Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

Elsie Backen Tandberg

Modeling of decisions made by test persons in a dice task

Master's thesis in Applied Physics and Mathematics
Supervisor: Håkon Tjelmeland
June 2022

NTNU
Norwegian University of Science and Technology

Elsie Backen Tandberg

Modeling of decisions made by test persons in a dice task

Master's thesis in Applied Physics and Mathematics
Supervisor: Håkon Tjelmeland
June 2022

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

NTNU
Kunnskap for en bedre verden

Preface

This master's thesis is the final work of a five year long master's degree programme in Applied Physics and Mathematics at the Norwegian University of Science and Technology. My main profile has been industrial mathematics, where I have specialized in statistics. This thesis considers the connection between schizophrenia and decision making. The prospect of using statistics to model decision making in a real life problem is what initially drew me to this topic.

I would like to thank my supervisor Håkon Tjelmeland who has had weekly meetings with me for almost a year, and has been a great resource. His thorough feedback has been highly valuable, and incredibly helpful. I have learnt very much from working closely with him. I would like to give thanks to Gerit Pfuhl for introducing me to the psychological aspect of the task, and for designing the dice task. This is the task I have used to consider the connection between schizophrenia and decision making. I would also like to thank Kristoffer Klevjer for providing me with the data from the participants of the dice task. I want to thank Fredrik E. Akre for being a great motivation and always helping me sort my thoughts out whenever I have felt stuck. Lastly, I want to thank all my classmates, I would never have been able to complete these five years if it had not been for all the fun we have had along the way.

Abstract

It has been hypothesized that there is a connection between hasty decision making and schizophrenia. Some claim that probabilistic reasoning tasks performed by both healthy controls, and schizophrenic participants have proven this connection. Others suggest that the previously performed tasks can not be used to draw the conclusion.

Professor Gerit Pfuhl has made an alternative task to the tasks previously used to test the hasty decision making in schizophrenic participants. This alternative task is called the dice task, and the aim of this task is to identify a loaded die among four dice. In this thesis we have data from 212 healthy controls and 41 people diagnosed with schizophrenia who have participated in the dice task. This data will be used to consider the participants decisions, and to investigate whether there are any differences between the two groups.

The decisions are evaluated by deriving a probability model used to calculate rewards for each decision. Next, the rewards are used to fitting a stochastic model to each participant. The stochastic model contains parameters which are estimated and studied in order to determine whether it is possible to distinguish the two groups of participants. The findings were that the parameter estimates are similar for the two groups. Thus, from the particular model presented in this thesis, used to evaluate the data, we can not conclude that the schizophrenic participants tend to make different decisions compared to the healthy controls.

Sammendrag

Det er antatt å være en sammenheng mellom å ta forhastede avgjørelser og å være diagnostisert med schizofreni. Noen hevder at oppgaver som omhandler sannsynlighetsresonementer utført av både friske og schizofrene deltakere, viser denne sammenhengen. Andre mener at disse oppgavene som tidligere har vært utført, ikke danner et godt nok grunnalg til å trekke den konklusjonen.

Professor Gerit Pfuhl har laget en alternativ oppgave til de oppgavene som tidligere har blitt brukt til å teste forhasted beslutningstakning hos schizofrene. Denne alternative oppgaven er kalt terningoppgaven og går ut på at man skal identifisere en urettferdig terning blandt fire terninger. I denne avhandlingen har vi data fra 212 friske og 41 schizofrene deltakere, som har gjennomført terningoppgaven. Disse dataene vil bli brukt til å betrakte deltakerenes beslutninger, og å undersøke om det er noen forskjell mellom de to gruppene.

Beslutningene evalueres ved å utlede en sannsynlighetsmodel som brukes til å regne ut belønninger for hver beslutning. Belønningene blir så brukt til å tilpasse en stokastisk model til hver deltaker. Den stokastiske modellen inneholder parametere som estimeres og evalueres med hensyn på hvorvidt det er mulig å se forskjell på de to gruppene av deltakere. Fra denne modellen, brukt til å evaluere dataene, kan vi ikke konkludere med at de schizofrene deltakerene tar annerledes beslutninger enn de friske deltakerene.

Contents

1 Introduction **1**

2 The dice task **4**
2.1 Data selection 7

3 Theory for construction of models **7**
3.1 Theorem of total probability 8
3.2 Bayes' theorem 9
3.3 Generalized Linear Models 10
3.4 Maximum Likelihood Estimation 11

4 Evaluating decisions **12**
4.1 Problem setup 13
4.2 Probability model 16
4.3 Stochastic model 28

5 Results **33**
5.1 Variable choices 33
5.2 Intervals considered for α 38
5.3 Analyzing $\hat{\alpha}$ and $\hat{\beta}$ 40

6 Closing Remarks **45**

1 Introduction

How we make decisions in our everyday life, and what reasoning lies behind is complex and complicated to understand. Can our decisions say something about our state of mind, and is it possible to determine based on the decisions a persons makes whether he or she is healthy or is suffering from a psychiatric disorder? This thesis aims to investigate if there exists a connection between patients suffering from schizophrenia, and how they make decisions in a task. This is done by constructing a probability model which gives rewards for each decision, and use this to fit a stochastic model to each participant. Further, the parameters in the stochastic model will be estimated for each participant, and used to evaluate if there is any difference between the the two groups.

Schizophrenia is a mental illness that effects how you think, feel and behave. Psychotic symptoms include hallucinations where you see, hear and experience things that are not there, and delusions which is having strong beliefs that are not true and may seem irrational to others (National Institute of Mental Health 2022). Previously, research on this topic has looked at the connection between schizophrenia and especially deluded patients and whether they are prone to making hasty decision. Hasty decision making is described as a jumping to conclusions (JTC) bias which is defined as hasty decision making based on little evidence (Balzan et al. 2017).

To this date, the probabilistic reasoning task used to investigated the JTC bias has mostly been the beads task (Moritz and Woodward 2004). In the beads task participants are presented two jars each containing different proportions of colored beads. Typically jar A contains 80% black beads and 20% white beads, whereas jar B contains the opposite amount, 20% black beads and 80% white beads. The jars are then removed from the participants view, and the participant is told that beads will be drawn from one jar at a time with replacement, such that the proportions remain unchanged. The task is for the participants to decide which of the two jars the beads are drawn from. We can then compare the performance of the participants by considering how many beads they ask to be drawn prior to concluding if the beads are drawn from jar A or jar B. The results from the beads task is that participants with schizophrenia tend to reach a conclusion after fewer draws compared to healthy controls (Moritz and Woodward 2004).

Despite the consistent findings that schizophrenic participants reach their conclusions after fewer draws than healthy test persons, it is argued that the mechanism contributing to the JTC bias is not clear. Three alternative interpretations of the results are mentioned by Moritz and Woodward (2004). They maintain the hypothesis that there is a JTC bias in schizophrenic participants, but are critical to the interpretation that fewer drawn beads is the driving mechanism. The first alternative interpretation is that a participant with schizophrenia will have a "winner takes it all" mechanism such that after drawing one black bead the participant will decide in favour of jar A. The second interpretation is that schizophrenic participants rule out alternative hypothesis too quickly, so after seeing one black bead they will conclude that it can not come from jar B. The third interpretation is that patients might over-adjust while faced with contradicting evidence, so after drawing three black beads, followed by one white bead, they might choose in favour of jar B because they are unable to retrieve past experiences. In conclusion, Moritz and Woodward (2004) enlighten aspects of the results of the beads task apart from fewer draws than test persons, that might be used to argue that the patients have a JTC bias.

In addition, the conclusion that the results from the beads task is evidence for a JTC bias in schizophrenic participants have been questioned. In Pfuhl and Tjelmeland (2019) they confirm the findings that the number of beads drawn is lower in patients than in healthy controls, but in an asymmetric sequence of beads, and with more even ratios, both groups evaluate more beads. Thus, Pfuhl and Tjelmeland (2019) argue that both groups are sensitive to the cognitive effort required to estimate the probabilities. This contradicts that deluded patients have a JTC bias, because reduced cognitive abilities have been linked to the JTC bias. The findings by Pfuhl and Tjelmeland (2019) shows that the patients do not have reduced cognitive abilities.

There is a need for an alternative task to consider if there actually is a JTC bias in deluded patients, which is more credible than the beads task. In the Master's thesis Skogvang (2021) the box task was considered as an alternative to the beads task. Here, the participants are shown grey boxes, they choose which box to open, and the box shows either the color blue or red. It is known to the participants that one of the two colors, blue or red, are in majority, and the participants task is to determine which color that is. The thesis derives a method for finding an ideal observer solution of the box task and analyses how 76 participants make decisions using a softmax model. The findings were that the model was a good fit for participants who make good choices, but not as good of a fit for participants who made bad choices.

Professor Gerit Pfuhl has developed another alternative to the beads task, which is the dice task. This task gives the participants four dice, and they are told that one of the dice is loaded, meaning that it has a higher probability of showing one of the sides. There are two versions of the dice task, a limited, and an unlimited version. In the limited version the participants have to make exactly a total of ten throws, and for each throw they are free to choose any of the four dice. In the unlimited version, the participants can throw one die at a time, and this die can be thrown as many times as they want. When they decide that they are finished throwing one die, they can choose a new die to throw as many times as they like. However, they can not go back to a previous die and throw it again once they have decided that they are finished with it.

The limited version of the dice task was studied in Tandberg (2021), where an optimal strategy of identifying the loaded die was presented. However, there was no available data from schizophrenic participants for the limited version. Thus this thesis will focus on the unlimited version of the dice task where we have available data from both healthy and schizophrenic participants.

With a focus on evaluating the decisions made by participants in the dice task, we can construct a probabilistic model which assigns rewards for each possible decision. The necessary theory for deriving the probability model and a stochastic model, which will later be used to fit to the data, will be given together with theory for parameter estimation.

The probability model should focus on maximizing the probability of identifying the loaded die. For each throw our model should give two expected rewards, one expected reward if you throw the die one more time and another expected reward if you stop throwing the die. Thus, the choices made by both healthy controls, and schizophrenic patients can be compared to the rewards from our model, and this can be used to evaluate their choices.

A stochastic model will be fitted to the data. The stochastic model contains parameters from the probability model, and the goal is to fit a model for each participant. Lastly, the parameter estimates from the stochastic model will be used to investigate if there are any differences between the two groups who have completed the dice task.

2 The dice task

The dice task is designed as a computer game where each participant is presented four dice of different colors. They are told that one of the four dice is loaded, however they do not know which number of eyes the loaded die will show more frequently. Neither do they know “how loaded” the loaded die is, meaning that they do not know the probability of the loaded die showing the number of eyes that is more probable than the others. The task is given in two rounds, so after completing the first round a second round starts. The second round is designed in the same way as the first round, but it is independent of the first round. So in the second round the loaded die can again be any of the four dice, and the loaded die does not necessarily have the same probability of showing the loaded side as in round one.

The dice task is explained using pictures from a trial run of the task. The participants choose a die to start throwing as shown in Figure 1, and as

Figure 1: Screenshot from the dice task. At the beginning of the task you choose a die by clicking on any of the four dice displayed. In this example the red die was chosen first.

they throw, a bar with the previous outcomes is displayed, this we can see in Figure 2. They throw the die by clicking or swiping the die on the screen. This triggers an animation of the die rolling, which takes roughly one second to complete. After throwing at least once, they can decide if they want to throw the same die again, or stop throwing this die and move on to one of the remaining dice. After each throw they have the same choice where they are asked if they want to continue, or to stop throwing the current die. In the example given in Figures 1 and 2, the red die was chosen first, and it was thrown five times. When they choose to stop throwing a die, they are asked if they think this die is loaded, and if so which number of eyes is the loaded

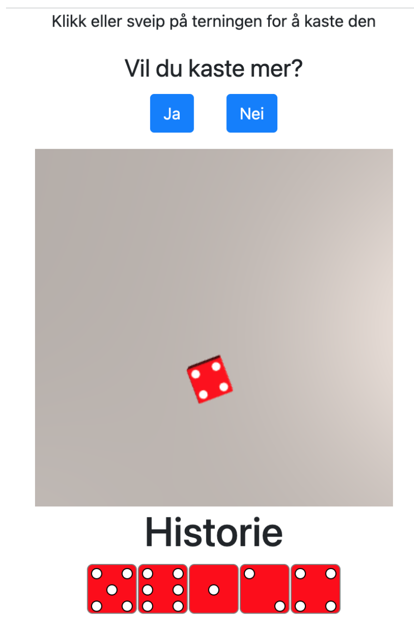

Figure 2: Screenshot from the dice task. Here we see that the red die has been thrown five times, and the history of what has been thrown is displayed underneath. At the top, we are asked if we want to continue throwing or to stop throwing.

