Random forest and deep learning models in the field of human decision making

Timothy Ntambala 503302
ERASMUS UNIVERSITY ROTTERDAM, Erasmus School of Economics
Supervisor: A. Baillon, Second Assesor: M. Zhelonkin

”03-07-2022”

Abstract

This paper aims to determine whether random forest models do a better job at predicting human decisions on risky gambles than deep learning models. The research uses responses to 13,006 risky choice problems. This data set has been used to estimate a random forest model and 2 deep learning models. The results show that, when comparing the predictive performance based on the out-of-sample mean squared prediction error, the most flexible deep learning model outperforms the random forest model. The more constrained deep learning model, which had the assumption of context independence, did not manage to outperform the random forest model.

Contents

1 Introduction **3**

2 Literature Review **4**

2.1 Economic theory 4

2.2 Machine Learning 9

3 Data **10**

4 Methodology **11**

4.1 Deep learning 11

4.2 Context-Dependent theory 13

4.3 Value-Based theory 14

4.4 Random Forest Model 15

4.5 Random Forest Regression - Applied 17

4.5.1 Feature importance 17

4.6 Model Evaluation 18

4.6.1 Mean Squared Error 18

4.6.2 Segmented Accuracy 18

4.6.3 Number of outcomes 18

5 Results & Discussion **19**

5.1 Predictive performance 19

5.1.1 Segmented Accuracy 21

5.2 Neural Networks 21

5.3 Feature importance Random Forest Model 23

6 Conclusion **24**

7 Appendix **27**

7.1 Code explanation 27

1 Introduction

People make decisions every single day. When our alarm goes off, we are faced with the question whether to snooze it or not. When we get to the kitchen, we have to decide what to eat. When we get dressed, we have to decide what clothes to wear. All our answers to these questions are based on a certain belief. This belief system has been studied for ages. For many researchers, it was of big interest to quantify the way that we people make decision. This is because the quantification of human decision making helps us understand why we make certain choices. The common belief was first that people make rational decisions as can be seen in Doucouliagos (1994). However, this theory was not always in line with the actual choices that people make and made researchers question whether this theory actually holds.

A major insight into human decision making was introduced by Bernoulli (1738). He resolved an interesting paradox named the St. Petersburg paradox. This paradox portrays that people are only willing to pay a small amount for a risky gamble with an infinite expected monetary value. Bernoulli argued that people do not make choices based on the expected monetary value of something but rather look at expected utility. This resulted in expected utility theory.

This was followed by the development of another widely known theory, namely prospect theory introduced by Kahneman and Tversky (1979). This theory is based on the fact that people value losses differently than gains, which is not incorporated in the expected utility theory.

More recently, we have gained access to more and more data. This gives a lot of possibilities. With this abundance of data it is for example possible to perform machine learning analysis, which is a fairly new field in the world of research but has already grown to be one of the main fields in the modern computing world (Shinde and Shah, 2018). These machine learning techniques can also be applied to human decision making theory. This has been done for example by Peterson et al. (2021). They evaluate several human decision theories on how well they perform at predicting human choices for different lottery gambles. They found that the context-dependent and value-based theory were one of the better predictors in human decision making. These two theories let the model freely determine how people evaluate rewards and probabilities, but the value-based model assumed that people's choices are independent of the context and the context-dependent model did not.

This analysis was done with deep learning. Neural networks were constructed in such way that they are in line with various human decision theories. However, deep learning is just a subset of machine learning itself. There are many other machine learning algorithms that can be applied to different problems. One very prominent model is called the random forest model.

This model is a learning method for classification and regression which operates by constructing a multitude of decision trees at training time. We want to investigate whether the random forest model does a better job at predicting human choices. The advantage of using random forest models instead of neural networks is that they are interpretable. It is possible to extract the most important variables used in the model. This information allows to determine what drives human choices and not just predict them. So, it would be very interesting to find out if random forest models do a good job at predicting human choices, due to their interpretable nature.

Therefore, the goal of this research is to answer the following question:

"Does the random forest model do a better job at predicting human choices on risky gambles than deep learning models?"

To answer this research we use a data set, also used in Peterson et al. (2021), of human decisions for almost 10,000 risky choice problems presented in a format that has been used in previous evaluations of models of decision-making. Each problem consists of 2 options, A and B. In our research, we focus on predicting the probabilities that a person chooses option A and B. We make use of neural networks based on the two best performing models in Peterson et al. (2021). Namely, the context-dependent and the value-based model. Thereafter, we estimate a random forest model in order to compare its performance to our previous models . We have found that the neural network which complies with the context-dependent theory has the highest predictive performance. Second comes the random forest model and third the value-based model.

The paper proceeds as follows: Section 2 reviews the relevant literature. Section 3 describes the data used in our models. Section 4 covers the methodology. Section 5 presents the results of our models and discusses them. Section 6 gives the conclusion of our research and in the Appendix we explain the used code.

2 Literature Review

2.1 Economic theory

A very important decision making theory where our machine learning models build up upon is the previously mentioned expected utility theory introduced by Bernoulli (1738). Expected utility consists of two components. Namely,

1. People use or should use the expected value of the utility of different possible outcomes of their choices as a guide for making decisions.

2. More of the same creates additional utility, but with a decreasing rate.

'Expected value' implies the weighted sum where the weights correspond with the probabilities of different possible outcomes. The first component was introduced by Blaise Pascal and Pierre de Fermat. The second component, even though it was already introduced by Bernoulli in 1738, set its real foot in the door during the marginalist revolution which took place in the late 19th century. In the fields of economics, Dupuit (1844) was the first to derive, from the general concept of decreasing marginal utility, the idea of a decreasing demand function. Besides Dupuit, there were a lot of other researches that have studied this topic extensively such as Jevons, Walras and Menger. These three authors are often also referred to as the most prominent figures in the field of marginalism.

With the theory of diminishing marginal utility they could resolve the Diamond-Water paradox which was described by Adam Smith. This paradox is about the fact it is very surprising that diamonds have a much higher value than water, even though water is of much more practical value to people and is necessary for human survival. Smith used this paradox to support his labor theory of value. Marginalist argued that Smith was in the wrong. They claim that people value items based on the specific uses that people have for each individual unit of a good. So, because the value of each additional unit decreases, the price of goods that are more plentiful will be lower, and the price of goods that are more scarce will be higher. This explains why diamonds are usually more expensive than water.

Combining the two previously mentioned components gives the powerful theory of expected utility as we know it today (Lengwiler, 2009).

Before this expected utility theory was used in the field of human decision making, behaviour was often described by the expected monetary value theory. This theory implies that if a person would have to choose between a 100% chance of receiving 50 euros or a 10% chance of receiving 600 euros, he would always opt for the second option as $1 * 50 < 0.1 * 600$. So the expected value of the second option is higher. However, this theory was a lot of times not in line with the actual choices made by people. This is where the expected utility theory came into play. If we now apply the expected utility theory to the previous example, the choice is between $1 * EU(50)$ and $0.1 * EU(600)$. Diminishing marginal utility could lead to people choosing the first option now.

