Canadian Journal of Law and Technology

Volume 17 | Number 1 Article 3

6-1-2019

Hardship and Hard Drives: Artificial Intelligence, Judicial Decision-Making, and the Discharge of Student Loan Debt

Forrest Finn
JD Candidate, University of Toronto

Follow this and additional works at: https://digitalcommons.schulichlaw.dal.ca/cjlt

Part of the Courts Commons, Jurisprudence Commons, and the Science and Technology Law Commons

Recommended Citation
Forrest Finn, "Hardship and Hard Drives: Artificial Intelligence, Judicial Decision-Making, and the Discharge of Student Loan Debt" (2019) 17:1 CJLT 59.

This Article is brought to you for free and open access by the Journals at Schulich Law Scholars. It has been accepted for inclusion in Canadian Journal of Law and Technology by an authorized editor of Schulich Law Scholars. For more information, please contact hannah.steeves@dal.ca.

Hardship and Hard Drives: Artificial Intelligence, Judicial Decision-Making, and the Discharge of Student Loan Debt

Forrest Finn*

ABSTRACT

Section 178(1.1) of the Bankruptcy and Insolvency Act allows individuals to apply for discretionary relief from the non-dischargeable nature of student loan debts. Subparagraph (b) of this relief establishes a "hardship" requirement. The elements for this hardship requirement have been developed and applied by judges in the form of standards. The issue addressed in this paper is whether these standards are applied predictably. Using both statistical analysis and machine learning algorithms, this paper demonstrates that judicial decision-making on the hardship requirement is predictable. This predictability has significant implications. Most importantly it suggests that predictive software could be created for s. 178(1.1) applications that could significantly reduce the uncertainty and cost of these applications.

TABLE OF CONTENTS

---

* JD, University of Toronto.

INTRODUCTION

In September 30, 1997, Parliament introduced s. 178(1)(g) into the *Bankruptcy and Insolvency Act*.[1] This enactment made student loan debts non-dischargeable in bankruptcy for a period of seven years.[2] However, Parliament was concerned that this provision could prejudice honest and unfortunate debtors. To provide some relief Parliament also introduced s. 178(1.1). This subsection allows a court to grant discretionary relief after five years provided that the following two prerequisites have been met:

(a) the bankrupt has acted in good faith in connection with the bankrupt's liabilities under the debt; and
(b) the bankrupt has and will continue to experience financial difficulty to such an extent that the bankrupt will be unable to pay the debt.[3]

While paragraph (a) is concerned with the good faith of the applicant, paragraph (b) was enacted with the express purpose of "recogniz[ing] that some students may find themselves in a hardship situation".[4] The interpretation of this "hardship" requirement was left to the judges.

In this paper, I examine 30 cases where the courts have decided on this hardship issue.[5] I argue that, in determining hardship, the courts act predictably. In proving this, I show two things. First, I use statistical analysis to show that there are observable trends in how important factors have been weighed by the courts. Second, I use machine learning to demonstrate that artificial intelligence can predict *ex ante* how a judge is likely to decide the hardship question with a high degree of certainty.

METHODOLOGY

Source Material

This study is based on data collected from court cases considering s. 178(1.1) applications. To build this dataset, I conducted a keyword search of s. 178(1.1) cases and decisions through WestlawNext. This search identified 138 cases and decisions. Of these cases, I decided to create a sample of 30. I used the convenience sampling method. This sampling method is a form of non-probability sampling often used for pilot studies.[6] Essentially, I started with

---

1 *Bankruptcy and Insolvency Act*, R.S.C. 1985, c. B-3.

2 The original provision was for two years. It was later increased to ten and subsequently decreased to seven.

3 *Bankruptcy and Insolvency Act, supra* note 1 at s. 178(1.1)(a)-(b).

4 *Ibid.*

5 For a full list of cases please see Appendix A.

6 For a full description of convenience sampling and other sampling methods, please see Mark Saunders, Philip Lewis & Adrian Thornhill, *Research Methods for Business Students*, 6th ed (Toronto: Pearson, 2012).

the first case in the search, going through the list until I had 30 cases. There were two central issues that disqualified a case from being incorporated in the dataset. First, I disqualified any cases in which the applicant was found to have acted in bad faith. Under s. 178(1.1)(a) bad faith automatically disqualifies the application and the judges do not need to consider the hardship question addressed in this paper. Second, when I first constructed the dataset, the information I collected far exceeded the factors discussed in this paper. I wanted to consider all the possible relevant factors that could influence a decision before deciding on the ones subject in this paper. For this reason, to be considered for the dataset a case was required to have the following information:

- education level and degree class;
- time in years since the applicant was a student;
- loan amount outstanding;
- yearly income of the applicant;
- age of the applicant;
- whether the applicant had children and, if so, how many;
- marital status of the applicant;
- whether the applicant pays child support;
- if the applicant owns a home, car, or other noted assets;
- whether the applicant completed the education; and
- whether the applicant used the education to earn income.

Any cases that did not have this information were not considered. In total, I read and considered 70 cases before finalizing the 30 cases that are the subject of this study. Finally, after looking at all the information collected, I focused on the factors with the highest degree of relevance and predictive value.