Figure 3: Screenshot from the dice task. The red die has been thrown, and we can now choose between the three remaining dice to throw next.

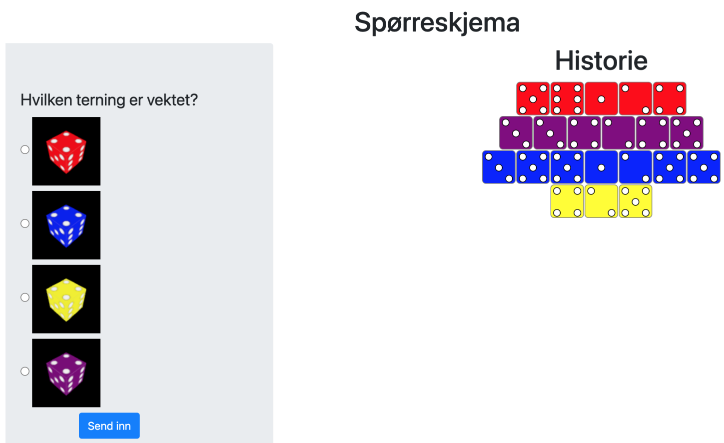

Figure 4: Screenshot from the dice task. All four dice have been thrown, the history is displayed to the right, and we are asked to choose one of the four dice to the left as the die we think is the loaded die.

side. Then they will return to the screen showing the dice that have not yet been thrown, displayed in Figure 3, and pick a new die to throw.

Each die must be thrown at least once. After all dice have been thrown they will get to view the history of what has been thrown on all the dice given in Figure 4. They are asked once again to decide which of the four dice that is loaded and which number of eyes on the chosen die that is more probable. After completing the two rounds, they are not told if they managed to identify the loaded die.

In the design of the experiment it does not matter which color the chosen die has. Independently of what color they choose, the first die will always show the same number of eyes in the same sequence for all participants, the same applies to the second, third and fourth die. This information is not known to the participants, all they know is that one of the four dice is loaded.

However, it will have some consequences that makes the implementation of the model a bit easier. With this design, all participants are shown the same sequences in the same order, so for all of them the third die they choose will be the loaded one in the first round. In the second round, the second die is loaded.

2.1 Data selection

The data that is considered in this thesis is data from 212 psychology students who participated in the dice task in 2021, which are used as healthy controls. There also exists student data collected in 2022, but these are not considered in this thesis. In addition to the data from the 212 students, we consider data from 41 participants who are diagnosed with schizophrenia.

Some of the available data will not be considered in this thesis. Due to time limitations only the data from the first round of the dice task is evaluated. Additionally, participants who have only thrown each die once, which was the minimum requirement, are not considered. This applies to 5 healthy participant, and 3 schizophrenic participants. Furthermore, some outlying results were excluded from consideration, for computational reasons, see Section 5.1. This applies to 22 healthy, and 2 schizophrenic participants. We are then left with considering 185 healthy participant, and 36 schizophrenic participants.

3 Theory for construction of models

The theory needed to construct a probabilistic model which is used to find expected rewards for each decision in the dice task is presented here. The stochastic model which will be fitted to the data from the dice task is also presented in this section. Lastly, the maximum likelihood for parameter estimation is stated.

When deriving a model for calculating rewards for each choice Bayes' theorem, and the law of total probability is needed. The probabilistic model, used to find rewards, is dependent on several parameters. Thus, Bayes' theorem, and the law of total probability will be formulated such that we can include three events. As the stochastic model we choose a generalized linear model, GLM, with binary response, hence GLM with binary regression will be presented here. Included in the GLM we have two parameters which are estimated, thus the theory of estimating parameters using log likelihood will be presented.

3.1 Theorem of total probability

The theorem of total probability will be presented using the definition from Härdle et al. (2015). Further, this definition will be used to expand the theorem of total probability such that it is valid for three events. From Härdle et al. (2015) the theorem of total probability is defined using a partition $A_1, A_2, ..., A_n$, satisfying

- $A_i \neq \emptyset \quad (i = 1, 2, ...n)$
- $A_i \cap A_k = \emptyset \quad (i \neq k; i, k = 1, 2, ..., n)$
- $A_1 \cup A_2 \cup ... \cup A_n = S$.

Then, for any event $B \subset S$, we can use the multiplication rule $P(A_i \cap B) = P(B|A_i)P(A_i)$, to find the law of total probability which is

$$\begin{aligned} P(B) &= P(B \cap A_1) + P(B \cap A_2) + ... + P(B \cap A_n) \\ &= P(B|A_1)P(A_1) + P(B|A_2)P(A_2) + ... + P(B|A_n)P(A_n) \\ &= \sum_{i=1}^{n} P(B|A_i)P(A_i). \end{aligned} \tag{1}$$

Now that we have an equation for the theorem of total probability from (1), this can be used to include a third event. Let the event $C \subset S$, and we are interested in $P(B|C)$. Using the multiplication rule and (1) we have that

$$
\begin{aligned}
P(B|C) &= \frac{P(B\cap C)}{P(C)} \\
&= \frac{1}{P(C)}(P(B\cap C\cap A_1) + P(B\cap C\cap A_2) + ... + P(B\cap C\cap A_n)) \\
&= \frac{1}{P(C)}\sum_{i=1}^{n} P(B|C\cap A_i)P(C\cap A_i) \\
&= \frac{1}{P(C)}\sum_{i=1}^{n} P(B|C\cap A_i)P(A_i|C)P(C) \\
&= \sum_{i=1}^{n} P(B|C\cap A_i)P(A_i|C).
\end{aligned}
$$

3.2 Bayes' theorem

Bayes' rule is defined in Härdle et al. (2015) as follows, let $A_1, A_2, ..., A_n$ be a partition. Then for any event $B \subset S$ with $P(B) > 0$ and given conditional probabilities $P(B|A_1), P(B|A_2), ..., P(B|A_n)$:

$$
P(A_j|B) = \frac{P(B|A_j)P(A_j)}{\sum_{i=1}^{n} P(B|A_i)P(A_i)} \quad \forall j = 1, ..., n. \tag{2}
$$

By using the multiplication rule and (2) we can include a third event. Letting the event $C \subset S$, with $P(C) > 0$, and given conditional probabilities $P(B|A_1\cap C), P(B|A_2\cap C), ..., P(B|A_n\cap C)$ we have

$$
\begin{aligned}
P(A_j|B\cap C) &= \frac{P(A_j\cap B\cap C)}{P(B\cap C)} \\
&= \frac{P(B|A_j\cap C)P(A_j\cap C)}{P(B|C)P(C)} \\
&= \frac{P(B|A_j\cap C)P(A_j|C)P(C)}{P(B|C)P(C)} \\
&= \frac{P(B|A_j\cap C)P(A_j|C)}{P(B|C)}.
\end{aligned}
$$

3.3 Generalized Linear Models

Linear models are well suited for regression analysis when the response variable is continuos and approximately normal. However, when the response variable does not necessarily follow a normal distribution, generalized linear models, GLM, can be used. GLMs still assume that the effect of covariates can be modeled through a linear predictor (Fahrmeir et al. 2013). With the aim of fitting a stochastic model to the data from the dice task, we consider a response variable which is binary. This is because the participants of the dice task have two choices, either to continue throwing the current die, or to stop throwing the current die. Thus, GLMs which unify many regression models, among them binary regression is a reasonable choice of stochastic model for the dice task.

The definition of binary regression models is given in Section 5.1.1 in Fahrmeir et al. (2013), and is given as follows. For the binary regression models we assume data on n individuals are given on the form $(y_i, x_{i1}, ..., x_{ik})$, $i = 1, ..., n$ with binary response $y_i \in \{0, 1\}$ and corresponding covariates denoted $x_{i1}, ..., x_{ik}$. The goal of binary regression is to estimate the effects of the covariates on the probability

$$\pi_i = P(y_i = 1) = E(y_i),$$

for the outcome $y_i = 1$ and given covariates $x_{i1}, ..., x_{ik}$.

Assuming we have a binary response, then the linear probability model

$$\pi_i = \beta_0 + \beta_1 x_{i1} + ... + \beta_k x_{ik},$$

has several disadvantages discussed in Section 2.3 in Fahrmeir et al. (2013). Particularly, the linear predictor

$$\eta_i = \beta_0 + \beta_1 x_{i1} + ... + \beta_k x_{ik} = \mathbf{x'}_i \beta,$$

with $\boldsymbol{\beta} = (\beta_0, \beta_1, ..., \beta_k)'$, and $\mathbf{x}_i = (1, x_{i1}, ..., x_{ik})'$ must lie on the interval $[0, 1]$ for all vectors $\mathbf{x}$. Thus, requiring restrictions on $\boldsymbol{\beta}$ that are difficult to deal with in the estimation process. In order to avoid this, all popular binary regression models combine π_i with the linear predictor η_i through a relation on the form

$$\pi_i = h(\eta_i) = h(\beta_0 + \beta_1 x_{i1} + ... + \beta_k x_{ik}),$$

where h is a strictly monotonically increasing cumulative distribution function on the real line. Now, instead we have that $h(\eta) \in [0, 1]$, which is ensured by the definition of h.

Additionally we can write

$$\eta_i = g(\pi_i),$$

with the inverse function $g = h^{-1}$, we have that h is called the response function, and g is the link function.

The two most widely used binary regression models are the logit and probit models. There are not any large advantages or disadvantages in choosing one of the models over the other. The choice fell on the logit model for the stochastic model to be fitted to the data from the dice task, but the probit model could also have been used.

The logit model, which uses the logistic response function is defined in Section 5.1 in Fahrmeir et al. (2013). It gives the probability

$$\begin{aligned} \pi_i = h(\eta_i) &= \frac{e^{\eta_i}}{1 + e^{\eta_i}} \\ &= \frac{e^{\beta_0 + \beta_1 x_{i1} + ... + \beta_k x_{ik}}}{1 + e^{\beta_0 + \beta_1 x_{i1} + ... + \beta_k x_{ik}}}. \end{aligned} \tag{3}$$

Equivalently we have the logit link function

$$g(\pi_i) = \log\left(\frac{\pi_i}{1 - \pi_i}\right) = \eta_i = \beta_0 + \beta_1 x_{i1} + ... + \beta_k x_{ik}.$$

3.4 Maximum Likelihood Estimation

Maximum likelihood estimation, MLE, is used for parameter estimation. In Fahrmeir et al. (2013) a definition of MLE is given in B.4.1. Let $Y_i, ..., Y_n$ be a random sample with $y_1, ..., y_n$ observations. For discrete $Y_i, ..., Y_n$, we have the joint probability

$$P(Y_1 = y_i, ..., Y_n = y_n | \boldsymbol{\theta}),$$

which is depending on an unknown vector $\boldsymbol{\theta} = [\theta_1, ..., \theta_p]$, which is what we want to estimate. The likelihood $L(\boldsymbol{\theta})$ is a the joint probability as a function of $\boldsymbol{\theta}$,

$$L(\boldsymbol{\theta}) = P(Y_1 = y_i, ..., Y_n = y_n | \boldsymbol{\theta}).$$

The MLE $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ is the value of $\boldsymbol{\theta}$ that maximizes the likelihood $L(\boldsymbol{\theta})$. For technical reasons the logarithm of the likelihood is commonly considered for maximization in stead of the likelihood. The logarithm is a strictly increasing function, so the log likelihood $l(\boldsymbol{\theta}) = \log L(\boldsymbol{\theta})$ attains its maximum at the same value of $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$ as $L(\boldsymbol{\theta})$. Assuming that the random variables $Y_i, ..., Y_n$ are independent, the log likelihood of the random variables is the sum of their joint probabilities,

$$\begin{aligned} l(\boldsymbol{\theta}) &= \log P(y_1|\boldsymbol{\theta}) + ... + \log P(y_n|\boldsymbol{\theta}) \\ &= \sum_{i=1}^{n} \log P(y_i|\boldsymbol{\theta}). \end{aligned} \tag{4}$$

The MLE $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ is the $\boldsymbol{\theta}$-value that maximizes the probability in (4). The $\boldsymbol{\theta}$-value can be found analytically by taking the first derivative of the log likelihood, setting it to zero and solving for $\boldsymbol{\theta}$. It can be found numerically by considering a set of $\boldsymbol{\theta}$-values, and finding which of the $\boldsymbol{\theta}$-values that maximizes the likelihood. It can also be found using profile likelihood where we have more than one parameter. Then we first find the value of one parameter which maximize the likelihood, and then find the value of the other parameter which maximize the likelihood already using the first parameter estimate.