Expected utility theory does not require people to be self-interested. It is well possible that people rather give money to a homeless person than buying a new pair of shoes. In other words, helping out a homeless person 'yields them more utility' than buying a new pair of shoes.

As mentioned previously, this theory was followed by the development of prospect theory by Kahneman and Tversky (1979). This theory is often also called the loss-aversion theory. It was developed due to inconsistencies found in expected utility theory. The expected utility theory assumes that the perceived utility of gaining 50 euros is of the same magnitude as losing 50 euros. However, in practice this is often not the case. Prospect theory assumes that people are more averse to losses than to gains. This usually results in 2 different utility functions. Namely, a utility function for gains and a utility function for losses. A simple example is the linear loss-averse representation of the utility function which can be seen below

$$u(x) = \begin{cases} x & \text{if } x \geq 0 \\ \lambda x & \text{if } x < 0 \end{cases}.$$

This utility representation allows people to react stronger to losses by setting λ bigger than one. Nowadays, prospect theory is still a widely used theory to explain phenomenons in the field of human decision making. Some intuitions of Kahneman and Tversky (1979) also contain that people focus heavily on the main outcome and neglect the probability of it actually happening.

We see this same phenomena happen with lotteries. When participating in a lottery, the odds are stacked highly against you. However, a lot of people still buy lottery tickets, so it must be that people focus more on the prize that they can win than the probability of it actually happening.

Ariyabuddhiphongs (2011) studied why people exactly buy lotteries. He concluded that people do not behave in a rational way while gambling on a lottery. People do not only join lotteries to 'make money', but entertainment purposes also play a huge role.

The existence of the context-dependent theory has already been proven in previous literature. For example in Kelman et al. (1996). This paper tested for context-independence in legal decision making. They examined whether the relative ranking of two options varies depending with the presence or absence of other options. They indeed found that the assumption of context-independence is violated in legal decision making. Our paper builds on this literature, by bringing the theory of context-dependence into a new field of decision making. Namely, the decision making of humans on risky gambles.

Another question that arises when thinking about context-dependency is what actually drives this context effect. Dhar et al. (2000) examine this exact phenomena. Researchers have identified two different frameworks to explain the existence of context effects.

Firstly, effort minimization. If context effects arise due to effort minimization, then it suggests that the low involvement of customers when making choices should make them more susceptible to context effects. This is because they are not completely focused on making 'rational' choices, thus they are more likely to follow their gut feeling which is vulnerable to all sorts of biases.

Secondly, perceptual contrast. It also could be that the context effect arises due to people taking into account relative characteristics of alternative options. Even though consumers work typically hard to derive at the best possible decision. Over analyzing the alternatives also adds difficulty and complication to the choice.

They investigated the source of context effects by looking at the effect of time pressure. If context effects origin from the preference of people for effort minimization, then this time pressure should increase the magnitude of context effects. On the other hand, if context effects are due to the relational characteristics of the alternatives provided, time pressure should actually reduce the magnitude of context effects.

Dhar et al. (2000) found that, under time pressure, consumers give less weight to comparative features, and, in addition, they also focus more heavily on the positive aspects of the alternatives. Thus, these findings support perceptual contrast being the source of context effects.

Even though this experiment was not conducted with risky gambles, it is fair to believe that if context dependency is actually found in our models, perceptual contrast is also the main driver of this context effect

The existence of context-dependence can alsobe applied to real life situations. Actually, a lot of companies use the phenomena context-dependence to shift consumer's preferences to higher priced items. They do this by adding a 'decoy product' to their product range. The purpose of this decoy product is to make the higher priced item seem as a very good deal. For example, when a cinema sells a small sized popcorn for 3 euros and a large sized popcorn for 7 euros. Adding a medium sized popcorn for 6.50 euros, makes the large popcorn seem as a much better deal. Decoys are usually asymmetrically dominated, which means that they are completely inferior to one option and only partially inferior to the other option. That is why the decoy effect is also sometimes referred to as asymmetric dominance effect. The decoy effect causes people to spend more than they really need. Because of this decoy product people do not tend to make decisions based on what suits them best, but more on what options seems most advantageous to them.

Figure 1: *Illustration of the decoy effect for pop corn at the cinema*

Another study which investigated the decoy effect was done by Simonson (1989). Researchers let participants choose from a set of various products. As expected, when a decoy option was included, people were more likely to choose the higher priced item. Simonson (1989) also added an extra element to this experiment. They also examined how people behaved when they were told that afterwards they had justify their selection. They found that adding this extra element made the decoy effect even stronger. This shows that when people make a choice, their goal is not the pick the correct option, but the goal is rather to justify the outcome of a choice they have already made. A decoy makes people comfortable in their choice by handing a ready-made justification for it.

Decoy products also exploit the aversion that most people have to losses. Research has shown that many people are more averse to lower quality than to higher prices Hendricks (2018). Exactly this is exploited by decoys, as they are usually introduced in order to push people to pick a higher priced item.

The decoy effect is exploited by a lot of big corporations. Josiam and Hobson (1995) for example did research on the use of the decoy effect in the travel and tourism industry and they have found that the introduction of decoy packages resulted in some consumers shifting their preferences to higher priced packages.

As can be read in Wible (2011), the psychologist Dan Ariely also conducted an experiment on the decoy effect.The news paper The Economist had 3 options. $59 for an online subscription, $125 for a print-only subscription, and finally, $125 for both print and online access. He let his students pick one of these options and 16% chose the cheapest subscription whilst 84% picked the print and online access version. Afterwards, he also let another group of students pick a subscription, but this time he removed the middle (decoy) option. The results were drastically different, 68% of students now chose the 59$ subscription whilst only 32% opted for the more expensive one.

All in all, it is pretty clear how the existence of the context-dependent theory in human decision making is used in the modern day and age.

2.2 Machine Learning

In this paper we heavily focus on the implementation of machine learning models, specifically deep learning and random forest models. Machine learning is a branch of computational algorithms that is designed to mimic the human brain as it gradually learns from the data that it receives (El Naqa and Murphy, 2015). Machine learning models have already proven to be very effective. Medeiros et al. (2021) illustrates that machine learning models do a better job at inflation forecasting than its benchmark models. They found that machine learning models with a large number of covariates are systematically more accurate. The random forest model seemed to outperform all machine learning models. Therefore, it is very interesting to see whether to random forest model also thrives in the field of human decision making.

Athey (2018) states that machine learning will have a dramatic impact on economics within a short time frame. A few of the key things that Athey predicts about the impact of machine learning on economics are the development of new econometric methods based on machine learning designed to solve traditional social science estimation tasks, increased emphasis on model robustness and other supplementary analysis to assess credibility of studies and increased use of data analysis in all levels of economics teaching; increase in interdisciplinary data science programs. A lot of the predictions which Athey made are already seen at the moment. So it is fair to say, that economics is becoming more and more important part in the field of economics. That is also what motivates us to perform this research. Human decision making has been studies for ages by general methods, but machine learning has not yet been applied that much in this field.