Limitations and Sources of Error

This study has several limitations caused by both its scope and the practical realities of preparing the dataset. The first and likely most important limitation is the small sample size. Thirty cases are not enough to fully determine how courts have weighed the factors. As a result, this study is not determinative. Rather, it is largely preliminary and opens the door to further, more exhaustive quantitative research on this subject.

The small sample size also had a significant impact on evaluating the accuracy of the machine learning programs. The low quantity of data points meant that the accuracy of the algorithms was likely dependant on where the program split the training and testing data. To avoid both overrepresenting and underrepresenting the accuracy of the algorithms I decided to average the accuracy results. For each algorithm (and on each type of data) I ran the machine learning program 20 times. Each time the programs split the data at random. I then averaged the results.

There may also be an issue of selection bias. Typically, convenience selection statistical inferences cannot easily be extrapolated to the general population. This is caused by the likely under-representation of subgroups in the sample as compared to the general population of interest.[7] This is likely mitigated in this study by the fact that the "convenience" of the sampling was in the availability of reported cases.[8]

An additional source of possible error is the ever-present risk of human error. As advanced as statistical analysis and machine learning have become over the last 20 years, they are still only as good as the data they use. Most data is either curated or constructed by humans who may unknowingly introduce mistakes or inconsistencies. This is possible here as well. I personally created the data points by reading the cases and making the entries into the dataset. One mitigating factor to this error might be that it was entirely created by a single person. This would reduce the possibility of inconsistencies in understandings or process that could lead to flaws in the data. This avoids a consistent point of tension within large data operations employing many people.

PART I: THE PREDICTIVE FACTORS

I will use two methods to test my hypotheses of the predictive factors. First, I have sorted the data into quantitative categories or "bins" to find and demonstrate any trends in the cases. While this process generally destroys the original detail of the data, it has the advantage of showing a clear overall picture the can elucidate otherwise unclear relationships.[9] The bins selected are largely arbitrary and have been selected to ensure a relatively equal amount of cases are placed in each bin. This was done to avoid outliers from skewing the results of a category. However, this does introduce the possibility of bias in the selection of the bins. This issue is dealt with at the second stage below. The results from this first analysis will be represented by histogram graphs.[10]

Second, to avoid the possible bias described above, I have also conducted a binary logistic regression using the R coding language on the raw data. This regression evaluates the significance of the variable. To evaluate the impact of the variable I have compared the null deviance to the residual deviance. The null deviance gives the accuracy of the model when the independent variable is only used as the intercept. In all the models, the null deviance was 27.392 on 30 degrees of freedom. The residual deviance uses the variable being tested. The

---

7 Marc H Bornstein, Justin Jager & Diane L Putnick, "Sampling in Developmental Science: Situations, Shortcomings, Solutions and Standards" (2017) 33:4 Developmental Rev 357.

8 Indeed, often convenience sampling involves interviewing people on the street. This would raise significant issues with, for example, the location or the time of the interviews adding bias to the sample. These criticisms should not apply here as my sample was drawn from the most relevant reported cases on the subject.

9 Murray R Spiegel & Larry J Stephens, *Schaum's Outline of Statistics,* 6th ed (New York: McGraw-Hill Education, 2017) at 36.

10 *Ibid* at 37 - 39: histograms are graphic representations of frequency distributions. They are important constructs intended to show the researcher how the data is distributed.

smaller the deviance, the better the model. In addition, I will use the Akaike Information Criterion ("AIC") to evaluate the quality of each model relative to each other.[11] The lower the AIC is, relative to the other models, the better.

Debt and Income

In determining the question of financial hardship, the courts have consistently emphasized two factors: the amount of debt and the amount of income. For example, in *Cook, Re* the court distinguished Ms. Cook's situation from that in *Rendely, Re* (2003) stating that Ms. Cook "does not have near the income nor expectation of substantial future income [as Rendely] . . . and has twice the debt".[12] Because of the high debt and low income, the court found that Ms. Cook was experiencing hardship under paragraph (b) and granted her application. In *Beaton, Re,* the court dismissed the application, contrasting Ms. Beaton's situation to those that *should* receive relief under s. 178(1.1) stating about the latter that "[t]here is no reasonable hope for them ever . . . being able to address the outstanding debt. *Many have no prospect of ever making significantly more money*. Many have children to support. *The outstanding loans of many are very substantial, $40,000 and more*" [my emphasis].[13] In contrast, "Ms. Beaton . . . has earned a good salary in the past. Her outstanding balance is only about $12,000".[14]

As we can see, both the size of the debt and the size of the income play an important role in determining financial hardship. I put forward two hypotheses therefore. First, as the loan amount increases, acceptance rates will increase. Second, as income increases, acceptance rates will decrease. I will test both of these in turn below.

Acceptance Rate Based on the Size of the Debt

Hypothesis # 1
Based on the cases above, I predict that as the size of the debt (loan amount) increases, so too will the acceptance rate.

---

[11] Jan deLeeuw, "Introduction to Akaike (1973) Information Theory and an Extension of the Maximum Likelihood Principle" in Samuel Kotz & Norman L Johnson, eds, *Breakthroughs in Statistics I* (New York: Springer, 1992) 599; Sadanori Konishi & Genshiro Kitagawa, *Information Criteria and Statistical Modelling* (New York: Springer, 2008): AIC is founded on information theory and serves as an estimator of the relative quality of statistic models for a given set of data. Where there is a collection of models for a dataset, AIC estimates the quality of each model, relative to the other models. This provides a means for model selection.