4 Evaluating decisions

With the aim of evaluating the decisions made by participants of the dice task, a probability model for calculating rewards for each decision will be presented. This probability model will be used to find expected rewards for all choices made by each participant, using real data from both healthy and schizophrenic participants of the dice task. When we have the expected rewards for all choices made by the participants, this will be used to fit a stochastic model, which in this case is a GLM with logistic response. While fitting the stochastic model we need parameter estimates both for $\boldsymbol{\beta}$ as described in Section 3.3, and for a cost parameter which is included in the expression for the expected reward. The method for estimating parameters is given at the end of this section.

4.1 Problem setup

The necessary variables for constructing the probability model for finding the expected reward for each choice, and for fitting the stochastic model will be presented.

Variables in the probability model

The goal of the dice task is to identify the loaded die. Hence, we introduce the reward of identifying the loaded die which is equal to 1 if you manage to do so. When performing the dice task the participants repeatedly have two choices, either to continue throwing the current die, or to stop throwing the current die. This gives two rewards, one rewards R_c if you continue to throw, and the other R_s, for the reward when stopping to throw. Thus, each of the two rewards, $R_c \leq 1$ and $R_s \leq 1$, describe how certain you are of having identified the loaded die, both if you continue and stop throwing. Additionally, the rewards aim to maximize the possible reward one can obtain by considering future throws. So we do not consider rewards of choices that have already been made in the past.

When constructing an expression for the reward when continuing to throw a die R_c, it is necessary to include a cost parameter which is denoted by α. The cost parameter is included because the design of the dice task is such that it takes time to roll a die, so logically there is a limit for how patient each participant is, and how many times they are willing to throw. The cost parameter α, will have a negative sign such that it reduces the expected reward of continuing to throw. Furthermore, α will be different for each participant, meaning that R_c will be a function of α, and we will estimate $\hat{\alpha}$ using MLE for each participant.

To express R_c and R_s several variables are needed. The history of the throws that have been made, is used to calculate the probability of having identified the loaded die, thus we need notation for representing the history. This introduces two vectors, the first vector $\boldsymbol{\rho} = [\rho_1, \rho_2, ...\rho_k]$ represents which die that has been thrown up to k throws, so for a throw $i \leq k$, $\rho_i \in \{1, 2, 3, 4\}$. The second vector $\boldsymbol{\omega} = [\omega_1, \omega_2, ..., \omega_k]$ represents what number of eyes each die has shown up to k throws, so for a throw $i \leq k$, $\omega_i \in \{1, 2, 3, 4, 5, 6\}$. Combined, $\boldsymbol{\rho}$ and $\boldsymbol{\omega}$ give k pairs that together say what die has been thrown, and what this die has shown. The length of the two vectors are either the same, or $\boldsymbol{\rho}$ can have length $k + 1$ when $\boldsymbol{\omega}$ has length k. When $\boldsymbol{\rho}$ contains one more element than $\boldsymbol{\omega}$, this means that we know what die will be thrown next, but the die has not yet been thrown, so ω_{k+1} is unknown.

Furthermore, we have the variable representing the true loaded die $Z \in \{1, 2, 3, 4\}$. We assume that when no information of what has been thrown is provided, all dice have the same probability of being the loaded die $P(Z = z) = \frac{1}{4}$ for any $z = \{1, 2, 3, 4\}$.

From the design of the task it is known that only one of the sides on the loaded die Z has a higher probability of facing up than the five other sides. The loaded side on the loaded die Z we denote by $\gamma \in \{1, 2, 3, 4, 5, 6\}$, which indicates how many eyes the loaded side is showing. We assume that without any information of what throws have been made, each side has the same probability of being the loaded side $P(\gamma) = \frac{1}{6}$, for any $\gamma = \{1, 2, 3, 4, 5, 6\}$.

The probability of the loaded side γ showing on the loaded die Z is denoted by p, and it is larger than $\frac{1}{6}$, so this probability is a continuos variable on the interval $p \in \langle \frac{1}{6}, 1 \rangle$. We assume that the probability p has a uniform probability density function on the interval, denoted

$$f(p) = \begin{cases} \frac{1}{1-1/6} & \text{for } \frac{1}{6} < p < 1 \\ 0 & \text{for } p \leq \frac{1}{6} \text{ or } p \geq 1. \end{cases}$$

In the probability model we will have to calculate the probability of making the exact combination of throws that has been observed. On the three fair dice, the probability of a die showing any number of eyes is $\frac{1}{6}$. On the loaded die the probability of throwing γ is equal to p, and the probability of not throwing γ is $\frac{1-p}{5}$. Thus, we have three categories describing what type of throws have been made, a throw on a fair die, a throw on the loaded die Z showing γ, and a throw on the loaded die Z not showing γ. To calculate the probability of having made a combination of throws, we need notation for counting how many of the throws made which belong to each of the three categories. Once this is known, we can use combinatorics to find the probability of having made any combination of throws. The number of throws belonging to the first category are throws made on the three unloaded dice, this is denoted by

$$n_u(\boldsymbol{\rho}) = \sum_{i=1}^{k} I(\rho_i \neq Z).$$

The second category are throws made on the loaded die Z with the loaded side γ showing. The number of throws belonging to this category is given by

$$n_{ll}(\boldsymbol{\rho}, \boldsymbol{\omega}) = \sum_{i=1}^{k} I(\rho_i = Z, \omega_i = \gamma).$$

The third category are throws made of the loaded die Z with other sides than the loaded side showing. The number of throws belonging to this category is given by

$$n_{lu}(\boldsymbol{\rho}, \boldsymbol{\omega}) = \sum_{i=1}^{k} I(\rho_i = Z, \omega_i \neq \gamma).$$

Variables for the stochastic model

When fitting the stochastic model, which in this case is a GLM with binary response, there are two components involved. We have the response vector $y_i \in \{0, 1\}$, and the vector of covariates $\mathbf{x_i} = (x_{i1}, ..., x_{ik})'$ as described in Section 3.3. The connection between the two components and what they represent when adapted to the stochastic model for evaluating decisions, is described here.

For the purpose of evaluating the participants choices, we consider the decisions they are making. In the dice task you can either continue to throw the current die, or stop throwing it. Hence, we have binary decisions which is our response. When a participant chooses to continue throwing the current die at the i'th throw, we have the response $y_i = 1$. If the participant chooses to stop throwing the current die, and moves on to the next one, the response is $y_i = 0$. If the current die is the fourth die, then $y_i = 0$ indicates that the participant chooses to finish the round.

The vector of covariates $\mathbf{x_i} = (x_{i1}, ..., x_{ik})'$ belongs to a decision y_i. The covariates contain information that explains the response, and from our probability model we have rewards for each decision, which should have an effect on the response. The rewards are dependent on what throws have been observed, thus we condition on the history of throws when expressing our covariate x_i. Hence, we define one covariate x_i, which is set to be the difference between the expected reward of continuing to throw the current die $E[R_c|\rho_1, ...\rho_{k+1}, \omega_1, ..., \omega_k]$, and the expected reward of stopping to throw the current die $E[R_s|\rho_1, ...\rho_k, \omega_1, ..., \omega_k]$. Each expected reward is conditioning on the history of throws up to throw number k. Note that the reward when continuing conditions on what die will be thrown next, $k+1$, this is because continuing means that the next throw will be on the current die. So we have $x_i = E[R_c|\rho_1, ...\rho_{k+1}, \omega_1, ..., \omega_k] - E[R_s|\rho_1, ...\rho_k, \omega_1, ..., \omega_k]$, which is found using the probability model. With this definition of x_i, a positive value of x_i means that the best choice is to continue throwing the current die, and consequently a negative value of x_i means that the best choice is to stop throwing this die. With this setup we only have one covariate, so our probability, using notation from (3), is

$$\pi_i = P(y_i = 1) = \frac{e^{\beta_0+\beta_1 x_i}}{1 + e^{\beta_0+\beta_1 x_i}}.$$

As mentioned earlier, R_c is a function of the cost parameter α, which means that x_i also is a function of α. To remember this dependency we write

$$x_i(\alpha) = E[R_c|\rho_1, ...\rho_{k+1}, \omega_1, ..., \omega_k] - E[R_s|\rho_1, ...\rho_k, \omega_1, ..., \omega_k], \tag{5}$$

such that it is easy to see that x_i depends on the cost parameter α.

4.2 Probability model

The model for calculating the rewards, R_s and R_c, after having made any combination of throws will be derived. Afterwards, the expectation of the two rewards $E[R_s|\rho_1, ...\rho_k, \omega_1, ..., \omega_k]$, and $E[R_c|\rho_1, ...\rho_{k+1}, \omega_1, ..., \omega_k]$ will be used to fit the stochastic model. The two models for finding R_s and R_c involves a recursion. Firstly, the initial situation of the recursion will be considered. That is when all throws have been made, so we are considering R_s when throwing the fourth die. Further, this will be used to construct a general equation for both R_s and R_c.

We start by considering R_s when throwing the fourth die. In this situation we have the history of what has been thrown up to throw number k. So R_s is a function of the history, denoted by $R_s(\rho_1, ..., \rho_k = 4, \omega_1, ..., \omega_k)$. The reward indicates how certain you are of having identified the loaded die. In this situation it is possible to calculate four probabilities of each die being the true loaded die as a function of the history, $P(Z = z|\rho_1, ..., \rho_k = 4, \omega_1, ..., \omega_k)$ for $z \in \{1, 2, 3, 4\}$. The highest one of the four probabilities is the die you should guess is the true loaded die, and this probability also represents how certain you are of having identified the loaded die. Hence, we have the following expression for R_s when considering the fourth die,

$$R_s(\rho_1, ..., \rho_k = 4, \omega_1, ..., \omega_k) \\ = \max_{z \in \{1,4\}} \{P(Z = z | \rho_1, ..., \rho_k = 4, \omega_1, ..., \omega_k)\}. \tag{6}$$

When considering R_s when throwing die 1, 2 or 3, the expression for the reward when stopping is slightly different from (6). This is due to the design of the task, which is constructed such that each die must be thrown at least one time. When stopping to throw die 1, 2 or 3, $\rho_k = \{1, 2, 3\}$, you must move on and throw the next die at least once. The reward when stopping is then equivalent to the reward of deciding to continue throwing the next die, where $\rho_{k+1} = \rho_k + 1$. The reward when continuing, R_c, also includes the history of what has been thrown, so we have $R_c(\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$. Additionally, R_c is dependent on ρ_{k+1} because when you choose to continue throwing, you know what die you are going to throw, but not what this die will show. The equation for the reward when stopping to throw die 1, 2 or 3 is

$$R_s(\rho_1, ..., \rho_k = \{1, 2, 3\}, \omega_1, ..., \omega_k) \\ = R_c(\rho_1, ..., \rho_{k+1} = \rho_k + 1, \omega_1, ..., \omega_k). \tag{7}$$

The general expression for the reward when stopping involves two equations. We use (6) when $\rho_k = 4$ and (7) when $\rho_k = \{1, 2, 3\}$, thus we have

$$R_s(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\ = \begin{cases} \max_{z \in \{1,4\}} P(Z = z | \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) & \text{for } \rho_k = 4 \\ R_c(\rho_1, ..., \rho_{k+1} = \rho_k + 1, \omega_1, ..., \omega_k) & \text{for } \rho_k = \{1, 2, 3\}. \end{cases} \tag{8}$$

Now, the reward when continuing to throw R_c, will be defined. The reward when choosing to continue has to take into account that there is more information to obtain for each future throw. When we choose to continue throwing, the next die ρ_{k+1} is known, but how many eyes this die will show, ω_{k+1}, is unknown. We know that the probability of not throwing γ is the same for all the five unloaded sides, so we just have to consider two outcomes of what the next throw can be. Even though we do not know the true value for γ, this will later be adjusted for when using the law of total probability to condition on each of the sides being γ. This means that when considering the next throw, this can either be γ or not γ. Meaning that we have two possible outcomes for throw number $k + 1$, $\omega_{k+1} = \gamma$, or $\omega_{k+1} \neq \gamma$. For each possibility we have two more possibilities $\omega_{k+2} = \gamma$ or $\omega_{k+2} \neq \gamma$. Hence, the expression for the reward of continuing to throw is a recursion, where you for each possible outcome of future throws can choose to continue throwing or to stop throwing.