Deep learning is a specific subset of machine learning which is heavily discussed in LeCun et al. (2015). Deep learning models consist of neural networks which resemble the human brain. They mention that deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. This implies that if we keep feeding the available data to our model, the model will optimize itself. So, if we adjust the models in such way that they comply with the respective human decision theories, deep learning will allow us to determine how well a specific theory performs at predicting the decisions of people.

Our research is also related to Hartford et al. (2016). In this paper they also use deep learning to predict strategic behaviour, but here they focus specifically on human strategic behaviour.

This paper really portrays the power of deep learning. Previous human decision models relied heavily on expert-constructed features and on all kinds of assumptions. Hartford et al. (2016) show that their deep learning model outperforms these previous models. This shows how relevant it is to apply deep learning on human decision making with risky gambles, as it has already been proven to be successful in another field of human decision making.

The research which we are trying to partly mimic is the research done by Peterson et al. (2021). They make us of neural networks to discover human decision theories. Neural networks are hard to interpret, but they tackle this problem by creating lots of neural networks and constrain them in such way that they represent a specific human decision theory. Then, it is possible by analyzing the mean squared prediction error, to determine which human decision theory explains people's choices best. As mentioned previously, they found that the value-based model and the context-dependent model performed best. Thus, we will mimic these 2 models in order to compare their predictive performance to a random forest model.

3 Data

In this paper we make use of the responses to risky choice problems which are contained in a data set named choices13k which is also used in Peterson et al. (2021). Each problem required a participant to choose between 2 gambles, namely gamble A and B. Each gamble consisted of possible rewards with their corresponding probability. For example,

- Gamble A: Reward of 26 with 95% probability or Reward of -1 with 5% probability
- Gamble B: Reward of 21 with 95% probability or Reward of 23 with 5% probability.

Each participant was given 20 risky choice problems and they were paid $0.75 plus a bonus which is proportional to their obtained rewards. After they made a choice, a reward was drawn from the gamble's payoff distribution and the participants were paid accordingly at the end of the experiment. The selections of the participants were all collected and we obtained them from the choices13k data set.

From this data set we only extracted specific variables, as not all variables have an impact on our research. The variables that we extracted are:

- problem: which problem the variables belonged to
- A_i: The i-th reward of gamble A
- pA_i: the probability to obtain the i-th reward of gamble A

- B_i: The i-th reward of gamble B
- pB_i: the probability to obtain the i-th reward of gamble B
- $bRate$: the fraction of people that selected gamble B.

We use $bRate$ as a proxy for the probability that a person chooses gamble B ($P(B)$). By subtracting $bRate$ of 1 we get the fraction of people that selected gamble A and we use this as a proxy for the probability that a person chooses gamble A ($P(A)$).

The variables which we ignored are $Feedback$, n, $Block$, $LotShape$, $LotNumB$, Amb and $Corr$ as these variables are not relevant for our research.

4 Methodology

The goal of our methods is to predict $P(A)$ and $P(B)$, that is the probability that a participant chooses option A and B respectively. We estimate these probabilities with deep learning models and with a random forest model. With our deep learning models we make use of neural networks to implement two economic decision theories, namely the context-dependent theory and the value-based theory. With the random forest model we are not able to constrain our model to a specific economic theory. Thus, this model is solely based on our most flexible context-dependent theory which takes all together as input. Our estimates of $P(A)$ and $P(B)$ will then be compared to the true $P(A)$ and $P(B)$ by means of a mean squared prediction error. Based on this value, we can determine which model performs best at predicting human choices on risky gambles. To train our models we make use of a training set, this set contains 90% of our data. To analyze the performance of our model, we evaluate the predictive performance on a test set which consists of the remaining 10% of our data. We split the data into these 2 sets, because this allows us to do out-of-sample forecasts, which give a better reflection of the predictive performance of our models.

4.1 Deep learning

Deep learning is a subset of machine learning which makes use of neural networks. Neural networks are a set of algorithms which work similarly as the human brain. The human brain consists of billions of neurons which are sensitive to inputs from the external world. After receiving similar inputs for some time, the human brain trains itself and will recognize patterns. Deep neural networks work similarly. In this paper we make use of deep neural networks to estimate $P(A)$ and $P(B)$.

A deep neural network consists of several layers, where each layer consists of nodes. A node receives input from the data and transforms this input based on a set of weights and a bias. After transforming the input, the node passes the information through a so-called activation function. This function determines whether a node should pass-on through the network and ultimately affect the outcome of the model. After letting the inputs pass through the whole neural network, the model will give an output. This is the estimated value of what it is trying to estimate. The estimation error is calculated by comparing it to the actual value. The neural network tries to minimize this error, so after an iteration, it calculates how it should adjust its weights and biases by means of gradient descent

Gradient descent is an optimization algorithm used to minimize some function by iteratively moving in the direction of steepest descent as defined by the negative of the gradient. So at each point, the algorithm calculates the gradient, scales it by a learning rate, and subtracts the obtained gradient from the current positive. The subtraction corresponds with a step towards the minimum of a function. Mathematically this looks as follows

$$C_{n+1} = C_n - \eta * \nabla f(C_n),$$

where η is the learning rate of the algorithm, $\nabla f(C_n)$ the gradient, C_n the current value of the error and C_{n+1} the new value of the error after the biggest possible decrease was made.

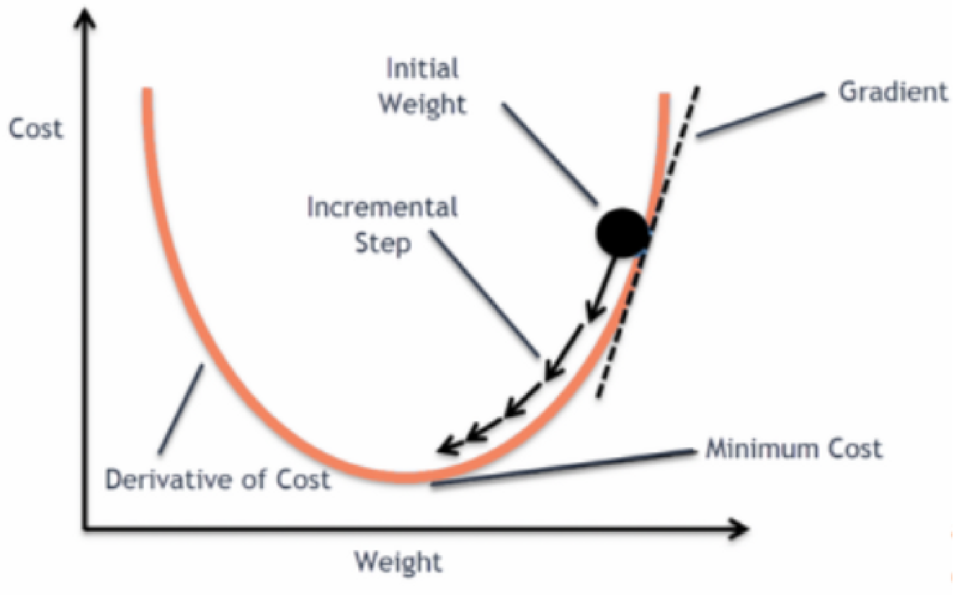

Figure 2: *Illustration of the gradient descent algorithm*

Source: `http://blog.clairvoyantsoft.com/`

So every time that the algorithm moves in the direction of steepest descent, the weights and biases are adjusted his procedure is called backpropagation. After the weights and biases have been adjusted, new input data is fed to the model and the weights will then be adjusted again accordingly. A visualization of a neural network can be seen in the figure below.