[12] *Cook, Re*, 2010 NSSC 224, 2010 Carswell NS 386, [2010] N.S.J. No. 333, 291 N.S.R. (3d) 380 (N.S. S.C.) at para 20; *Rendely, Re*, 2003 CarswellOnt 4392, [2003] O.J. No. 4678, 3 C.B.R. (5th) 136 (Ont. S.C.J.).

[13] *Beaton, Re,* 2012 NSSC 281, [2012] N.S.J. No 403, 319 N.S.R. (2d) 191 (N.S. S.C.) at para 12.

[14] *Ibid* at para 13.

1) < $30,000;
2) $30,000 - $60,000; and
3) > $60,000.

The acceptance rates by category can be seen in Figure 1 below:

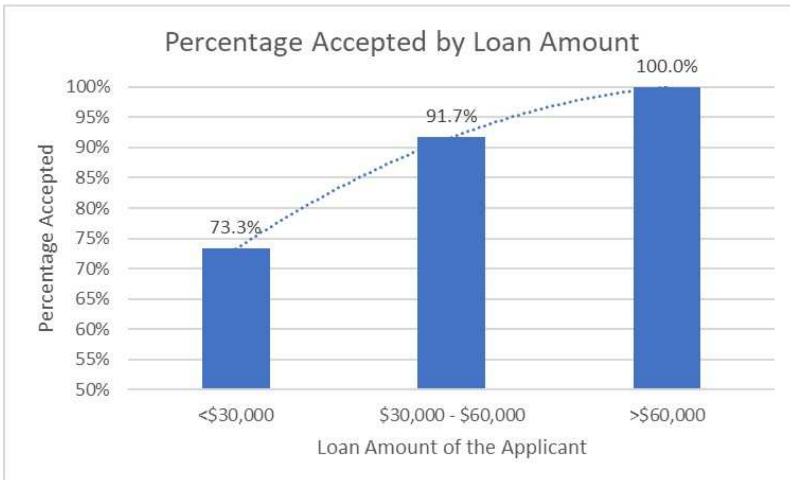

Figure 1: Percentage Accepted by Loan Amount.

As we can see from the graph, as loan amounts increase, so too does the likelihood of acceptance. This seems to flow logically as a smaller debt would take less income surplus to pay off and could, therefore, be paid off faster. These results are associated with less financial hardship.

This was also reflected in the binary logistic regression on whether debt alone could predict the result. The regression found a residual deviance of 24.263 on 29 degrees of freedom. This is lower than the null deviance of 27.392 on 30 degrees of freedom. The variable clearly had an impact on the deviance lending it predictive value. In addition, its AIC was 28.263

Acceptance Rate Based on the Size of the Income

Hypothesis # 2
Similarly, based on the statements in the case law, I predict that as income increases, the application acceptance rate will decrease. To test this hypothesis, I have separated the cases into four income groups:

1) < $20,000;
2) $20,000 - $30,000;
3) $30,001 - $50,000; and
4) > $50,000.

The results can be seen in Figure 2 below:

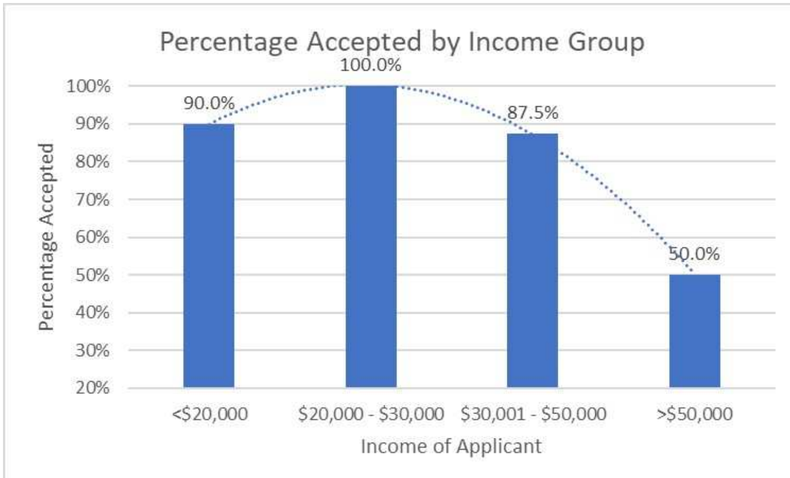

Figure 2: Percentage Accepted by Income Group.

As we can see, the general trend seems to prove my hypothesis. Generally, as income increases, the acceptance rate falls. However, there is an interesting increase in the acceptance rate of 10% from the first to the second category. There are several possible explanations for this. First, this may be a distortion from the small sample size. It is possible that rejected applications are overrepresented in the < $20,000 income group and underrepresented in the $20,000 – $30,000 income group.

Second, this may imply that judges are willing to evaluate the reasonableness of an applicant's employment decisions in evaluating the applicant's hardship. In *Godfrey, Re,* for example, the applicant's income was very low. However, she only worked 25.5 hours a week. In finding bad faith, the court bluntly noted that "[s]he is underemployed".[15] While *Godfrey* was decided on the issue of good faith, it does raise the interesting possibility that this consideration may be impliedly considered in determining hardship. More research is required on this subject.