In the probability model for finding rewards we assume that future decisions are the best decisions, that is the decisions which maximizes the reward. So, for each possible outcome of the next throw $\omega_{k+1} = \gamma$ or $\omega_{k+1} \neq \gamma$ we have two rewards. When $\omega_{k+1} = \gamma$, we have the reward for continuing $R_c(\rho_1, ..., \rho_{k+2}, \omega_1, ..., \omega_{k+1} = \gamma)$, and the reward for stopping $R_s(\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_{k+1} = \gamma)$. Consequently, we have two rewards when $\omega_{k+1} \neq \gamma$, and our assumption is that future throws, meaning throws after the $k+1$'th throw, are the choices that give the highest reward. This is denoted by the variables ϕ^1 and ϕ^0, which for any future combination of throws gives the maximum of the two rewards R_s and R_c in that situation. Where ϕ^1 is the maximum reward when the next throw is γ, and ϕ^0 is the maximum reward when the next throw is not γ, their equations are

$$\begin{aligned}\phi^1 = \max\{&R_s(\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_{k+1} = \gamma),\\ &R_c(\rho_1, ..., \rho_{k+2}, \omega_1, ..., \omega_{k+1} = \gamma)\}\\ \phi^0 = \max\{&R_s(\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_{k+1} \neq \gamma),\\ &R_c(\rho_1, ..., \rho_{k+2}, \omega_1, ..., \omega_{k+1} \neq \gamma)\}.\end{aligned} \tag{9}$$

The reward of continuing, R_c, includes the probability of the next throw being γ and the probability of the next throw not being γ. Both probabilities are functions of the history, $P(\omega_{k+1} = \gamma | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$, and $P(\omega_{k+1} \neq \gamma | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$. In each situation for the outcome of the next throw, we can choose to continue or to stop throwing. Hence, the two probabilities of the next throw being γ and not being γ, are multiplied with ϕ^1 and ϕ^0 respectively Additionally we include the cost parameter α which will have a negative sign, reducing the reward if you choose to continue. This parameter value is individual for each participant and is later estimated using MLE. This results in the following expression for the reward when continuing to throw,

$$\begin{aligned}&R_c(\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) = -\alpha\\ &+ P(\omega_{k+1} = \gamma | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)\phi^1\\ &+ P(\omega_{k+1} \neq \gamma | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)\phi^0.\end{aligned} \tag{10}$$

Since we only consider two options for the next throw $\omega_{k+1} = \gamma$ and $\omega_{k+1} \neq \gamma$, the two probabilities of the outcome of the next throw being either γ or not, must sum to 1. That means that there is only one probability expression we need to derive in (10).

Considering the two equations (9) and (10), we see that they are depending on each other. The reward of continuing to throw the $k+1$'th throw from (10) includes ϕ^1 and ϕ^0, which both are expressions of the reward of continuing for throw number $k+2$. Hence, we have a recursion, where ϕ^1 and ϕ^0 are dependent on (10) where one more throw has been made.

With the aim of fitting a stochastic model to our data, we must express the covariate $x_i(\alpha)$. In order to do that we need the conditional expectations of the rewards $E[R_s|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k]$, and $E[R_c|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k]$. The equations for the rewards, found by (8) and (10) involves two probability expressions, which we need to find explicit exspressions for in order to find the conditional expectations of the rewards. Next, the model for finding these probabilities will be derived. Firstly, we will consider the probability of having identified the loaded die $P(Z = z|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k)$, which is the reward when stopping when throwing the fourth die $\rho_k = 4$ included in (8). The second probability expression that will be derived is the probability of the next throw being γ, denoted by $P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$ which is included in (10).

The probability of having identified the loaded die

Now we have two general equations for the rewards given by (8), and (10). Next, we will consider the components of the two equations. Starting with the reward of stopping to throw R_s when $\rho_k = 4$, which includes the four probabilities that each of the four dice is the true loaded die given what you have observed,

$$\begin{aligned} &P(Z = 1|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\ &P(Z = 2|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\ &P(Z = 3|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\ &P(Z = 4|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k). \end{aligned} \tag{11}$$

It is easy to calculate the four equations in (11) if p and γ are known, because then we can categorize the throws as described in Section 4.1, and count how many throws which belong to each category, $n_u(\boldsymbol{\rho})$, $n_{ll}(\boldsymbol{\rho}, \boldsymbol{\omega})$ and $n_{lu}(\boldsymbol{\rho}, \boldsymbol{\omega})$. After that, we are left with a combinatorics problem which can be solved by multiplying the probabilities of making the throws that have been observed using $n_u(\boldsymbol{\rho})$, $n_{ll}(\boldsymbol{\rho}, \boldsymbol{\omega})$ and $n_{lu}(\boldsymbol{\rho}, \boldsymbol{\omega})$.

Since the two variables p and γ are unknown, we use the law of total probability to condition on both p and γ, starting with p. As p is uniformly distributed on the interval $\langle \frac{1}{6}, 1 \rangle$ using the law of total probability gives an integral. Thus, for computational purposes p will be discretized on the interval, choosing a resolution, with s steps evenly distributed on the interval. Hence, the law of total probability conditioning on p can be approximated to a sum using the s discrete values of p,

$$\begin{aligned}
&P(Z = z|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\
&= \int_{1/6}^{1} P(Z = z|p, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) f(p|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) dp \\
&\approx \sum_{p} P(Z = z|p, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) f(p|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \Delta p.
\end{aligned} \tag{12}$$

Where we have that $\Delta p = \frac{1-1/6}{s}$. The two first factors of the sum in this expression are considered separately, beginning with the first factor. This factor is found using the law of total probability again, this time conditioning on γ. We know that $\gamma \in \{1, 2, 3, 4, 5, 6\}$ is discrete, hence the law of total probability includes a sum,

$$\begin{aligned}
&P(Z = z|p, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\
&= \sum_{\gamma=1}^{6} P(Z = z|p, \gamma, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) P(\gamma|p, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k).
\end{aligned} \tag{13}$$

Just like for (12), the two factors in (13) will be considered separately starting with the first factor.

The first factor of (13) is rewritten using Bayes' theorem, which gives

$$\begin{aligned}
&P(Z = z|p, \gamma, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\
&= \frac{P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z) P(Z = z|p, \gamma)}{P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma)} \\
&= \frac{P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z) P(Z = z|p, \gamma)}{\sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z) P(Z = z|p, \gamma)} \\
&= \frac{P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)}{\sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)},
\end{aligned} \tag{14}$$

where we have used the law of total probability to rewrite the denominator, such that we condition on Z. Note that the probability of any die being the true loaded die Z, is equal when we do no have any information about what throws have been made. Therefore $P(Z = z|p, \gamma) = \frac{1}{4}$, and $P(Z = z|p, \gamma)$ cancel in(14).

In (14) both the numerator and denominator express the probability of observing what we have observed when we know p, γ and Z. These probabilities can be calculated by multiplying the probabilities of making each throw that has been made. We know that a throw can belong to three categories, and we have notation for how many of the throws that belong to each category. We have n_{ll} which is the number of throws made on the loaded die showing the loaded side γ, we have n_{lu} which is the number of throws made on the loaded die showing an unloaded side, and we have n_u which is the number of throws made on any of the three unloaded dice. Thus, we can find the probability of the particular combination of throws

$$P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z) = p^{n_{ll}(\boldsymbol{\rho},\boldsymbol{\omega})} \left(\frac{1-p}{5}\right)^{n_{lu}(\boldsymbol{\rho},\boldsymbol{\omega})} \left(\frac{1}{6}\right)^{n_u(\boldsymbol{\rho})}. \tag{15}$$

Next, the second factor of (13) is considered. When solving the second factor of (13) Bayes' theorem is used giving

$$P(\gamma|p, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) = \frac{P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma)P(\gamma|p)}{P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p)}. \tag{16}$$

The numerator and denominator of (16) will be considered separately before combining them to give a solution to (16). We know that $P(\gamma|p) = \frac{1}{6}$, because we do not have any information of what throws have been made. We need to consider the first factor of the numerator in (16). The law of total probability is used to condition on Z

$$\begin{aligned} &P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma) \\ &= \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)P(Z = z|p, \gamma) \\ &= P(Z = z|p, \gamma) \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z), \end{aligned} \tag{17}$$

here $P(Z = z|p, \gamma)$ is a constant and can be moved out of the summation. The expression in the summation is equal to (15), so this probability we know how to find. Further, we have the denominator of (16), and to solve this we must use the law of total probability twice to condition on both Z and γ

$$\begin{aligned}
&P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p) \\
&= \sum_{\gamma=1}^{6} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma)P(\gamma|p) \\
&= P(\gamma|p) \sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)P(Z = z|p, \gamma) \\
&= P(\gamma|p)P(Z = z|p, \gamma) \sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z).
\end{aligned} \tag{18}$$

Here the same logic as above is applied, where the constants $P(\gamma|p)$ and $P(Z = z|p, \gamma)$ are set outside of the summation. Inserting (17) and (18) into (16), we see that the factors $P(\gamma|p)$ and $P(Z = z|p, \gamma)$ cancel, resulting in the following expression for the second factor of the sum in (13)

$$\begin{aligned}
&P(\gamma|p, \rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\
&= \frac{\sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)}{\sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)}.
\end{aligned} \tag{19}$$

Now we have the necessary equations to find the probability in (13) which is equal to the first factor of (12), the remaining problem is solving the second factor of (12). This is done similarly to how we found (19) using both Bayes' theorem and the law of total probability. Again we start by using Bayes' theorem, and then the numerator and denominator are found separately. Bayes' theorem applied to the second factor of (12) gives

$$f(p|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) = \frac{f(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p)f(p)}{f(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k)} \tag{20}$$

Because we are considering discrete values of p, the first factor of the numerator of (20) is $f(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p) = P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p)$, for a given p. This is equal to (18), so this we know how to find. Since p is uniformly distributed we know that $f(p) = \frac{6}{5}$, which is just a constant. So we need to consider the denominator of (20), which is solved using the law of total probability to condition on p, γ and Z,

$$\begin{aligned}
&f(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\
&\approx \sum_p P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p) f(p)\Delta p \\
&= f(p)\Delta p \sum_p \sum_{\gamma=1}^{6} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma) P(\gamma|p) \\
&= f(p)\Delta p P(\gamma|p) \\
&\cdot \sum_p \sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z) P(Z = z|p, \gamma) \\
&= f(p)\Delta p P(\gamma|p) P(Z = z|p, \gamma) \\
&\cdot \sum_p \sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z).
\end{aligned} \tag{21}$$

Inserting (18) and (21) into (20) we get

$$\begin{aligned}
&f(p|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) \\
&= \frac{\sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)}{\Delta p \sum_p \sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\rho_1, ..., \rho_k, \omega_1, ..., \omega_k|p, \gamma, Z = z)},
\end{aligned} \tag{22}$$

where the probability in both the numerator and denominator is found using (15). The second factor of (12) is solved using (22), thus we have all necessary equations for solving (12) which is equal to the four probabilities of each die being the true loaded die in (11). It is now possible to find the reward of stopping to throw the fourth die R_s given by (8).

The probability of throwing γ

The second probability expression we need to derive is the probability of throwing γ, $P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$, which is included in (10). As mentioned earlier, (10) also includes the probability of not throwing γ, and the two probabilities must sum to one. So by finding the probability of throwing γ we have all we need since

$$P(\omega_{k+1} \neq \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) = 1 - P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k).$$

When deriving a model for solving $P(\omega_{k+1} = \gamma | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$ we use the law of total probability to condition on p, γ and Z. When p, γ and Z are known, the probability of throwing γ can be found. Then we know if the next throw is on the loaded die Z or not, and we know the probability of the die showing γ. Meaning that the probability of throwing γ is $\frac{1}{6}$ if we throw a fair die, and the probability is p if we are throwing the loaded die. Conditioning on p, γ and Z gives

$$
\begin{aligned}
&P(\omega_{k+1} = \gamma | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \\
&\approx \sum_p P(\omega_{k+1} = \gamma | p, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \\
&\cdot f(p | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \Delta p \\
&= \sum_p \sum_{\gamma=1}^{6} P(\omega_{k+1} = \gamma | p, \gamma, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \\
&\cdot P(\gamma | p, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) f(p | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \Delta p \\
&= \sum_p \sum_{\gamma=1}^{6} \sum_{z=1}^{4} P(\omega_{k+1} = \gamma | p, \gamma, Z = z, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \\
&\cdot P(Z = z | p, \gamma, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \\
&\cdot P(\gamma | p, \rho_1, ..., \rho_{k+1}, , \omega_1, ..., \omega_k) f(p | \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \Delta p.
\end{aligned}
\tag{23}
$$

Observing that the second, third and fourth factors in the equation above are equal to (14), (19) and (22) respectively, so the only equation left to consider in order to solve (23) is the first factor. The first factor $P(\omega_{k+1} = \gamma | p, \gamma, Z = z, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$, is the probability that the next throw, ω_{k+1}, is on the loaded side γ. Additionally, the probability is now conditioning on that we know the true loaded die Z, the loaded side γ, the probability p, and we know what die we are throwing ρ_{k+1}. This means that we have all the information needed to find $P(\omega_{k+1} = \gamma | p, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)$. We do not need to consider the previous throws, because each throw is independent of previous throws. When we know p, γ, and Z, there is no more information to get from the previous throws made. Hence, the first factor can be written as follows

$$
\begin{aligned}
&P(\omega_{k+1} = \gamma | p, \gamma, Z = z, \rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k) \\
&= P(\rho_{k+1}, \omega_{k+1} | p, \gamma, Z = z),
\end{aligned}
\tag{24}
$$

and with the notation given in this equation we can use (15) to evaluate the necessary probabilities.