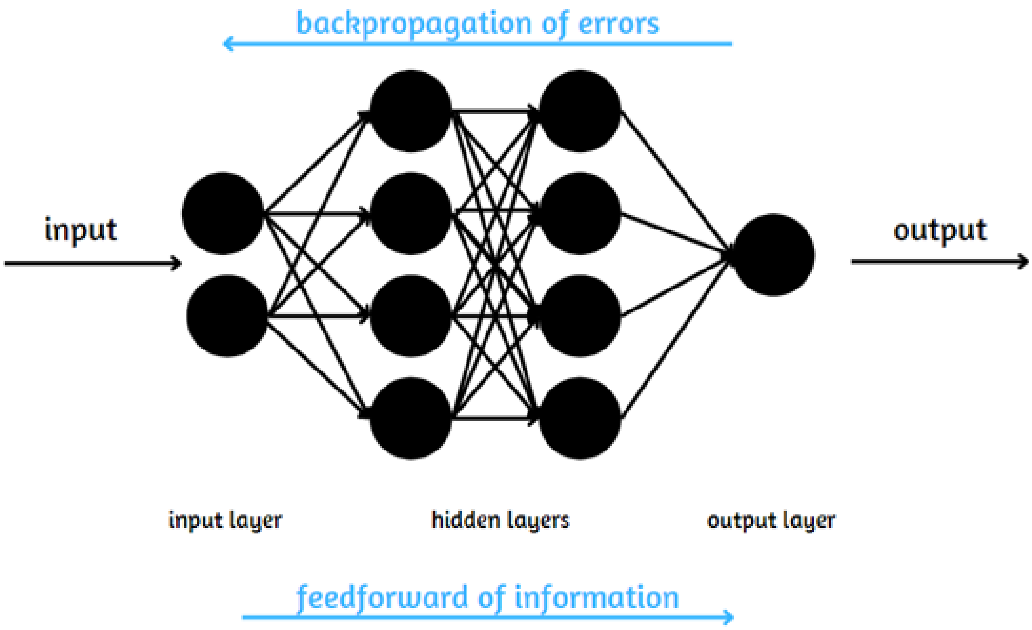

Figure 3: *Basic representation of a neural network*

4.2 Context-Dependent theory

The first theory that we will use to predict human decision making is the context-dependent (CD) theory. This theory is based on the assumption that when valuing a gamble, people also take the other gamble into account. So the value of gamble A also depends on the content of gamble B. This theory contradicts most economic theories which assume that the value that people assign to a certain gamble does not depend on the context, but instead people value each gamble independently from each other. The CD model is the most flexible model as there are no restrictions and it lets the model freely determine the impact of the rewards and their associated probabilities.

For the structure of the neural network, we follow the previous research that has been done by Peterson et al. (2021). Namely, we construct a neural network which takes all information about both gambles as inputs and consists of 2 hidden layers of 32 neurons.

Formally, we have a neural network g that outputs $P(A)$ and $P(B)$ such that

$$(P(A), P(B)) = g(x_A, p_A, x_B, p_B),$$

where x_i is a vector containing the rewards of options i and p_i is a vector containing its respective probabilities. The figure below portrays a visualization of the CD model.

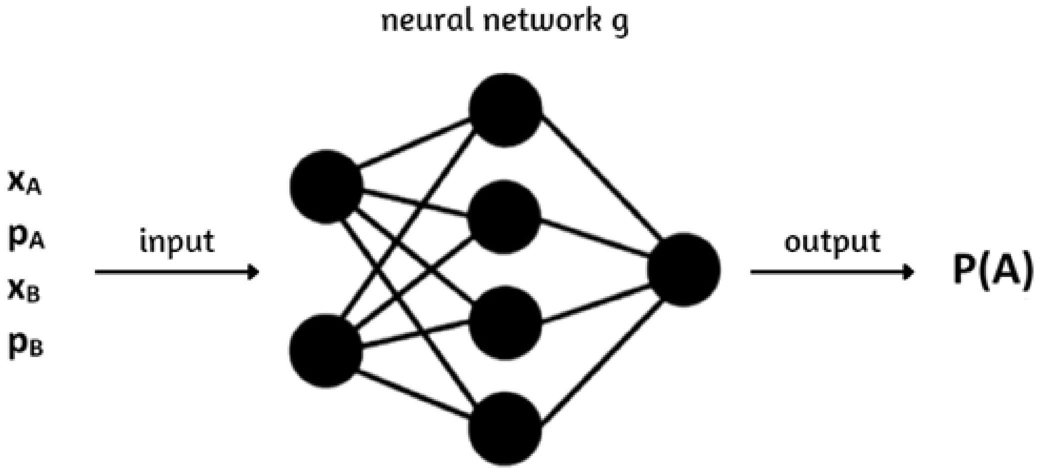

Figure 4: *Visualization of the context-dependent model*

4.3 Value-Based theory

The value-based (VB) theory is a subset of the context-dependent theory. This theory has the assumption of context-independence. This implies that when people value gamble A, they do not take into account the content of gamble B. Just as with the CD model, this model is completely free in determining how the rewards and their associated probabilities have an effect of $P(A)$.

For this model, we again follow the model proposed by Peterson et al. (2021). So, we have a big neural network g which consists two sub-networks f. These sub-networks consist of 1 hidden layer with 64 neurons. The first sub-network takes x_A and p_A as input, and the second sub-network takes x_B and p_B as input. The networks put out $V(A)$ and $V(B)$ respectively, which portrays the value that a person assigns to gamble A and B. These 2 values are then forwarded to a softmax function, which transforms these values into $P(A)$ and $P(B)$. The function looks follows

$$P(A) = \frac{e^{V(A)}}{e^{V(A)} + e^{V(B)}}$$

and

$$P(B) = \frac{e^{V(B)}}{e^{V(A)} + e^{V(B)}}.$$

An epoch in deep learning means training the neural network with all the training data for one cycle. Peterson et al. (2021) runs their models for 200 epochs. We will follow the same procedure. However, as our models are slightly different from theirs, we will also run the neural networks for 500 epochs, for both models, to make sure that they converge. Also, we will use the first 90% of our data to estimate our models, and the remainder to examine their predictive performance

The figure below portrays a visualization of the VB model.