Finally, in some cases the debt may have been small enough as to make the low income irrelevant. If this is true then neither evaluating the debt nor the income alone is sufficient. Rather, we should look at debt as a percentage of income.

The binary logistic regression analysis, however, had a different finding. The residual deviance from the model was 4.889 e^{-10} on 1 degree of freedom. This is a particularly bad score. In addition, the AIC was 60. This was also a bad score, especially when compared to the score of debt alone, which had an AIC of 28.263. These results indicate that the factor is significantly less predictive than the histogram may suggest. Indeed, while my hypothesis predicted income would

---

[15] *Godfrey, Re*, 2017 NBQB 5, [2016] N.B.J. No. 297 (N.B. Q.B.) at para 21.

be a determinative factor, this appears wrong. Still, it will be important to evaluate whether the two variables may be predictive when working together. This is evaluated in debt as a percentage of income.

Debt as a Percentage of Income

Hypothesis # 3
I hypothesize that where debt as a percentage of income increases, the acceptance rate of applications will also increase.

To test this hypothesis, I have separated the cases into three groups according to debt as a percentage of income:

1) < 50%;
2) 50% - 100%; and
3) > 100%.

The results can be seen in Figure 3 below:

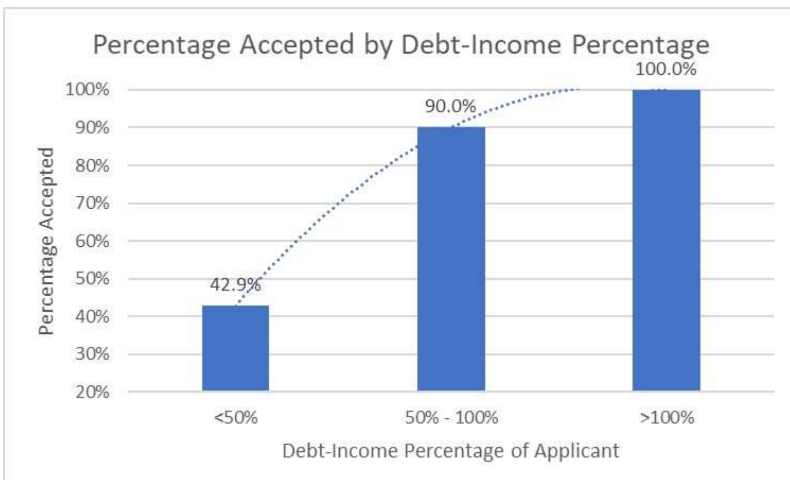

Figure 3: Percentage Accepted by Debt-Income Percentage

The histogram shows that there is a dramatic increase in acceptance rates once the debt crosses the 50% of income threshold. Above 100%, all applications were accepted. This dramatic increase likely extends from the interpretation of "continue to experience" financial hardship under paragraph 178(1.1)(b). While there is no hard and fast rule that courts can use to determine the appropriate time period,[16] a high debt-to-income percentage likely results in a longer repayment schedule than most judges are comfortable accepting.

[16] In *Minto, Re,* [1999] S.J. No. 798, 191 Sask. R. 1 (Sask. Q.B.) at para 88, the court held

The binary logistic regression analysis provided similar results. Indeed, the residual deviance was 14.930 on 28 degrees of freedom. This was a significantly better score than debt alone. In addition, the AIC was 20.93, the best score of the three variables and combinations discussed so far. As a result, it appears likely that judges are evaluating these two factors together.

Age of the Applicant

Another factor that may limit the time period of repayment and thereby increase the likelihood of an application's success is the age of the applicant. The case law seemed to indicate that judges are open to considering a longer repayment schedule for younger applicants[17] and a shorter repayment schedule for older applicants.[18] This follows logically. The older a person is the fewer working years they have left in their life. Furthermore, as people age they need to provide more of their resources towards retirement.[19] The opposite has been noted with young applicants. For example, in *Re Burke,* Saunders J. stated that "Mr. Burke is 28 years of age. He has a long life ahead of him, and, potentially, considerable earning capacity".[20]

Hypothesis # 4
I hypothesize that as age increases, the acceptance rate of applications will also increase.

To test this hypothesis, I have separated the cases into three age categories:

1) 35 and Under;
2) 36 — 45; and
3) > 45.

The results can be seen in Figure 4 below:

that the appropriate period will depend on the facts of each situation. In *Wood, Re*, [1998] M.J. No. 466, 133 Man.R. (2d) 230 (Man. Q.B.), the appropriate time period was two to three years. In *Pyke, Re,* 2005 NSSC 33, [2005] N.S.J. No. 54, 230 N.S.R. (2d) 104 (N.S. S.C.) at para 55, the court found that, based on the facts, "I think the future that one need consider is not that long."

17 *Burke, Re*, [1992] N.S.J. No. 236, 114 N.S.R. (2d) 160 (N.S. S.C.).

18 *Cook, Re, supra* note 12 at paras 23, 31.

19 *Ibid.*

20 *Burke, Re, supra* note 17.