The expected rewards

Now that we have expressions for the probabilities used to express the rewards R_s and R_c from (8) and (10) respectively, we can consider the expectations of the conditional rewards. These will later be used in the stochastic model. As before we start with considering $E[R_s|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k]$. We take the expectation on both sides of equation (8) and get

$$\begin{aligned}
&E[R_s|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k] \\
&= \begin{cases} E[\max_{z \in \{1,2,3,4\}} P(Z = z|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k)] & \text{for } \rho_k = 4 \\ E[R_c|\rho_1, ..., \rho_{k+1} = \rho_k + 1, \omega_1, ..., \omega_k] & \text{for } \rho_k = \{1, 2, 3\}. \end{cases} \qquad (25) \\
&= \begin{cases} \max_{z \in \{1,2,3,4\}} P(Z = z|\rho_1, ..., \rho_k, \omega_1, ..., \omega_k) & \text{for } \rho_k = 4 \\ E[R_c|\rho_1, ..., \rho_{k+1} = \rho_k + 1, \omega_1, ..., \omega_k] & \text{for } \rho_k = \{1, 2, 3\}. \end{cases}
\end{aligned}$$

Observing that in the above equation $\max_{z=\{1,...,4\}} P(Z = z|\rho_1, ..., \rho_k = 4, \omega_1, ..., \omega_k)$ is just a number which we know how to find. The expectation is then simply equal to this value. We will consider the expected reward when continuing to throw the fourth die. So the second part of the above equation for $\rho_{k+1} = \{1, 2, 3\}$ will be evaluated afterwards.

Since the expressions for the rewards are recursions, we must define a limit of maximum number of throws. If no limit exists we could in theory continue to throw infinitely many times. A recursion requires that you have an end point, so we must limit the maximum number of throws. We first consider the expectation of the rewards by continuing to throw the fourth die, and we must limit the number of throws that we are allowed to make in order to find the expectation. Let n_4 denote the maximum number of throws allowed to make on die 4, and we have already made $k = n_4 - 1$ throws on the die, so we are only allowed one more throw. Considering the expected reward by continuing to throw, we take the expectation on both sides of (10) and move the constants outside of the expectation,

$$
\begin{aligned}
&E[R_c|\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_k] \\
&= E[-\alpha + P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_k)\phi^1 \\
&+ P(\omega_{k+1} \neq \gamma|\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_k)\phi^0] \\
&= -\alpha + P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_k)E[\phi^1] \\
&+ P(\omega_{k+1} \neq \gamma|\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_k)E[\phi^0].
\end{aligned} \tag{26}
$$

We are left with the two expectations $E[\phi^1]$ and $E[\phi^0]$. The two equations ϕ^1 and ϕ^0 include the reward when continuing to throw the current die again. However, in our example this would be the reward for throw number $k + 2 = n_4 + 1$, which is above our set limit. Thus, we do not consider $E[R_c|\rho_1, ..., \rho_{n_4+1}, \omega_1, ..., \omega_{n_4}]$ in this situation when finding the expectation of ϕ^1 and ϕ^0. Taking the expectation on both sides of (9), we have in this situation

$$
\begin{aligned}
E[\phi^1] &= E[R_s|\rho_1, ..., \rho_{n_4} = 4, \omega_1, ..., \omega_{n_4} = \gamma] \\
E[\phi^0] &= E[R_s|\rho_1, ..., \rho_{n_4} = 4, \omega_1, ..., \omega_{n_4} \neq \gamma]
\end{aligned} \tag{27}
$$

which are found using (25). This is the initial situation for the recursion where we have reached our maximum limit for the number of throws we allow to make on the fourth die. Next, the situation where we have made less than $n_4 - 1$ throws on the fourth die, and situations where we are considering the three first dies will be considered.

When we are throwing the fourth die, and we have made fewer than $n_4 - 1$ throws, $k < n_4 - 1$, the expected reward by stopping is unchanged. However, the expected reward of continuing has to include all possible combinations of future throws up to throw n_4. Since we have not reached our maximum limit of throws on the fourth die, the expectations of ϕ^1 and ϕ^0 include both rewards. Only the method for finding $E[\phi^1]$ will be discussed, as the only difference between $E[\phi^1]$ and $E[\phi^0]$ is if the next throw is γ or not. So by finding a way of solving $E[\phi^1]$ we can easily adapt this method to find $E[\phi^0]$. Again, taking the expectation on both sides of (9), we get

$$
\begin{aligned}
E[\phi^1] = E[\max\{&R_s(\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_{k+1} = \gamma), \\
&R_c(\rho_1, ..., \rho_{k+2} = 4, \omega_1, ..., \omega_{k+1} = \gamma)\}] \\
&= \max\{E[R_s|\rho_1, ..., \rho_{k+1} = 4, \omega_1, ..., \omega_{k+1} = \gamma], \\
&E[R_c|\rho_1, ..., \rho_{k+2} = 4, \omega_1, ..., \omega_{k+1} = \gamma]\}.
\end{aligned} \tag{28}
$$

If $k+2 = n_4$ we are at the end of the recursion and $E[R_c|\rho_1, ..., \rho_{k+2} = 4, \omega_1, ..., \omega_{k+1} = \gamma]$ is found using (26) and (27). However, if $k+2 < n_4$ we have to repeat the procedure by considering (26) for $E[R_c|\rho_1, ..., \rho_{k+2} = 4, \omega_1, ..., \omega_{k+1} = \gamma]$, until the next throw is n_4.

The expectation of the conditional rewards have been considered for the fourth die, next we look at the general expression for the expected rewards where we do not specify which die is considered. We have that the general expression for the expected reward of continuing is given by

$$\begin{aligned}
&E[R_c|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k] \\
&= E[-\alpha + P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)\phi^1 \\
&+ P(\omega_{k+1} \neq \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)\phi^0] \\
&= -\alpha + P(\omega_{k+1} = \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)E[\phi^1] \\
&+ P(\omega_{k+1} \neq \gamma|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_k)E[\phi^0],
\end{aligned} \tag{29}$$

For $\rho_{k+1} = 4$ this is equal to (26). However, when $\rho_{k+1} = \{1, 2, 3\}$ we must define a limit for how many throws we are allowed to make on each die n_1, n_2, and n_3 for die 1, 2, and 3 respectively. If we have made fewer throws on a die than the maximum number of throws that are allowed, we have that the expectations of ϕ^1 and ϕ^0 include both rewards of stopping to throw the current die, and to continue to throw the current die. Again we only consider $E[\phi^1]$ which for any die has the general equation

$$\begin{aligned}
E[\phi^1] = E[\max\{&R_s(\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_{k+1} = \gamma), \\
&R_c(\rho_1, ..., \rho_{k+2}, \omega_1, ..., \omega_{k+1} = \gamma)\}] \\
= \max\{&E[R_s|\rho_1, ..., \rho_{k+1}, \omega_1, ..., \omega_{k+1} = \gamma], \\
&E[R_c|\rho_1, ..., \rho_{k+2}, \omega_1, ..., \omega_{k+1} = \gamma]\}.
\end{aligned} \tag{30}$$

When considering the last throw allowed on a die so $k+1 = \{n_1, n_2, n_3\}$, the equations $E[\phi^1]$ and $E[\phi^0]$ only consider the reward of stopping to throw. As long as we have made fewer throws on a die than what is allowed we use (29) until we reach the limit of maximum throws allowed.

We have to set four limits n_1, n_2, n_3, and n_4 when finding the expectations of the conditional rewards for die 1 in our setup. Because we only consider future throws, when considering the second die, we need three limits n_2, n_3 and n_4, and the same argument holds for the third and fourth die. However, the limits n_1, n_2, n_3 and n_4 does not need to be the same when considering the expected rewards for each die. For computational purposes the limit on the die considered will be set higher than the limit on the dice that have not yet been thrown. Meaning that when considering the expected rewards when throwing die 1, we set $n_1 > n_2, n_3, n_4$. When moving on to consider the expected rewards for die number 2, we set $n_2 > n_3, n_4$, and so on.

4.3 Stochastic model

The probability model describing how to find the rewards when stopping and continuing to throw a die in the dice task, was defined in Section 4.2. In our stochastic model this will be used to express the covariate x_i for each decision, y_i, for a participant. Firstly, we define the particular stochastic model which will be used. Afterwards, we will estimate $\hat{\alpha}$ for each participant using MLE, but in order to do so we will fit the model, for different values of α. The setup for the stochastic model for different values of α will be explained through an example where we will focus on what data is used to fit each model. Lastly, the implementation for finding MLE of α is given, and we describe what programs were used to implement the entire model for finding probabilities and fitting the stochastic model.

In our stochastic model we only have one covariate $x_i(\alpha)$ defined by (5). Further, this means that the logit model from (3) has two additional parameters, the intercept, β_0, and β_1. After having tested and fitted some of the data, we discovered that the intercept, β_0, was never significantly different from zero. In addition, setting the intercept to zero caused a lower variance for β_1, hence we prefer to use the logit model with the intercept $\beta_0 = 0$. The stochastic model which will be used to fit the data for a given value of α is then given by

$$P(y_i = 1|\alpha) = \frac{e^{\beta x_i(\alpha)}}{1 + e^{\beta x_i(\alpha)}}. \tag{31}$$

Additionally, we assume that all choices are independent. Consequently, we only consider the probability that the next decision is to throw the current die again, without including past or future choices.

Parameter estimation

When considering parameter estimates for α and β, both are estimated using maximum likelihood estimation, MLE. With the aim of estimating the two parameters we need to fit the stochastic model using different values of α.

Table 1: An example of decisions made in the dice task. Ten decisions are made, and for each decision we have n covariates with different values of α.

i	y_i	$x_i(\alpha_1)$	$x_i(\alpha_2)$	...	$x_i(\alpha_n)$
1	1	$x_1(\alpha_1)$	$x_1(\alpha_2)$	...	$x_1(\alpha_n)$
2	0	$x_2(\alpha_1)$	$x_2(\alpha_2)$	...	$x_2(\alpha_n)$
3	1	$x_3(\alpha_1)$	$x_3(\alpha_2)$	...	$x_3(\alpha_n)$
4	0	$x_4(\alpha_1)$	$x_4(\alpha_2)$	...	$x_4(\alpha_n)$
5	1	$x_5(\alpha_1)$	$x_5(\alpha_2)$	...	$x_5(\alpha_n)$
6	1	$x_6(\alpha_1)$	$x_6(\alpha_2)$	...	$x_6(\alpha_n)$
7	0	$x_7(\alpha_1)$	$x_7(\alpha_2)$	...	$x_7(\alpha_n)$
8	1	$x_8(\alpha_1)$	$x_8(\alpha_2)$	...	$x_8(\alpha_n)$
9	1	$x_9(\alpha_1)$	$x_9(\alpha_2)$	...	$x_9(\alpha_n)$
10	0	$x_{10}(\alpha_1)$	$x_{10}(\alpha_2)$	...	$x_{10}(\alpha_n)$

The setup for which variables we will use to fit the model will be explained through an example which shows how the stochastic models for different α-values are fitted to one example participant. Afterwards, we will proceed to estimate α using profile likelihood.

The setup of the stochastic model involves the participants decisions which are the response y_i, and the covariate $x_i(\alpha)$. Since α is unknown, we will fit n models for each participant with n different values of α. The setup of the stochastic model is visualized in Table 1, for an example of choices, y_i. In Table 1 ten decisions have been made, so the response vector y_i is of length 10. Since you must roll each die at least one time, the first throw is not considered as a decision. After the first throw has been made, we can find n covariates, $x_1(\alpha_1), ..., x_1(\alpha_n)$. The second choice made in the example was to throw the first die again $y_1 = 1$. After this throw has been made, we can calculate n new covariates, $x_2(\alpha_1), ..., x_2(\alpha_n)$.

This example of a combination of decisions corresponds to Figure 5. Where we see that die 1 and 2 are thrown two times, and die 3 and 4 are thrown three times. What number of eyes the dice in Figure 5 show is from the first round of the dice task. So if you perform the dice task, the first round will show the same as Figure 5 for the same number of throws on each die.