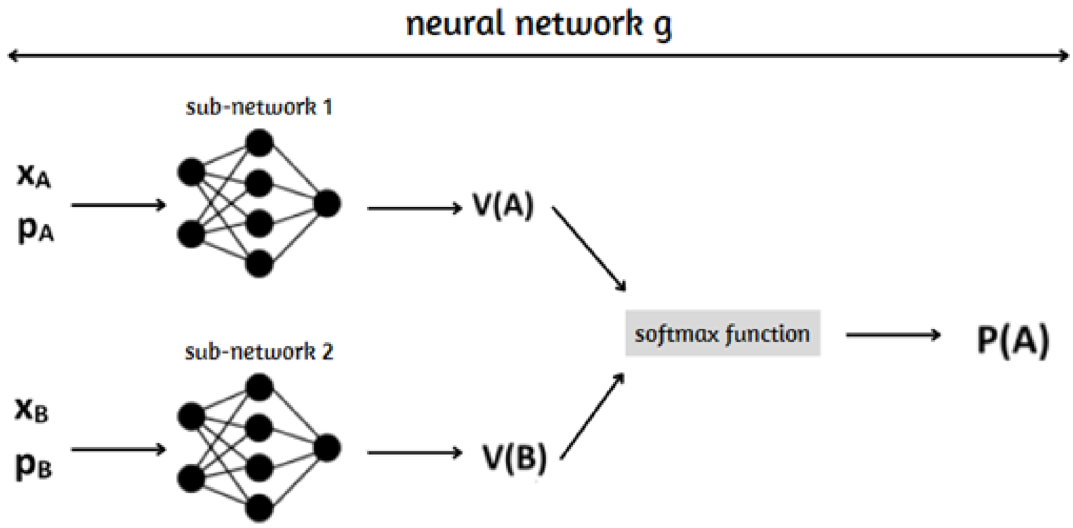

Figure 5: *Visualization of the value-based model*

4.4 Random Forest Model

Random forest is a machine learning algorithm which is widely used for classification and regression problems. It makes use of decision trees. In this paper we use a random forest regression model. Thus, we make use of decision tree regression models. A decision tree is a model which asks questions to the data narrowing our possible values until it gets confident enough to make an appropriate prediction. It consists of decision nodes and leaf nodes. On a decision node a question is asked to split the data. On a leaf node the decision tree stops and makes a prediction of what it thinks that y should be. The questions are all in true or false form. For example, if our data consists of numeric values x_1, x_2 and x_3 which are used to predict y. The questions would all look something like: $x_1 < 1, x_2 > 5$ and $x_3 > 2$. It is possible to extract which variables play the most important role in the decision trees, this makes these models more interpretable compared to neural networks. This procedure is explained in section 4.5.1.

In this paper, the decision trees decide which questions it asks based on variance reduction. A decision node goes over all possible questions and determines the variance of its child nodes. Then the weighted average variance of both child nodes is calculated and it chooses which specific question gives the lowest variance in its child nodes. This process goes on and on until the ends of the tree consists of just leaf nodes. The figure below gives a visualization of a simple decision tree.

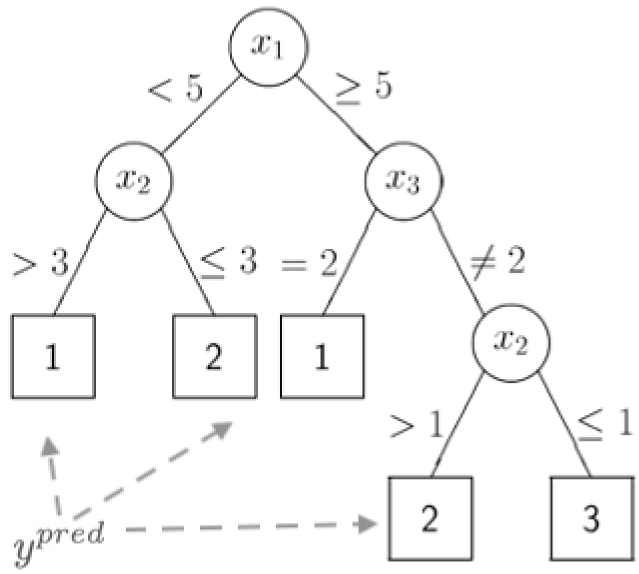

Figure 6: *Basic representation of a decision tree model*
Source: `http://bookdown.org/`

There is one downside to a decision tree model. Its structure is heavily dependent on the data that it is fed. Therefore, it can have a very high variance with out-of-sample estimates. The random forest model is often found to have a better predictive performance than a decision tree (Banfield et al., 2006). This model consists of multiple decision trees. It works as follows:

1. It first selects a random sub-sample of a data set
2. It uses this sub-sample to estimate a decision tree
3. After being told how many trees this random forest model should consist of, it repeats step 1 and 2 N times.
4. For a new data point, it lets each of its N trees predict y and it takes the average of all predictions as its estimated value.

The process that happens at step 1 is called bootstrapping. As this model consists of lots of decision tree which are all based on a different data set, this model is much more general which helps with out-of-sample forecasts.

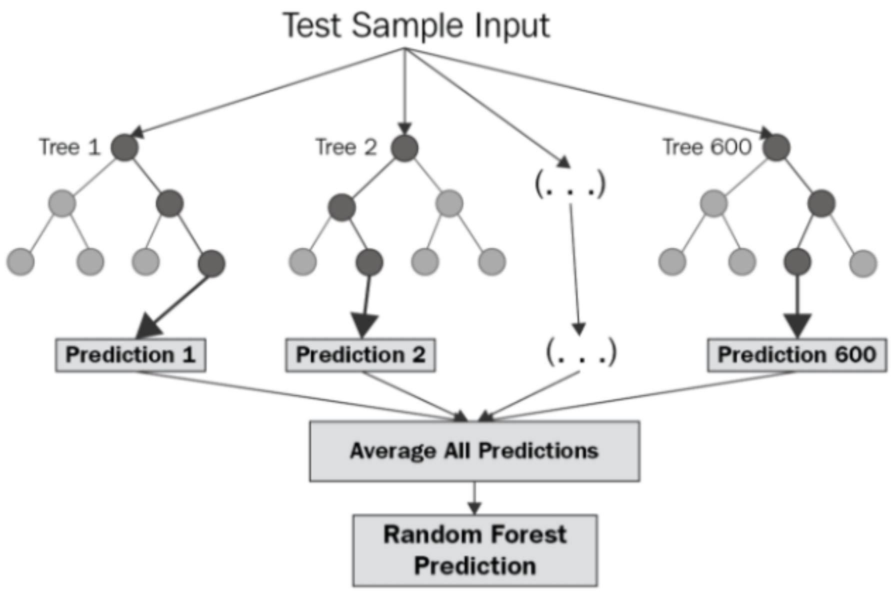

Figure 7: *Basic representation of a Random Forest model*

Source: `http://oreilly.com/`

4.5 Random Forest Regression - Applied

In this paper, we feed the random forest model with all possible outcomes and probabilities of both gambles, thus it is a context dependent model. For the number of trees we use the default amount in the sklearn.ensemble package which is set at 100. The output of this model is both the $P(A)$ and $P(B)$. Also, just as with our previous model, we use the first 90% of our data to train our model and the remainder to estimate its predictive performance.

4.5.1 Feature importance

In order to understand the model and to see which variables are most important, we analyse the feature importances of the model. Feature importances are computed as the mean accumulation of the impurity decrease within each tree. By comparing the values of each variable, we can see which variables are most important for the model. The interpretability of the random forest model is the major advantage of the random forest model compared to the neural networks.