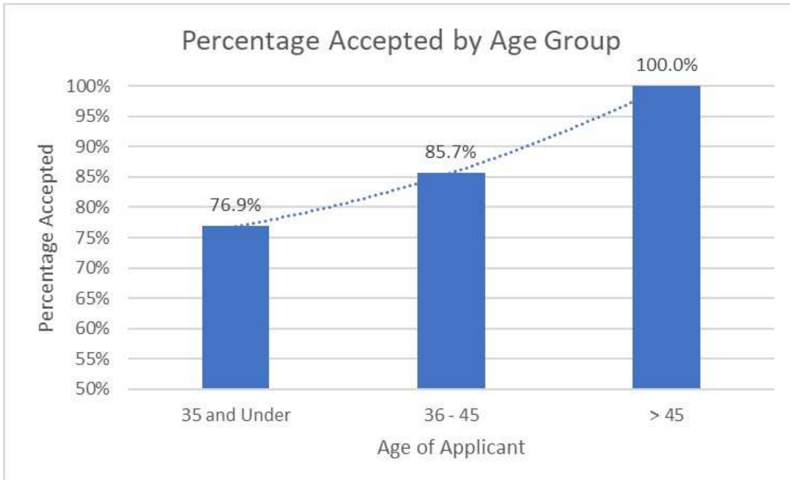

Figure 4: Percentage Accepted by Age Group

As we can see from the histogram above, the data meets our expectations. As age increases, so too does the acceptance rate.

The logistic regression also showed predictive value but was less dramatic than the histogram results. While the null deviance is 27.392 on 30 degrees of freedom, the residual deviance was only 26.953 on 29 degrees of freedom. Moreover, the AIC was 30.953. This is not a bad score but is less than debt alone (28.263) and significantly less than debt and income together (20.93). Age, therefore, has some predictive value, but the trend viewed in the histogram above may be subject to some selection bias.

Debt, Income, and Age

While not necessarily appropriate to be represented in a histogram, I have also done a binary logistic regression on the combination of the three above factors: debt, income, and age. The residual deviance was 14.756 on 27 degrees of freedom, scoring marginally better than the previous best score of debt and income together (14.930). However, the AIC, at 22.756, did worse than the debt and income model (20.93).[21]

Part I Conclusion

As has been demonstrated by the histograms and binary logistic regressions, the data shows clear patterns in how judges in the aggregate determine hardship

[21] Spiegel & Stephens, *supra* note 9: The representation of the process that generated the data in a statistical model is never exact and some information is lost. The AIC estimates the relative amount of information lost. It is possible that this reduction in AIC was caused by the fact that an increase in variables may increase the likelihood of information being lost.

under s. 178(1.1)(b). While income was interestingly not a predictive factor alone, it increased the accuracy of the models significantly when combined with other factors. The predictive value discovered in these factors has significant implications for machine learning on this area of law. Indeed, it suggests that judges are acting consistently based on several discrete factors. In Part II below I apply these findings in the context of machine learning to determine whether these discrete factors can be used to predict judicial decision-making *ex ante*.

PART II: MACHINE LEARNING

The question the above conclusion raises is whether we can take these descriptive results and use them to accurately predict *ex ante* how a court will rule given a particular factual scenario.

Hypothesis # 5
I hypothesize that, considering the clear trends shown above, artificial intelligence can accurately predict judicial decision-making on questions of hardship in s. 178(1.1) applications.

To test this hypothesis, I wrote three machine learning programs using three different machine learning algorithmic techniques. I then trained the algorithms on the data using different combinations of the above factors and evaluated their accuracy.

The Algorithms

To test my hypothesis, I used three different types of algorithms: K-Nearest Neighbour, Decision Tree, and Random Forest. Before discussing the results, I will briefly explain how each algorithm interacted with the data.

K-Nearest Neighbour

K-Nearest Neighbour is one of the simplest classification algorithms. It is incredibly useful when there is either little or no knowledge about the distribution of the data.[22] It operates by using the distance method. It plots the training data based on the various factors represented in the training dataset. Each training point has a class value expressed as a 1 or a 0. 1 means the application was accepted, whereas a 0 means the application was rejected. Then the test points are similarly plotted. The algorithm does not know the class of the test points. The distance between the test points and the training points are measured using Euclidean distance.[23] After this, the nearest neighbours to the test data are determined using the kth minimum distance.[24] After all the values

---

22 Suresh Chandra Satapathy & Jyotsna Kumar Mandal, *Information Systems Design and Intelligent Applications: Proceedings of Third International Conference INDIA* (New Delhi: Springer India, 2016) at 190.

23 *Ibid* at 190: Euclidean distance formula: $d = \sqrt{(x2 - x1) + (y2 - y1)}$.