When fitting a GLM to this example we first calculate all the different values of the covariate $x_i(\alpha)$, such that all components of Table 1 are known.

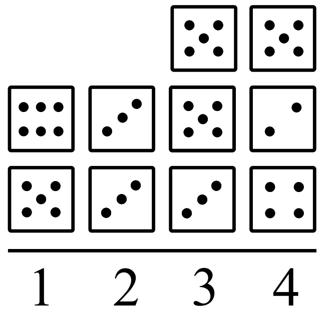

Figure 5: A possible combination of throws belonging to the decisions made in Table 1. Here we have first thrown a five on die one, and then a six, before moving on to die two which showed two threes and so on.

After that, we fit n models to the response y_i, such that a model consists of the response column with the values y_i and a column for the covariate $x_i(\alpha)$ for one α-value. When all n models are fitted, we can find the likelihood such that the parameters β and α can be estimated for each participant.

The MLE of β is returned from an R-function, and α is found using profile likelihood. Thus, we will focus on finding the MLE of α since this requires construction and implementation of the likelihood function.

In order to maximize the likelihood with respect to α we use the profile likelihood where we have

$$L(\alpha) = \max_{\beta} L(\alpha, \beta; y). \tag{32}$$

Information about the profile likelihood was found using Sprott (2000). Thus, we can use the β-estimate from R, and then consider which value of α that maximizes the likelihood. The log likelihood of our random variable y is found using (4). In the following, we first set up the expression for the log likelihood, and then consider the probability $P(y_i|\alpha, \beta)$ included in (4), before deriving the method for finding the log likelihood and thus the estimate for α.

In the expression for the log likelihood we have the probability $P(y_i|\alpha, \beta)$, and it is either $P(y_i = 1|\alpha, \beta)$ or $P(y_i = 0|\alpha, \beta)$. When summing over all y_i

decisions for this one participant, the log likelihood expression is

$$
\begin{aligned}
l(\alpha) &= \sum_{i=1}^{n} \ln P(y_i|\alpha,\beta) \\
&= \sum_{i=1}^{n} \ln \left(P(y_i = 1|\alpha,\beta)^{y_i} (P(y_i = 0|\alpha,\beta))^{1-y_i} \right) \\
&= \sum_{i=1}^{n} y_i \ln P(y_i = 1|\alpha,\beta) + (y_i - 1)\ln(1 - P(y_i = 1|\alpha,\beta)).
\end{aligned} \tag{33}
$$

Replacing the probability $P(y_i = 1|\alpha, \beta)$ in (33) with the expression for the probability from (31), we obtain an expression for the log likelihood,

$$
\begin{aligned}
l(\alpha) &= \sum_{i=1}^{n} y_i \ln \frac{e^{\beta x_i(\alpha)}}{1 + e^{\beta x_i(\alpha)}} + (1 - y_i) \ln \left(1 - \frac{e^{\beta x_i(\alpha)}}{1 + e^{\beta x_i(\alpha)}} \right) \\
&= \sum_{i=1}^{n} y_i \ln \frac{e^{\beta x_i(\alpha)}}{1 + e^{\beta x_i(\alpha)}} + (1 - y_i) \ln \frac{1}{1 + e^{\beta x_i(\alpha)}} \\
&= \sum_{i=1}^{n} y_i(\beta x_i(\alpha) - \ln(1 + e^{\beta x_i(\alpha)})) + (1 - y_i)(-\ln(1 + e^{\beta x_i(\alpha)})) \\
&= \sum_{i=1}^{n} y_i \beta x_i(\alpha) - \ln(1 + e^{\beta x_i(\alpha)}).
\end{aligned} \tag{34}
$$

Now that we have an expression for the log likelihood, we can find the α-value which maximizes our (34). The implementation of the method for finding this α-estimate will be presented.

Algorithm for estimating α

The stochastic model will be fitted to each participant for different α-values. Each of these models will have a MLE of $\hat{\beta}$ such that we have an expression for the profile likelihood, for each α as given in (34). We want to find the α-value which maximizes (34), and thus we have the α-estimate for one participant. Additionally, we assume that α only has one maximum.

The algorithm used to estimate $\hat{\alpha}$ considers n values of α evenly distributed on an interval, and for each α-value we calculate the log likelihood from (34). The MLE $\hat{\alpha}$, was set to the value of α which maximized (34).

The algorithm used to estimate $\hat{\alpha}$ is given in Algorithm 1, where we use a for-loop to fit a GLM which returns an estimate for $\hat{\beta}$. This β-estimate is used to calculate (34) for each α, and store this value in a vector. The MLE $\hat{\alpha}$, is found by considering the largest value in the vector of log likelihoods, and return the corresponding $\hat{\alpha}$. A detailed description on how the intervals were chosen is given in Section 5.2.

Algorithm 1 Implementation for estimating α for one participant.

```
loglikelihood = [ ]
for i in 1 to n do
    Use y_i and x_i(α_i) to fit a GLM, which returns β̂_i.
    Calculate l(α_i) using (34) and β̂_i, and add it to the loglikelihood vector.
end for
Which index k in the loglikelihood vector has the largets value l(α_k)
return α_k
```

Implementation of the probability and stochastic model

For implementation of the probabilistic model for finding expected rewards, the model was first implemented using R. As we have a recursion, the possible combinations of future throws increases rapidly so, the code for calculating the expected rewards should be efficient. The implementation in R was not very fast, so instead the model was implemented in Python using the package "Numba". This packages allows you to write Python code, and it translates the code into fast machine code which approaches the runtime of traditional compiled languages such as C or FORTRAN (Lam et al. 2015). The Python code created files with all $x_i(\alpha)$ for all participants.

The stochastic model was fitted using the programing language R. The files containing the response and covariates for all participants were read and used to fit the GLMs. When fitting a GLM, the R-function returns the MLE $\hat{\beta}$. The built in function "glm" in R was used. It calls the "glm.fit"-function which uses iteratively reweighed least squares (R Core Team 2022). Given a response vector y_i and a covariate vector $x_i(\alpha)$, the glm-function returns the estimate for $\hat{\beta}$.

The algorithm described in Algorithm 1 was implemented in R, and used to estimate $\hat{\alpha}$.

5 Results

With the aim of fitting a stochastic model to the data we need to decide the maximum number of throws allowed on each die n_1, n_2, n_3 and n_4, and the number of steps, s, to consider on the interval for the probability p. The choices of algorithmic parameters, and the reasons for choosing them will be given. Next, we will define the approach for deciding on α-values to consider such that we can estimate $\hat{\alpha}$ for all participants of the dice task. Finally, we will analyze the data from the two groups who have performed the dice task by comparing the estimated α-and β-estimates for each participant. We will consider if there are any differences between the estimates and the group the participants belong to.

5.1 Variable choices

When finding all the different values of $x_i(\alpha)$, needed to fit stochastic models, we need to define several algorithmic parameters. Which are the number of throws we are allowed to make on each die, and how many values, s, of the probability p we will consider.

When finding $x_i(\alpha)$ when considering die 1, we have a maximum number of throws allowed on the first, second, third and fourth die. When finding $x_i(\alpha)$ when considering die 2, we have a maximum number of throws allowed on the second, third and fourth die, and so on. However, it requires a lot of computational time if the maximum number of throws allowed are set high. Therefore, when considering $x_i(\alpha)$ for throws on the first die, we can allow more throws on die 1 than on the following dice. Afterwards, when considering the second die, we set the limit n_2 higher than n_3 and n_4. The actual values that were used in the analysis of the data are listed in Table 2. How these were determined will be explained by comparing the log likelihood for different values of α, found using (34). For testing we will use data from the dice task from one healthy participant. We choose an interval for α where we know that the log likelihood has its maximum for this participant. Within this framework we can change the algorithmic parameters n_1, n_2, n_3, n_4 and s, to see how they effect the log likelihood.

For a set of defined values like the ones in Table 2, we can fit models for different values of α as described in Section 4.3. For each model we fit, we estimate $\hat{\alpha}$ using (34), and the log likelihood can be plotted as a function of the α-values. First we evaluate the n_1, n_2, n_3 and n_4. Second, we evaluate

Table 2: Variables used when calculating $x_i(\alpha)$ for each decision made by each participants of the dice task. When calculating $x_i(\alpha)$ for decisions made when throwing die 1, we used the variables on the top row. When calculating $x_i(\alpha)$ for decisions made when throwing die 2, we used the variables on the second top row.

Considered die	n_1	n_2	n_3	n_4	s
die 1	31	5	2	2	15
die 2	-	20	2	2	15
die 3	-	-	20	2	15
die 4	-	-	-	20	15

Table 3: Variables tested for calculating $x_i(\alpha)$ for each decision made by each participants of the dice task. When calculating $x_i(\alpha)$ for decisions made when throwing die 1, we used the variables on the top row. When calculating $x_i(\alpha)$ for decisions made when throwing die 2, we used the variables on the second top row.

Considered die	n_1	n_2	n_3	n_4	s
die 1	10	1	1	1	15
die 2	-	10	1	1	15
die 3	-	-	10	1	15
die 4	-	-	-	10	15

the number of steps, s, we consider on the interval for p.

We fit two models both using ten steps of values for α evenly distributed on the interval $\alpha \in [0, 0.01]$. One model uses data from Table 2 and the other has lower limits for the maximum number of throws allowed, stated in Table 3. In Figure 6, we see the log likelihood as a function of α for models using values of n_1, n_2, n_3, n_4 and s given in Tables 2 and 3. Observing that the difference in the log likelihood calculated using variables from Table 3 and Table 2 is almost 0.2. Using data from Table 2 requires more computational time, and increasing the values further would be too time consuming. Thus, we do not consider higher values for n_1, n_2, n_3 and n_4, and we decide to use the values for n_1, n_2, n_3 and n_4 as given in Table 2.

We use the same values for n_1, n_2, n_3 and n_4, and double the number of

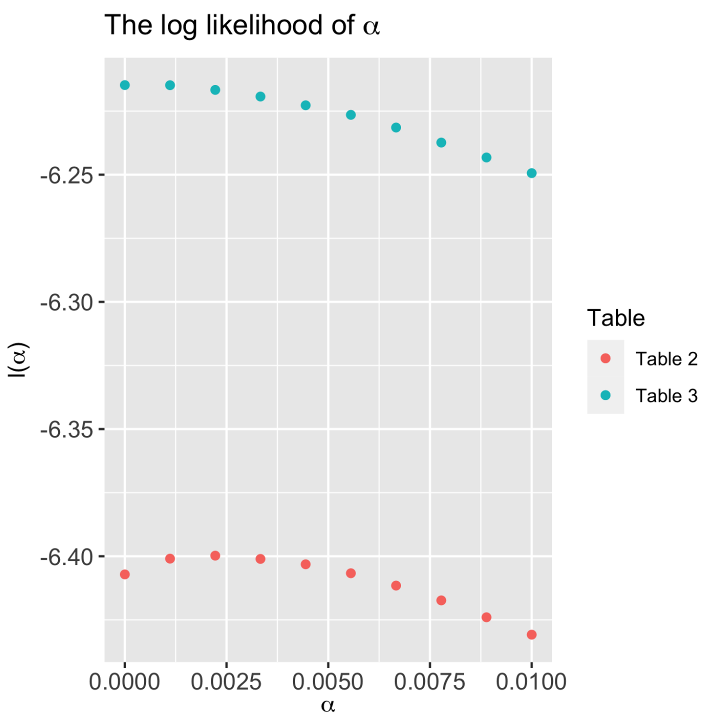

Figure 6: Plot of the log likelihood as a function of α for one participant, using ten values of α on the interval $\alpha \in [0, 0.01]$. The red plot is found using variables from Table 2, and the blue plot is found using variables from from Table 3. The difference is that Table 3 has a lower limit for n_1, n_2, n_3 and n_4 compared to Table 2.

Table 4: Variables tested for calculating $x_i(\alpha)$ for each decision made by each participants of the dice task. When calculating $x_i(\alpha)$ for decisions made when throwing die 1, we used the variables on the top row. When calculating $x_i(\alpha)$ for decisions made when throwing die 2, we used the variables on the second top row.

Considered die	n_1	n_2	n_3	n_4	s
die 1	31	5	2	2	30
die 2	-	20	2	2	30
die 3	-	-	20	2	30
die 4	-	-	-	20	30

values we consider on the interval for p, $s = 30$, to see what effect that has on the log likelihood. The variables that will be considered are given in Table 4. Using the variables in Tables 2 and 4 we can make the same type of plot as in Figure 6 to investigate the choice of s. The plot of the log likelihood is displayed in Figure 7. In Figure 7 we see that both plots have the same shape and the difference between the log likelihood using variables from Tables 2 and 4 is 0.01. By considering half as many values of s which is the difference between the variables in Table 2 and 4, we halved the computational time. Meaning that using Table 2 runs on half the time compared to using Table 4. Inspecting the plots in Figure 7, the difference between the two choices of s is not unreasonably large considering the time saved by choosing a lower resolution for p, hence s was set to 15 for analysis of the data.

