then enables to rank the students based on risk. The three students with the highest scores are eventually checked by the teacher.

#### 7.4 Recommendations based on fuzzy sets

Recently, Meza et al. [76] propose to use fuzzy sets and recommendation systems to improve taxation. Their idea is to represent some key concepts for taxation, in particular the compliance of a taxpayer and the riskiness of tax debt, as fuzzy sets. Subsequently, a tax administration may adapt its approach based on the membership degrees to these fuzzy sets. The approach has been tested on small scale by the authors for the municipal taxation of the city of Quite (Ecuador).

Although the ideas about the applicability of fuzzy sets for tax administrations are still developing, we have tested some aspects of it at the Netherlands Tax and Customs Administration, in cooperation with the authors of [76]. Two aspects have been investigated: first the applicability of the fuzzy sets approach. This resulted in Figure 7.9. Second, the predictive power of the approach, resulting in Table 7.1. Both investigations are explained below.

To test the framework, we selected 100,000 Dutch tax debts that were open on 1 July 2015 and were fully completed by 1 August 2018. We have ensured that all selected debts were before the phase that a judicial notice had to be sent. We selected these 100,000 debts in such a way that 50,000 were paid before the payment deadline (or a payment arrangement was requested), and 50,000 at least 10 days after the deadline. In all other respects, the debts are selected at random.

For each debt, we looked at the citizen involved and constructed a Citizen Behavior Ranking. We looked at historic payment data of the citizen in the period from 1 July 2013 to 1 July 2015. For each debt of the citizen that was closed in that two-year period, we recorded whether the debt was paid before the payment deadline or not. Then we calculated the ratio of the number of debts that were paid in time to all debts. Figure 7.5 shows a histogram of the CBR for our population belonging to the 100,000 debts.



Figure 7.5: Citizen Behavior Ranking for the 100,000 citizens belonging to the 100,000 debts

Based on the Citizen Behavior Ranking, fuzzy sets for Citizen Behavior can be easily created. For instance by using Figure 7.6. Applying the membership functions of Figure 7.6, we find that a citizen with a Citizen Behavior Ranking of, say, 0.85 has a membership degree to the fuzzy set 'Excellent' of 0.4, a membership degree of 0.6 to the fuzzy set 'Good' and 0 for all other fuzzy sets.

Note that to calculate the Citizen Behavior Ranking we made use of a Central Data Repository that is daily refreshed using ETL processes. To respect the privacy of the taxpayers involved, the analysis has been performed using data where all identifying records have been removed.

For the 100,000 debts mentioned before, we calculated the days till the payment deadline. In some cases this number was negative, which means that the payment deadline had been passed already. Inspired on Figure 3 of [76], here reproduced for reference as Figure 7.7, we come to a membership degree for the fuzzy set 'Payment Deadline'. A membership degree of 1 is given to debts that were over the payment deadline and for debts that were before the deadline a linear relation was used, starting with a membership degree of 0 on the day the debt was communicated to the citizen.

Similarly, a membership degree for the fuzzy set 'risk of debt' is computed. We



Figure 7.6: Fuzzy sets for Citizen Behavior for the Dutch Citizens belonging to the 100,000 selected debts.



Figure 7.7: Reproduction of Figure 3 of [76], showing the relation between the payment deadline of a tax debt and the membership function of a fuzzy set related to late payment.

ordered the amounts of the 100,000 debts and took the rank number and divided that by 20,000 in order to get a risk relevance scale between 0 and 5. Subsequently, we applied the curve shown in Figure 4 of [76], here reproduced for reference as Figure 7.8, to come to a membership degree to the fuzzy set 'risk of debt'.



Figure 7.8: Reproduction of Figure 4 of [76], showing the relation between the risk relevance and the membership function of a fuzzy set related to extreme riskiness payment.

The Citizen Behavior Ranking, membership functions of 'payment deadline', and 'risk of debt' can be plotted in two two-dimensional Figures as shown in Figure 7.9. Note that we have plotted a random sample of size 200 of all 100,000 debts. The dashed line separates the risky debts from the less risky debts. It is up to the management of the tax authority to define this line. After the separation, different treatments (recommendations) for the risky and non-risky debts can be considered. For instance, it can be decided to do nothing for the non-risky debts, while the risky debts for payment deadline are sent a reminder letter, while the risky debts for risk relevance are receiving additional monitoring.

To assess the predictability of the separations in Figure 7.9, we focused on the plot at the top of Figure 7.9, i.e. on the risk of non-payment before the payment deadline. The 100,000 debts are separated a priori according to the separation line. Subsequently, it was checked a posteriori whether the debts were actually paid before the payment deadline. This gives the confusion matrix shownn as Table 7.1. The Chi square test for independence has been applied to the confusion matrix, resulting in a value of 9,104.6 for the chi square statistic. Since the degrees of freedom is equal to



Figure 7.9: Citizen Behavior Ranking versus Payment Deadline and Risk Relevance for a random sample of 200 debts. The dotted lines separate risky debts from non-risky debts.

A Priori Risk	Paid	Not Paid
Not Risky	38,830	24,269
Risky	11,170	25,731
Total	50,000	50,000

Table 7.1: Confusion matrix. Riskiness versus Paid on time

1, this leads to a p-value of  $2.2 \cdot 10^{-16}$ , signifying a clear dependence between a priori risk predictions and actual payment problems.