24 *Ibid* at 191: kth is determined by the author of the algorithm. For example, if the author

are collected, the computer will check which class values is higher.[25] For example, if k = 7 and there are 4 values for accepted applications and 3 for rejected applications, the computer would classify that point as an accepted application. The rationale behind this method assumes that the features used to describe the points are relevant to their labels.[26] As a result, close-by points are likely to have the same label. Things that look alike must be alike.[27]

Decision Tree

A decision tree is a classification algorithm that predicts the label associated with an instance. Each question is a node and each answer is a leaf that travels down the tree.[28] The questions and nodes are determined by the algorithm. To determine the result, the decision tree classifier follows the tree. The benefit of this classifier is that it is very simple to understand and interpret.[29]

Random Forest

The random forest is related to the decision tree algorithm. It operates by constructing a multitude of decision trees. The output is the mode of the classes as determined by all of the decision trees.[30] This algorithm solves one of the primary issues with using decision trees: overfitting. Because random forests average multiple deep decisions trees, trained on different parts of the same training set, they reduce this risk of overfitting.[31]

Factors Considered

The results in Part I indicated that the two most predictive models were those that used debt and income, and those that used debt, age, and income. Considering this, I have decided to train the machine learning algorithms on these combinations of factors, rather than debt and income separately. In addition, due to the small size of the dataset I have decided to train and test the algorithms on both the intervals used in the histograms and the integer values of the data points themselves.

---

determines that k should equal 5, then all the sorted distances greater than five are ignored.

25 *Ibid.*

26 Shai Shalev-Shwartz, *Understanding Machine Learning: From Theory to Algorithms* (New York: Cambridge University Press, 2014) at 219.

27 *Ibid* at 225.

28 *Ibid* at 212.

29 *Ibid* at 213.

30 Tin Kam Ho, "The Random Subspace Method for Constructing Decision Forests" (1998) 20:8 IEEE Transactions on Pattern Analysis and Machine Intelligence 832.

31 Trevor Hastie, Robert Tibshirani & Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed (New York: Springer, 2009) at 596.

Training the Algorithms

To train the algorithms, I wrote a program in the python coding language for each respective algorithm. In the first step I loaded the csv file of my collected data. The classification (i.e. whether the application is accepted or rejected) was recorded as the label. The data was then split at random into two parts. 70% of the data became the training data. This is what the computer trained the machine learning algorithms on. Then, the algorithm was tested on the remaining 30% of the data: the testing data. The algorithm created a prediction for each test point.

Testing the Accuracy

These predictions were then tested against the training data's labels. This resulted in an accuracy score reflected as a percentage. Because the data set I worked with is very small (only 30 data points), I tested each type of algorithm twenty times. I then determined the accuracy for each type of algorithm based on the average of these twenty attempts. I did this because, due to the small size of the data set, the location of the split between the training and testing data may have a significant impact on the predictive ability. In addition, I have included frequency polygons for each of the results. These demonstrate both the spread of the accuracy results as well as the clusters, if any, in the accuracy of the algorithms.

Machine Learning Results

First, I ran each of the three algorithms on just the debt as a percentage of income intervals (discussed in Figure 3). I had the following results:

DECISION TREE	RANDOM FOREST	K NEAREST NEIGHBOUR
84.5%	82.2%	77.8%

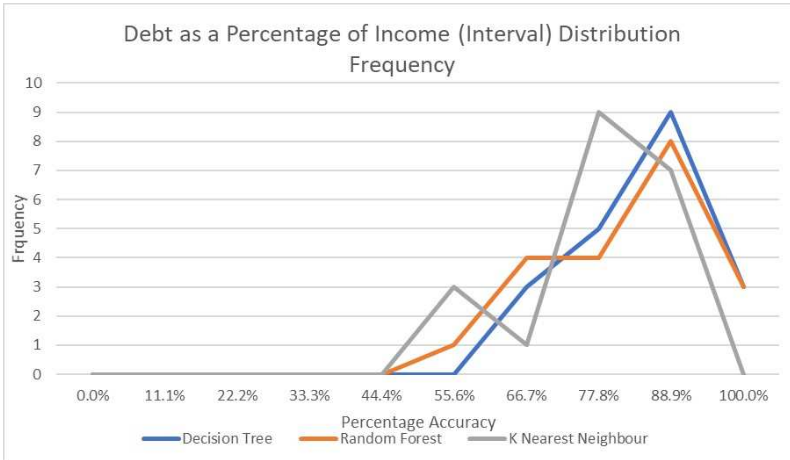

Figure 5: Debt as a Percentage of Income (Interval) Distribution Frequency

Using the same data without the intervals (i.e. in their original integer form), both the decision tree and random forest algorithms performed slightly worse. K Nearest Neighbour, in contrast, performed the same:

DECISION TREE	RANDOM FOREST	K NEAREST NEIGHBOUR
81.7%	79.5%	77.8%

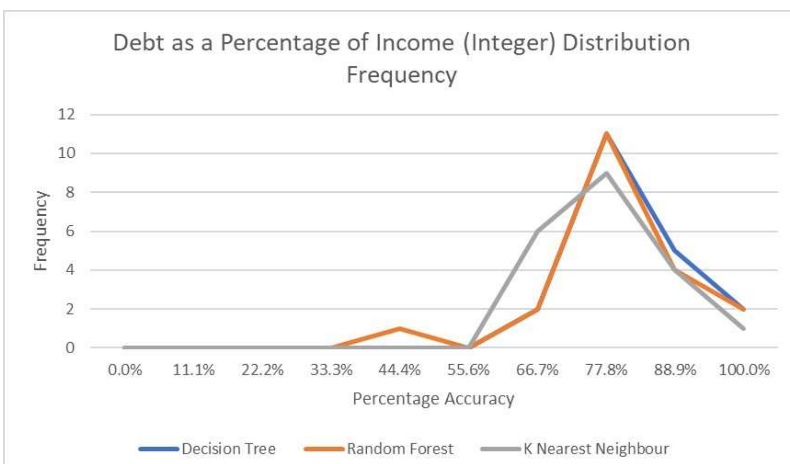

Figure 6: Debt as a Percentage of Income (Integer) Distribution Frequency

When I added the age factor as an interval, the decision tree's accuracy fell slightly more. The random forest algorithm performed slightly better than the second time, but still worse than the first. In contrast, k nearest neighbour scored its highest accuracy result.