4.6 Model Evaluation

4.6.1 Mean Squared Error

To evaluate the predictive performance of our human decision models we compare the mean squared prediction error of each model of their output. Specifically, this can be computed as

$$\mathcal{L}_{MSE} = \frac{1}{N} \sum_{n} \left[\hat{P}_n(A) - P_n(A) \right]^2,$$

where n ranges over the different risky choice problems, N is equal to the total amount of problems, $\hat{P}_n(A)$ is equal to the estimated probability of choosing A in problem n, and $P_n(A)$ represents the actual proportion of people choosing A in problem n.

4.6.2 Segmented Accuracy

It is also very interesting to inspect for what kind of gambles our model performs best at predicting. We split the gambles into 3 segments:

1. Extreme gambles
2. Similar gambles
3. Other

Extreme gambles represent the gambles where people expressed a strong preference for gamble A or B, we specify this by putting all gambles, where the fraction of people who chose A is smaller than 0.2 or bigger than 0.8, into this segment. Similar gambles represent the gambles where the preference of people in favor of a gamble was close to indifferent. We specify this by putting all gambles, where the fraction of people who chose A is between 0.4 and 0.6, into this segment. Other is all the left over gambles. So, the gambles where the fraction of people who chose A is between 0.2-0.4 or 0.4-0.6.

To determine the accuracy of each segment we calculate the mean squared prediction error as explained in section 4.6.1.

4.6.3 Number of outcomes

Another aspect of the gambles which could have an impact on people's choices is the amount of outcomes a gamble consists of. Gamble A always consists of 2 outcomes. The number of outcomes in gamble B can range from 2 until 8. A high number of outcomes could have an impact on a gamble's perception. Therefore, for each of our 3 models, we make a second model which consists 2 extra variables. Namely, the number of outcomes in gamble A, and the number of outcomes in gamble B. The mean squared error of these models is then compared to our original models to see whether there is a significant difference between them.

5 Results & Discussion

5.1 Predictive performance

To compare the predictive performance of our various models we have to look at the mean squared prediction errors which can be found in table 1. Here we see that the results are quite similar for both 200 and 500 epochs. We see that the mean squared error improved a bit when we increased the amount of epochs, but one model did not suddenly pass the other one in predictive performance.

We find that the context-dependent neural network model has the best results in terms of predictive performance with a mean squared error of 0.0123. Our random forest model comes second with a mean squared error of 0.0156 and last we have the value-based neural network model with a mean squared error of 0.021. The neural network values are drawn from the model that ran for 500 epochs. Thus, the findings of Medeiros et al. (2021) that the random forest model is the best performing machine learning model can not be transferred to the field of human decision making on risky gambles.

Table 1: Predictive performance of models for out-of-sample forecasts

	200 epochs		
	Random Forest	Context-Dependent	Value-Based
Mean Squared Error	0.0156	0.0130	0.0230
	500 epochs		
	Random Forest	Context-Dependent	Value-Based
Mean Squared Error	0.0156	0.0123	0.0208

	Model containing number of outcomes		
	Random Forest	Context-Dependent	Value-Based
Mean Squared Error	0.0154	0.0127	0.0207

This different result could be caused by the fact that for risky gambles the data is structured completely different than the data used in Medeiros et al. (2021). Here we are dealing with all sorts of gamble rewards and their respective probabilities. Also, all the probabilities are related with each other, because when you only have 2 outcomes with probabilities p and q, we know that $p = 1 - q$. This data is clearly much different than the explanatory variables and the lagged variables used in the inflation model of Medeiros et al. (2021), which is a time series model.

Adding the the number of outcomes in a gamble to the model does not have a significant impact on the predictive performance of our models. Looking at table 1, we find that the mean squared errors are very similar, 0.0154 vs 0.0156, 0.0127 vs 0.0123, and 0.0207 vs 0.0208. Also when looking at the feature importances, in figure 8, of our random forest model containing the number of outcomes. We find that the number of outcomes in gamble A (nA), the number of outcomes in gamble b (nB) are not of high importance.

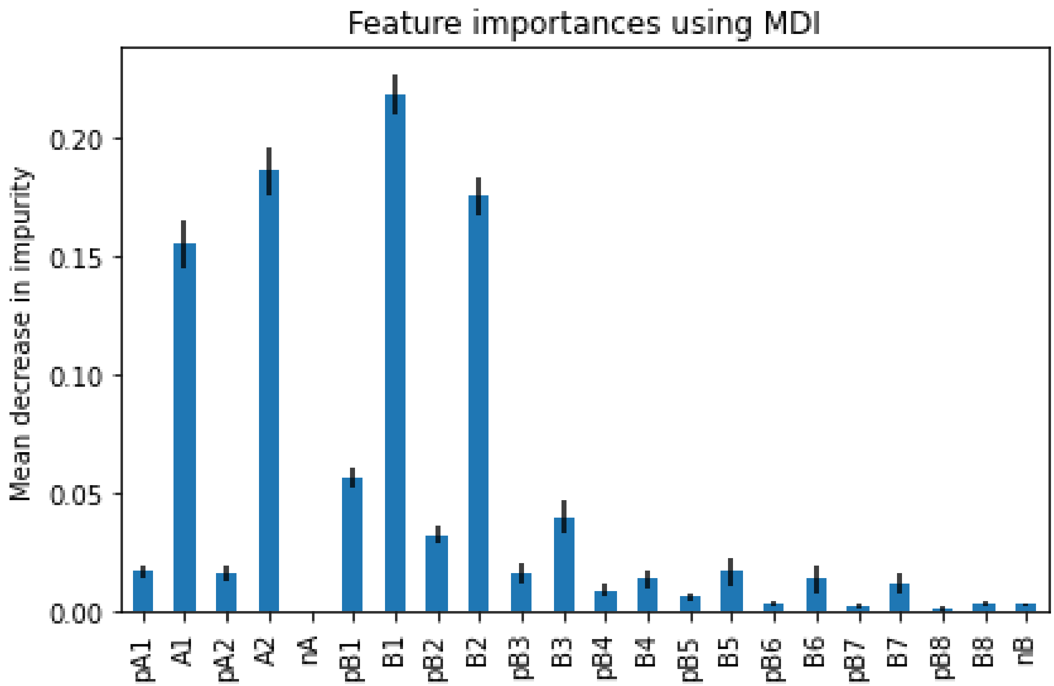

Figure 8: *Feature importances of all the variables in the random forest model including the number of outcomes*

5.1.1 Segmented Accuracy

Looking at table 2 we see that for all three models, the best performance can be found with similar gambles. Especially for the value-based model, we see a steep decline in performance for extreme gambles. A probable explanation for this is that the gambles mostly consist of the segment "Similar" and "Other". So, when a model is not properly fitted, it could be harder for the model to predict extreme values. This results in a lower predictive accuracy for these possibly "not properly fitted" models.

Looking at these results and at the Mean squared errors presented in section 5.1, it is fair to say that context-dependent model is undoubtedly the best performing model.