By comparing the differences in the log likelihood in Figures 6 and 7, it seems that having higher values for n_1, n_2, n_3 and n_4 have a bigger effect on the log likelihood than a higher value of s. Hence the variables given in Table 2 were chosen for the analysis of all the participants.

The maximum number of throws for die 1 was set higher than for the three other dice. Because of the design of the dice task is such that no matter which color die you choose in what order, the first die will always show the same sequence, the second die you chose will show the same sequence and so on. Therefore all $x_i(\alpha)$ are the same for all decisions made while throwing the first die. To save time when calculating $x_i(\alpha)$, we first find all values of $x_i(\alpha)$ for 31 throws on the first die, and store them in a file. When finding all $x_i(\alpha)$ for each participant, we can read $x_i(\alpha)$ when considering die 1, from the file. This can not be done as easily for the second, third and fourth die,

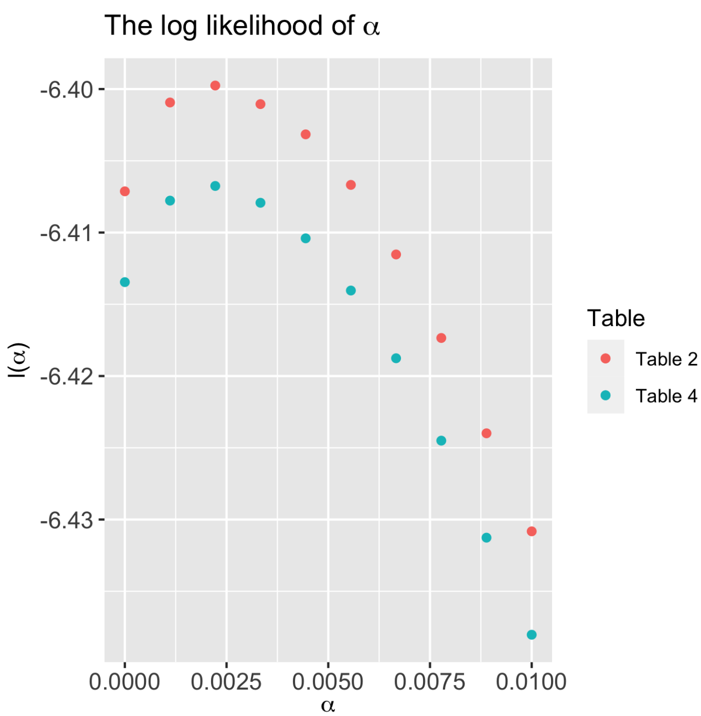

Figure 7: Plot of the log likelihood as a function of α for one participant, using ten values of α on the interval $\alpha \in [0, 0.01]$. The red plot is found using variables from Table 2, where $s = 15$. The blue plot is found using variables from from Table 4, where $s = 30$.

because the $x_i(\alpha)$ is dependent on the history of throws made. The history of throws made varies from participant to participant, because they do not make the same number of throws on each die. Hence, we can afford to set the algorithmic parameter values higher when calculating $x_i(\alpha)$ when considering die 1, since this only has to be done once for an interval of α-values. Due to time limitations the maximum number of throws were set lower when considering the three other dice. After having the file with $x_i(\alpha)$ for die 1, the runtime for calculating all $x_i(\alpha)$ for all choices made by one participant for one value of α using data from Table 2 was around 1 minute.

The algorithmic parameters n_1, n_2, n_3 and n_4 are defined as the maximum number of throws allowed. Meaning that if $n_1 = 20$, and you have made one throw on die 1, you have 19 possible throws left. If you make a second throw, you now have 18 possible more throws. As a consequence of the choices of maximum number of throws, the participants who have made more throws than our limit are excluded from the evaluation. Even though more throws are allowed on the first die, we make it even by excluding all participants who have made 20 throws or more on at least one die. The limit of maximum throws allowed must be one larger than the throws made by the participants. If we allow 20 throws when our limit is 20, we would calculate $x_{20}(\alpha)$ which would only allow us to stop, so this would not be correctly reflecting the situation the participant is in, because they do not have to stop at any limit. Excluding all participants who have made 20 throws or more on at least one die, means that we are left with 185 healthy controls and 36 patients. Where we have excluded 22 healthy participants and 2 patients.

5.2 Intervals considered for α

The cost parameter α is different for each participant, and we want to estimate α for each participant in order to see if there is a connection between the estimated values and the group the participant belongs to. Up until now, α has been treated as a parameter which our model is dependent on. The goal is to find the α-value which maximizes the log likelihood from (34). In order to fit the stochastic model, we need to insert values for α. Thus, we need to decide what values of α we will consider, such that the log likelihood has its maximum. The method for deciding which α-intervals to consider when fitting the models, is presented.

Firstly, some testing was performed for a few participants to identify for which α-values the log likelihood had its maximum. The same type of plots

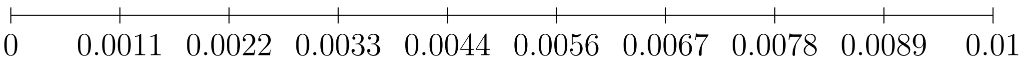

Figure 8: The interval for α first considered to sort participants by their α-value. The steps are rounded off, so the actual α-values used in the calculations were more precise.

as shown in Figures 6 and 7 were used for testing different intervals of α-values. The result from the testing was that many of the participants had their maximum on the interval $\alpha \in [0, 0.01]$.

Thus, the interval $\alpha \in [0, 0.01]$ is selected, and we consider ten evenly spaced points on this interval as shown in Figure 8. We find that 79 healthy, and 12 schizophrenic participants have their maximum value of the log likelihood in one of the α-values given in Figure 8, that were not on the edges, 0 or 0.01. We found that, 98 healthy, and 22 schizophrenic participants had their maximum of the log likelihood at zero, meaning that the true values lies between 0 and the first point 0.0011. Lastly, 8 healthy and 2 schizophrenic participants had their maximum at $\alpha = 0.01$, meaning that the true maximum most likely lies above 0.01.

In order to get more precise values for $\hat{\alpha}$, the participants were sorted according to what interval of α they had their maximum log likelihood on. The calculations of $x_i(\alpha)$ were run again for new intervals with ten values of α evenly distributed. For example if a participant had its maximum log likelihood for $\alpha = 0.0022$ the calculations of $x_i(\alpha)$ were run again for $\alpha \in [0.0011, 0.0033]$. This was done for all participants who had their maximum log likelihood for $\alpha > 0$ and $\alpha < 0.01$.

For the 98 and 22 participants who had their maximum log likelihood for $\alpha = 0$, the new interval for α was set to $\alpha \in [0, 0.0011]$ with ten point evenly distributed on the interval. The maximum log likelihood was then evaluated for each participant, and for the participants who had their maximum for $\alpha \in [0.00011, 0.0011]$, no further precision of $\hat{\alpha}$ was done. For participants who still had their maximum for $\alpha = 0$, we again considered a new interval $\alpha \in [0, 0.00011]$ with ten point. This scheme was done down to the interval $\alpha \in [0, 1.1 \cdot 10^{-6}]$, meaning that the participants having α-values $\alpha \in [0, 1.1 \cdot 10^{-7}]$ were not considered for smaller intervals of α. The participants who had α-values smaller than $1.1 \cdot 10^{-7}$ were defined to have $\alpha = 1.1 \cdot 10^{-7}$. This applies to 87 healthy and 20 schizophrenic participants.

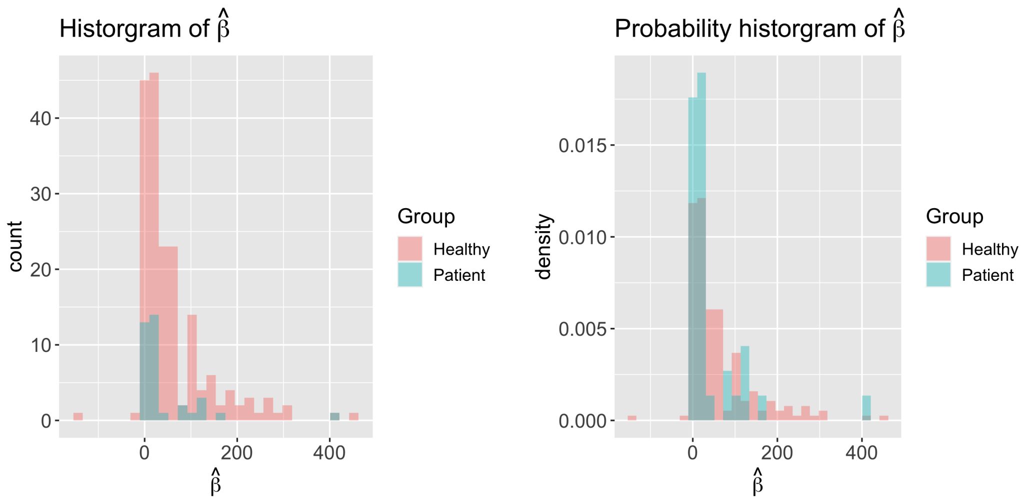

Figure 9: To the left we have the histogram of the estimates for $\hat{\beta}$, and to the right we have the probability histogram for $\hat{\beta}$ for all 185 healthy controls, and 36 patients considered. The two groups in both plots overlay.

Lastly, the 8 participants who had α-values above 0.01, were considered for ten values on the interval $\alpha \in [0.01, 0.1]$.

5.3 Analyzing $\hat{\alpha}$ and $\hat{\beta}$

After all the data has been considered, we have maximum likelihood estimates for both $\hat{\alpha}$ and $\hat{\beta}$ for all participants. These estimates will be presented by considering histograms for both parameters, probability histograms, and scatterplots of $\hat{\beta}$ and $\ln(\hat{\alpha})$. The focus while evaluating the plots of the estimates is to consider whether it is possible to see any differences between the two groups. The values of the estimates will we discussed, and compared to how many throws the participants made.

The histogram of the β-estimates are plotted in Figure 9, together with the probability histogram of the β-estimates. Here, from the histogram to the left, we see that the count for the healthy participants are higher, but this is due to our data which considered a total of 185 healthy controls, and 36 patients. By considering the probability histogram, we adjust for the difference in the sample size. From the probability histogram we see that

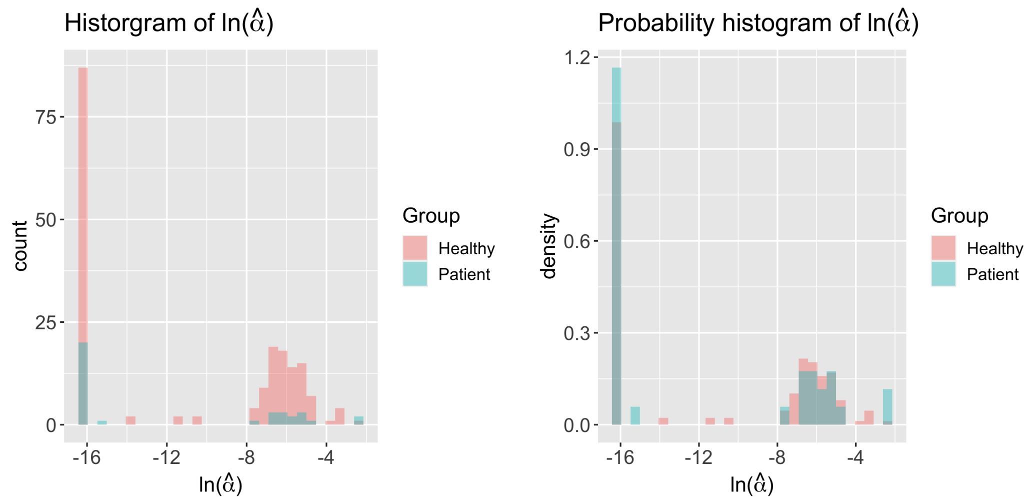

Figure 10: To the left we have the histogram of the estimates for $\ln(\hat{\alpha})$, and to the right we have the probability histogram of $\ln(\hat{\alpha})$ for all 185 healthy controls, and 36 patients considered. The two groups in both plots overlay.

the shapes of the histograms are similar for the two groups. The probability histogram shows a higher spike for β-estimates close to zero for the schizophrenic participants.This is however not entirely unexpected from the normalization. When normalizing the results of a scarcer data-set, it is likely to observe more amplified peaks and steeper curves. It is therefore difficult to determine whether this spike would be higher had we considered a larger sample of patients. Both the histograms have the same shape, and there are not any obvious difference to be observed between the two groups.