DECISION TREE	RANDOM FOREST	K NEAREST NEIGH-BOUR
81.1%	80.6%	83.4%

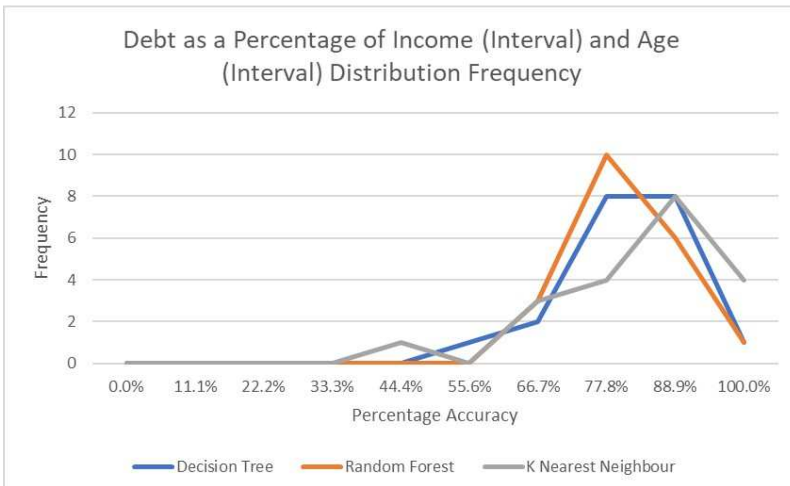

Figure 7: Debt as a Percentage of Income (Interval) and Age (Interval) Distribution Frequency

Finally, when both age and debt as a percentage of income are expressed in their integer form, the decision tree and random forest algorithms performed at their lowest accuracy. K nearest neighbour scored only slightly lower than its best.

DECISION TREE	RANDOM FOREST	K NEAREST NEIGHBOUR
75.0%	75.6%	82.2%

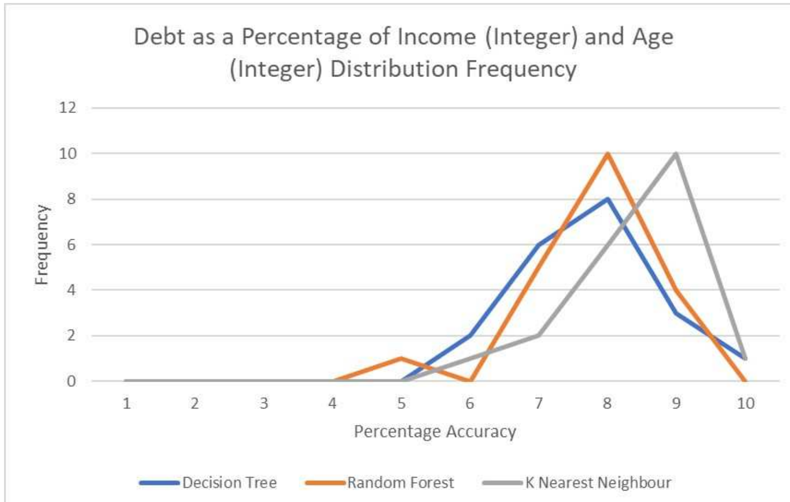

Figure 8: Debt as a Percentage of Income (Integer) and Age (Integer) Distribution Frequency

Due to the variations in percentage accuracy and the close relative scores, it is not possible from these results to conclude that any one algorithm or combination of factors is best suited to this area of law. What I can conclude, however, is that, for a dataset of only 30 datapoints, all had higher than expected accuracies. In addition, the frequency distributions were highly clustered in most of the frequency polygons. This suggests that there may be some outliers or pairs of dissonant cases in the sample that are skewing the results depending on whether they make it into the training or testing data. With a larger data set, incorporating more cases, the accuracy is likely to increase.

Part II Conclusion

In 1897, Oliver Wendell Holmes Jr. predicted that "[f]or the rational study of the law the blackletter man may be the man of the present, but the man of the future is the man of statistics and the master of economics".[32] With the recent advances in statistical theory and machine learning, these predictions are growing. Indeed, in a 2016 article, Professor Ben Alarie predicted that future technological advances will lead to a "legal singularity" in which legal standards will be overtaken by complex systems of rules that will greatly increase reliability, predictability, and accessibility of the law.[33]

The machine learning results above have several corroborating implications in the context of s. 178(1.1)(b) applications. First, it suggests that we can predict

---

[32] Oliver Wendell Holmes Jr., "The Path of the Law" (1897) 10:8 Harvard L Rev 457 at 469.