Table 2: Predictive performance of models for out-of-sample forecasts

	Extreme Gambles	Similar Gambles	Other
N	310	419	648
		Random Forest Model	
Mean squared error	0.0199	0.00886	0.0177
		Context-Dependent Model	
Mean squared error	0.0112	0.00927	0.0142
		Value-Based Model	
Mean squared error	0.0388	0.0128	0.0169

5.2 Neural Networks

Figure 9 and 10 represent the model loss of our neural networks based on the training set. We we can see that for that for both amount of epochs the value-based model starts of with the lowest loss, but after a couple iterations this model is quickly surpassed by the context-dependent model. This shows that the the context-dependent theory does a better job at explaining human behaviour than the value-based theory. We can draw the same conclusion when looking at the mean squared error of the predictions of both models given in table 1 in Section 5.1. These results are also in line with the findings in Peterson et al. (2021). So, humans indeed do not value gambles independently. Instead, the value of something does also depend on the context.

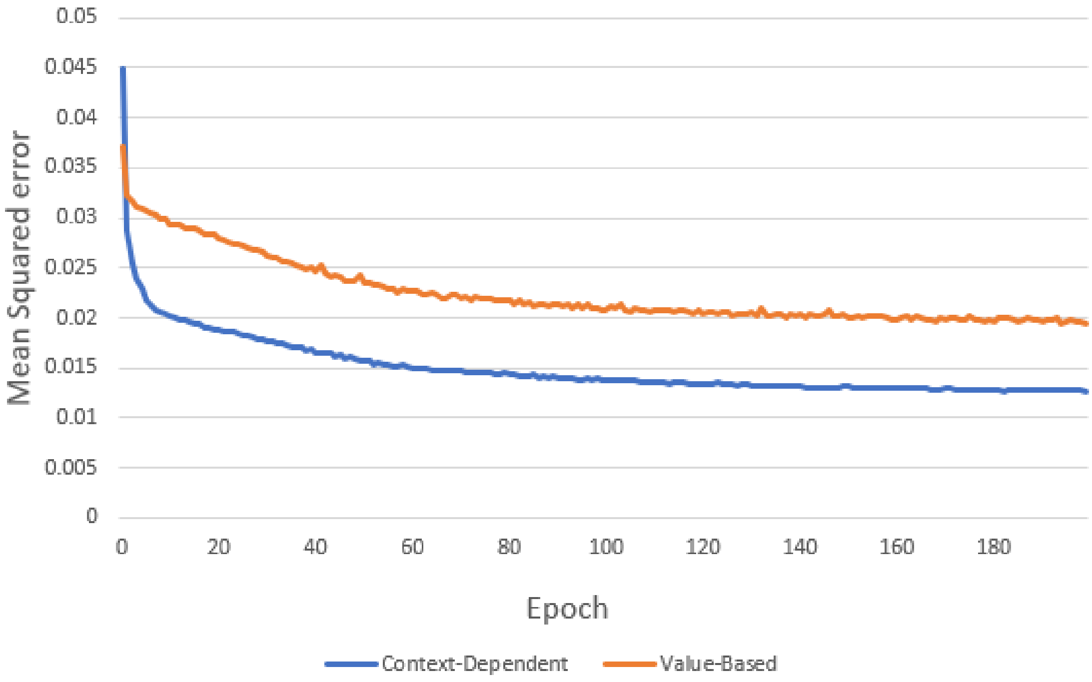

Figure 9: *Mean squared error of the neural networks overtime for 200 epochs*

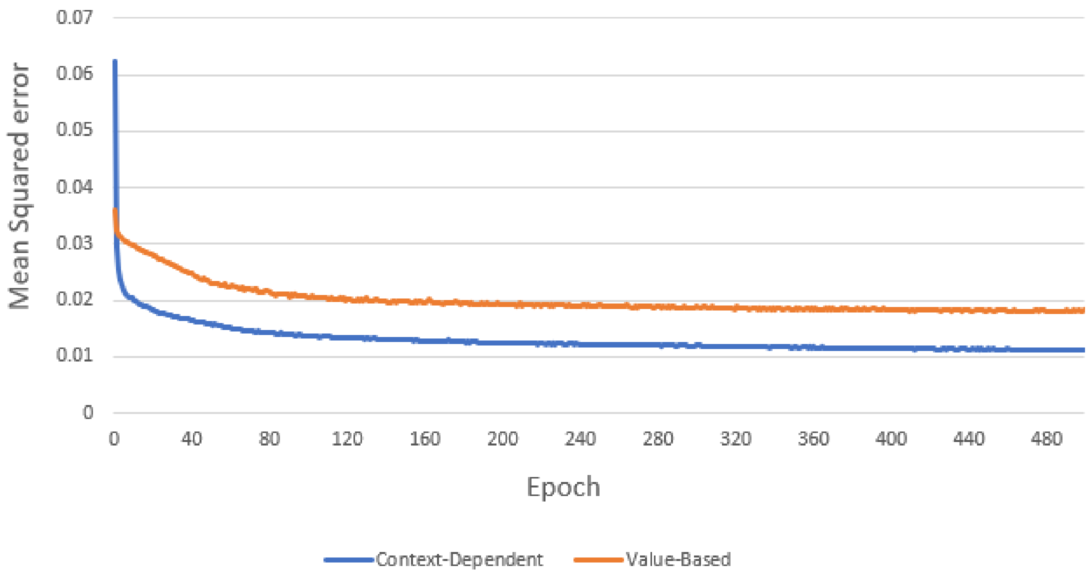

Figure 10: *Mean squared error of the neural networks overtime for 500 epochs*

So, as we can see in our results, when people are valuing a gamble, they are also taking the characteristics of the other gamble into account. Thus, they are not valuing each gamble independently.

It is interesting to see that the value-based model starts off with the lowest loss. A probable explanation for this is that the value-based model is slightly restricted, thus the model has already some idea what the network should look like in the beginning in comparison to the completely flexible context-dependent model. The context-dependent model starts of freshly with no reference point so it is expected that this model will also perform worse in the beginning, but after a couple iterations it catches up as the complete flexibility also allows the model to reflect the true theory behind human decision making.

5.3 Feature importance Random Forest Model

Looking at figure 11, we see that the values of the outcomes are much more important than their respective probabilities. This implies that people are much more focused on the money that they can possibly win than on the chance of it actually happening. Which is in line with some intuitions of Kahneman and Tversky (1979).

An explanation for this could be the explanation given in Ariyabuddhiphongs (2011). It might be that the people, who participated in the risky gamble survey where our data is derived from, also saw choosing between gamble A and B more as a fun game. Which could explain why our random forest model puts so little attention to the probabilities of the gambles, as the surveyed people were more focused on the money that they can win.