When considering the estimates $\hat{\alpha}$, we have that most of the values lie close to zero. Therefore, we consider the log-transformed estimates such that is is easier to distinguish the different α-estimates. In Figure 10 the log-transformed histogram of $\hat{\alpha}$ is plotted together with the probability histogram of the log-transformed α-estimates. There is a high count to the left end of the histogram. This is due to the 87 healthy participants, and 20 patients who have their MLE $\hat{\alpha}$ between 0 and $1.1 \cdot 10^{-7}$, and thus their α-estimate was set to $1.1 \cdot 10^{-7}$. The other estimates we see lie mostly around e^{-6}, where we see that both groups have higher counts. The estimates for the healthy controls spread out more than the estimates for the patients, but this is

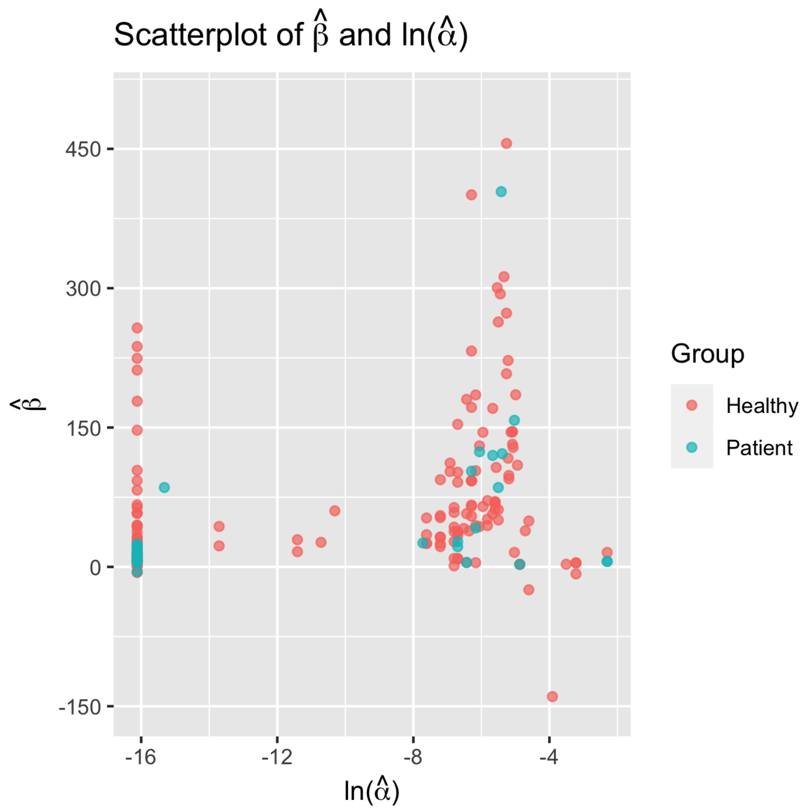

Figure 11: Scatterplot of the estimated β-values and the log-transformed estimated α-values, $\ln(\hat{\alpha})$, for all 185 healthy controls, and 36 patients considered. The plotted points overlay, so a stronger color indicates that several points have the same value.

likely due to the higher number of participants considered for the healthy controls. From considering the probability histogram, we see that the plots have the same shape for both groups. If we exclude the α-estimates that were set to $1.1 \cdot 10^{-7}$, we see that both groups are most likely to have their α-estimates around e^{-6}. From evaluating the histograms in Figure 10 there are no obvious differences that separate the two groups.

Both the estimated values $\hat{\alpha}$, and $\hat{\beta}$ are considered in the scatterplot in Figure 11. Here we also consider the log-transformed α-estimates such that it is easier to separate the estimated values. Again we see that many participants have α-estimates at $\hat{\alpha} = e^{-16}$, which is expected as we saw this

Table 5: Average number of throws made in the first round of the dice task for different categories of α and β-estimates. The average number of throws for healthy participants are given without parenthesis, and the average number of throws for schizophrenic participants are given in parenthesis.

	$\hat{\beta} < 150$	$\hat{\beta} > 150$
$\ln(\hat{\alpha}) < -16$	22,7 (20,6)	65,4 (-)
$-16 < \ln(\hat{\alpha}) < -8$	27,3 (44,0)	-
$-8 < \ln(\hat{\alpha}) < -4$	29,2 (26,9)	43,5 (39,0)
$-4 < \ln(\hat{\alpha})$	7,5 (7,0)	-

in the histogram, and probability histogram in Figure 10. Along the left edge we see that the healthy controls have β-estimates that spread out more and are higher than for the patients. However, we can not determine if this is happening because more healthy participants are considered, and therefore we observe a larger variation in the estimated values. The estimates that have $\hat{\alpha} > e^{-16}$, follow the same trend for both groups. Both groups have a larger variation in β-estimates for α-estimates, $e^{-8} < \hat{\alpha} < e^{-4}$. We see that the β and α-estimates for the healthy participants spread out more than for the patients. This might be because of the fewer participants considered for the patient group. Overall it is not possible to see any differences between the two groups of participants from the plot in Figure 11.

We can compare the estimated parameters to the average number of throws made by each participant. If we consider the scatterplot in Figure 10 a possible categorization could be to consider two categories of β-estimates and four categories of α-estimates. We consider β-estimates above and below 150, and four groups of α-estimates, $\ln(\hat{\alpha}) < -16$, $-16 < \ln(\hat{\alpha}) < -8$, $-8 < \ln(\hat{\alpha}) < -4$, and $-4 < \ln(\hat{\alpha})$. Combined we have eight categories, but the two categories with $\hat{\beta} > 150$ and $-16 < \ln(\hat{\alpha}) < -8$, and $-4 < \ln(\hat{\alpha})$ do not have any participants. The average of the total number of throws made by participants belonging to each of the six catagories, are presented in Table 5. We have that the average for the healthy participants are given first, and the schizophrenic participants are given in parenthesis.

From this we see that the β-estimate seems to have a great influence on the number of throws made. The participants having $\hat{\beta} > 150$ make around double the amount of throws as participants with β-estimates under 150.

This is reasonable if we consider our stochastic model from (31) and rewrite it as

$$P(y_i = 1) = \frac{1}{e^{-\beta x_i(\alpha)} + 1},$$

we see that a larger β-estimate increases the probability of throwing the current die again.

By considering the average number of throws made for the participants belonging to the four categories for $\hat{\beta} < 150$, we see that the number of throws made for participants having α-estimates $-4 < \ln(\hat{\alpha})$, are lower than for the three other categories. If we consider the plot in Figure 10, we see that these participants also have low β-estimates, which might be the reason that these participants make fewer throws. The number of throws made by the schizophrenic participants in the category, $-16 < \ln(\hat{\alpha}) < -8$, are higher than the number of throws made by both the healthy participants in the same group, and in the other categories with $\hat{\beta} < 150$. From the plot in Figure 10 we see that this is only one participant, and the β-estimate is closer to 150 than the other β-estimates for participants in the same category. So it is not unreasonable that this participants high number of throws can be explained by the β-estimate.

The average number of throws, for the three categories with $\hat{\beta} < 150$ and $\ln(\hat{\alpha}) < -4$, are quite similar, so it does not seem like the α-estimate is necessarily as connected to the number of throws as the β-estimate. We do see a slight increase in the number of throws made for higher values of $\hat{\alpha}$. By looking at the plot in Figure 10 this might be explained by the β-estimate. Where we see that the participants with $\ln(\hat{\alpha}) < -16$ and $\hat{\beta} < 150$, have β-estimates closer to zero compared to the two groups with $-8 < \ln(\hat{\alpha}) < -4$ and $\hat{\beta} < 150$. It is difficult to have a clear interpretation of the α-estimate as it is involved in the stochastic model in a less explicit way than $\hat{\beta}$. However, it does look like the combination of a high $\hat{\beta}$ and low $\hat{\alpha}$ estimate increases the number of throws made. The results suggest that if you have a high β-estimate and a low cost parameter $\hat{\alpha}$, it is likely that you will make many throws. However, if you have a larger cost parameter $\hat{\alpha}$, while maintaining the same high $\hat{\beta}$, this seem to only result in a somewhat reduced expected number of throws. Consequently, a possible interpretation of the results is that the β-estimate is the most influential on the probability of throwing the same die again.

6 Closing Remarks

A probability model for finding expected rewards for each decision made by participants in the dice task has been derived, and used to fit GLMs to each participant. The two parameters in the GLM were estimated using maximum likelihood estimation. The parameter estimates were plotted as histograms, probability histograms, and as a scatterplot of the two parameters. From the plots there was no clear difference between the two groups of participants. We saw in all the plots that the parameter estimates overlap, and follow the same trends. From this analysis of the data from the dice task, it is not possible to distinguish the two groups of participants.

Furthermore, we compared the parameter estimates to the average number of throws made by each participant. This showed that a higher β-estimate corresponded to an increased number of observed throws. The α-estimate does not seem to have as large of an influence on the number of throws made as the β-estimate, especially when the β-estimate is low. However, a low α-estimate combined with a high β-estimate seems to result in a high number of throws, wheres a high α-estimate seems to reduce the number of throws when maintaining a high β-estimate.

A recursion limit and a resolution for the discretization of p were set during calculations. These limitations were used when deriving the expectations of the reward when stopping and continuing to throw R_s, and R_c. The rewards were then employed to express the covariate $x_i(\alpha)$ in the stochastic model. It would be possible to redo the analysis with more precision if the maximum number of throws allowed on each die were set higher, and to consider a higher resolution on the interval of p. This might contribute to different parameter estimates. It would also be interesting to consider more even limits of the maximum number of throws between the current and succeeding dice. In this thesis the maximum number of throws was set higher on the die that was considered than the dice not yet thrown. We might have obtained a different result if we allowed the same, or nearly the same number of throws on both the die considered and future dice. Additionally, in this analysis we have defined the maximum limit such that it does not change when a new throw is added. Another approach could be to always allow n more throws independently of how many throws have been made. By doing that, we would remove the limitation that we can not consider participants who have made more throws than our limit. This would also mean that you could reduce the maximum number of throws allowed, and thus reducing the computational time. This allows a higher precision for the number of points to consider on p, and the number of throws allowed on future dice. Based on the results found in this thesis a higher number for the limits n_1, n_2, n_3, and n_4, seem to be more important than considering a higher resolution for the interval of p.

The minimum interval of α-values that were considerd, was set to $\alpha \in [0, 1.1 \cdot 10^{-6}]$. There were 87 healthy participants and 20 patients who had their α-estimate in this interval. It would be possible to apply the same model as used in this thesis to consider even smaller values of α.

In conclusion, from the analysis done and the results presented in this thesis, there is no observed difference between the two groups who have performed the dice task. Thus, we can not conclude that the dice task shows a difference in the decision making between healthy and schizophrenic participants.

References

Balzan, R. P., Ephraums, R., Delfabbro, P., and Andreou, C. (2017). "Beads task vs. box task: The specificity of the jumping to conclusions bias." *Journal of Behavior Therapy and Experimental Psychiatry*, 56, 42–50.

Fahrmeir, L., Kneib, T., Lang, S., and Marx, B. (2013). *Regression*. Springer.

Härdle, W. K., Klinke, S., and Rönz, B. (2015). *Introduction to Statistics*. Springer.

Lam, S. K., Pitrou, A., and Seibert, S. (2015). "Numba: A LLVM-based Python JIT compiler." In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, LLVM '15. New York, NY, USA: Association for Computing Machinery.

Moritz, S. and Woodward, T. S. (2004). "Plausibility judgment in schizophrenic patients: Evidence for a liberal acceptance bias." *The German Journal of Psychiatry*, 7, 66–74.

National Institute of Mental Health (2022). "Schizophrenia." `https://www.nimh.nih.gov/health/topics/schizophrenia` [Accessed: 2022/05/6].

Pfuhl, G. and Tjelmeland, H. (2019). "Probabilistic reasoning in schizophrenia is volatile but not biased." 441–444. Presented at Conference on Cognitive Computational Neuroscience in Berlin, September 13-16, 2019.

R Core Team (2022). *glm: Fitting generalized linear models*. R Foundation for Statistical Computing, Vienna, Austria.

Skogvang, J. (2021). "Modelling decisions in the box task." Master's thesis, Department of Mathematical Sciences, Norwegian University of Science and Technology.

Sprott, D. (2000). *Statistical Inference in Science*. Springer.

Tandberg, E. B. (2021). "An optimal strategy of identifying a loaded die." Project thesis in the course TMA4500, Department of Mathematical Sciences, Norwegian University of Science and Technology. Available upon request form the author.

NTNU

Norwegian University of
Science and Technology