[33] Benjamin Alarie, "The Path of the Law: Toward Legal Singularity" (2016) 66:4 UTLJ 443.

how a judge will decide on the question of hardship *ex ante*. This could open a market for predictive tools that could reduce the transaction costs involved in filing such an application. For example, predictive tools could be created that not only predict a result but provide a measure of confidence. If the program predicts that an application would fail with a high degree of confidence, the individual should not bring the application. If the program predicts success or a low confidence of failure, the individual could decide for themselves whether the claim is worth the cost. Similar tools already exist in other areas of law.[34]

Second, if such a tool were made available to the public, this may promote access to justice in the field of insolvency law. The average person does not actually know the law or understand how provisions such as s. 178(1.1) function. By providing access to tools like this that show the operation of the provision, the average person would be better able to understand how they apply to concrete situations.

APPENDIX A: TABLE OF CASES

Case#	Case Citation
1.	*Aymar, Re*, 2017 NSSC 328, [2017] N.S.J. No. 493 (N.S. S.C.)
2.	*Bauckman, Re*, 2013 NSSC 34, [2013] N.S.J. No. 51 (N.S. S.C.)
3.	*Beaton, Re,* 2012 NSSC 281, [2012] N.S.J. No. 403 (N.S. S.C.)
4.	*Beaupre, Re*, 2016 NBQB 36, [2016] N.B.J. No. 21 (N.B. Q.B.)
5.	*Chrisara, Re*, 2017 BCSC 811, [2017] B.C.J. No. 930 (B.C. S.C.)
6.	*Cook, Re*, 2006 CarswellOnt 688, [2006] O.J. No. 493, 20 C.B.R. (5th) 192 (Ont. S.C.J.)
7.	*Cook, Re*, 2010 NSSC 224, 2010 CarswellNS 368, [2010] N.S.J. No. 333 (N.S. S.C.)
8.	*Cote, Re*, 2010 BCSC 490, 2010 CarswellBC 868, [2010] B.C.J. No. 652 (B.C. S.C.)
9.	*Cusack, Re*, 2014 SKQB 136, 2014 CraswellSask 361, [2014] S.J. No. 322 (Sask. Q.B.)
10.	*Dobbelsteyn, Re*, 2016 NBQB 58, [2016] N.B.J. No. 49 (N.B. Q.B.)
11.	*Dunfield, Re,* 2013 NBQB 195, 2013 CarswellNB 273, [2013] N.B.J. No. 163 (N.B. Q.B.)
12.	*Dunn, Re,* 2012 NSSC 240, [2012] N.S.J. No. 338 (N.S. S.C.)
13.	*Field-Currie, Re*, 2010 NSSC 41, 2010 CarswellNS 45, [2010] N.S.J. No. 40 (N.S. S.C.)
14.	*Furey, Re,* 2016 NSSC 47, [2016] N.S.J. No. 67 (N.S. S.C.)
15.	*Hankinson, Re,* 2009 NSSC 211, 2009 CarswellNS 381, [2009] N.S.J. No. 312 (N.S. S. C.)

[34] For example, Blue J Legal, a Canadian legal technology company, currently has commercially successful machine learning classifiers in the areas of tax and employment law that have a similar function.

16. *Hughes, Re*, 2018 NSSC 189, 2018 CarswellNS 592, [2018] N.S.J. No. 303 (N.S. S.C.)
17. *Lundrigan, Re,* 2012 NSSC 231, [2012] N.S.J. No. 326 (N.S. S.C.)
18. *McFarlane, Re,* 2015 NSSC 263, [2015] N.S.J. No. 397 (N.S. S.C.)
19. *McNutt, Re,* 2008 NSSC 166, 2008 CarswellNS 283, [2008] N.S.J. No. 228 (N.S. S.C.)
20. *Miller, Re*, 2016 BCSC 787, [2016] B.C.J. No. 898 (B.C. S.C.)
21. *Minto, Re,* [1999] S.J. No. 798, 191 Sask. R. 1 (Sask. Q.B.)
22. *Pyke, Re,* 2005 NSSC 33, 2005 CarswellNS 60, [2005] N.S.J. No. 54 (N.S. S.C.)
23. *Rendely, Re,* 2003 CarswellOnt 4392, [2003] O.J. No. 4678, 3 C.B.R. (5th) 136 (Ont. Sp. Ct. J.).
24. *Roach, Re,* 2008 NSSC 15, [2008] N.S.J. No. 19 (N.S. S.C.)
25. *Roy, Re,* 2016 BCSC 1845, [2016] B.C.J. No. 2094 (B.C. S.C.)
26. *Slanina, Re,* 2009 BCSC 1881, 2009 CarswellBC 3807, [2009] B.C.J. No. 2811 (B.C. S.C.)
27. *Swann, Re,* 2001 BCSC 1175, [2001] B.C.J. No. 1703 (B.C. S.C.)
28. *Taylor Re*, 2017 CarswellNfld 413, [2017] N.J. No. 357, 53 C.B.R. (6th) 177 (N.L. S.C.)
29. *Westwood, Re,* 2005 BCSC 1575, [2005] B.C.J. 2416 (B.C. S.C.)
30. *Wood, Re,* [1998] M.J. No. 466, 133 Man. R. (2d) 230 (Man. Q.B.)