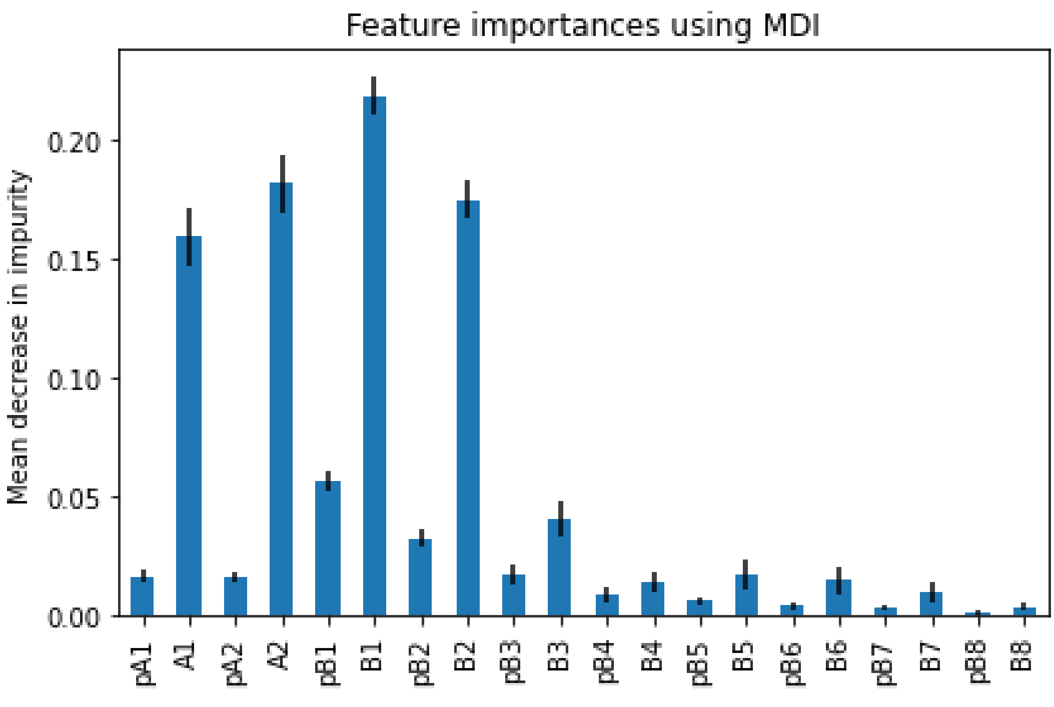

Figure 11: *Feature importances of all the variables in the random forest model*

6 Conclusion

In this paper we aimed to answer the following research question:

"Does the random forest model do a better job at predicting human choices on risky gambles than deep learning models?"

Looking at the mean squared errors of our out-of-sample predictions, we can safely say that the most flexible deep learning model, which took context dependency into account, does a better job at predicting human choices on risky gambles than the context-dependent random forest model. On the contrary, the value-based model performed worse.

The downside of the random forest model is that you can not specify the model in such way to test different human decision theories. The model consists of lots of decision trees which can not be structured in such way to comply with different theories. Thus, the only theory which we could implement with this model is the most flexible context-dependent theory which takes both gambles together as input.

On the contrary, neural networks allowed us to structure them in such way to represent specific human decision theories. As mentioned previously, the value-based model performed worse than the context-dependent model. So, we can conclude that context dependency is a key feature of human decision making.

After computing the feature importances of the random forest model, we noticed the values of the outcomes are much more important than their respective probabilities. From this we can conclude that people are much more focused on the money that they can possibly win than on the chance of it actually happening. Also, we found that the number of outcomes a gamble consists of does not have a significant impact on people's choices.

Unfortunately, due to the black box nature of neural networks, we were not able to interpret these models. So, we could compare the predictive performance of different models, but we could not interpret the importance of various inputs of a specific model. More recently, there has been some research done on the interpretation of neural networks. Zhang et al. (2018) for example found a way to interpret neural networks in clinical applications. For future research, it would be very interesting to not just study the predictive performance of these neural networks in human decision making, but also to try to understand which variables really drive these models.

References

Ariyabuddhiphongs, V. (2011). Lottery gambling: A review. *Journal of Gambling Studies*, 27(1):15–33.

Athey, S. (2018). The impact of machine learning on economics. In *The economics of artificial intelligence: An agenda*, pages 507–547. University of Chicago Press.

Banfield, R. E., Hall, L. O., Bowyer, K. W., and Kegelmeyer, W. P. (2006). A comparison of decision tree ensemble creation techniques. *IEEE transactions on pattern analysis and machine intelligence*, 29(1):173–180.

Bernoulli (1738). Evolution and economics under risk.

Dhar, R., Nowlis, S. M., and Sherman, S. J. (2000). Trying hard or hardly trying: An analysis of context effects in choice. *Journal of Consumer Psychology*, 9(4):189–200.

Doucouliagos, C. (1994). A note on the evolution of homo economicus. *Journal of Economic Issues*, 28(3):877–883.

Dupuit, J. (1844). On the measurement of the utility of public works. *International Economic Papers*, 2(1952):83–110.

El Naqa, I. and Murphy, M. J. (2015). What is machine learning? In *machine learning in radiation oncology*, pages 3–11. Springer.

Hartford, J. S., Wright, J. R., and Leyton-Brown, K. (2016). Deep learning for predicting human strategic behavior. *Advances in neural information processing systems*, 29.

Hendricks, K. (2018). The decoy effect: Why you make irrational choices every day (without even knowing it). *Kent Hendricks*, 7.

Josiam, B. M. and Hobson, J. P. (1995). Consumer choice in context: the decoy effect in travel and tourism. *Journal of Travel Research*, 34(1):45–50.

Kahneman, D. and Tversky, A. (1979). On the interpretation of intuitive probability: A reply to jonathan cohen.

Kelman, M., Rottenstreich, Y., and Tversky, A. (1996). Context-dependence in legal decision making. *The Journal of Legal Studies*, 25(2):287–318.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436–444.

Lengwiler, Y. (2009). The origins of expected utility theory. In *Vinzenz Bronzin's Option Pricing Models*, pages 535–545. Springer.

Medeiros, M. C., Vasconcelos, G. F., Veiga, Á., and Zilberman, E. (2021). Forecasting inflation in a data-rich environment: the benefits of machine learning methods. *Journal of Business & Economic Statistics*, 39(1):98–119.

Peterson, J. C., Bourgin, D. D., Agrawal, M., Reichman, D., and Griffiths, T. L. (2021). Using large-scale experiments and machine learning to discover theories of human decision-making. *Science*, 372(6547):1209–1214.

Shinde, P. P. and Shah, S. (2018). A review of machine learning and deep learning applications. In *2018 Fourth international conference on computing communication control and automation (ICCUBEA)*, pages 1–6. IEEE.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of consumer research*, 16(2):158–174.

Wible, A. (2011). It's all on sale: Marketing ethics and the perpetually fooled. *Journal of business ethics*, 99(1):17–21.

Zhang, Z., Beck, M. W., Winkler, D. A., Huang, B., Sibanda, W., Goyal, H., et al. (2018). Opening the black box of neural networks: methods for interpreting neural network models in clinical applications. *Annals of translational medicine*, 6(11).

7 Appendix

7.1 Code explanation

The data and code which we used can be found in the provided zip file. All models are created in python. For each model we imported the data from an excel file, split the data into a test and training set, and finally fed this data to our machine learning models. Then we let our trained models predict $P(A)$ and evaluated their performance by means of their mean squared error.