Johanne Skogvang

Modelling Decisions in the Box Task

Masteroppgave i fysikk og matematikk
Veileder: Håkon Tjelmeland
Juni 2021

Masteroppgave

NTNU
Norges teknisk-naturvitenskapelige universitet
Fakultet for informasjonsteknologi og elektroteknikk
Institutt for matematiske fag

NTNU
Kunnskap for en bedre verden

Johanne Skogvang

Modelling Decisions in the Box Task

Masteroppgave i fysikk og matematikk
Veileder: Håkon Tjelmeland
Juni 2021

Norges teknisk-naturvitenskapelige universitet
Fakultet for informasjonsteknologi og elektroteknikk
Institutt for matematiske fag

NTNU
Kunnskap for en bedre verden

Modelling Decisions in the Box Task

Johanne Skogvang

Spring 2021

TMA 4900 - MASTER'S THESIS

Industrial Mathematics
Department of Mathematical Sciences
Norwegian University of Science and Technology

Preface

This is the TMA4900 - Industrial Mathematics, Master's Thesis, which is part of my Master of Science in Applied Physics and Mathematics with an Industrial Mathematics major. I wish to thank my supervisor, Håkon Tjelmeland, for following up on me helping me all the way through. He has answered all my questions with great patience and has been vital in helping me complete this work. I would also like to thank Kristoffer Klevjer and Gerit Pfuhl for introducing me to this interesting project and providing me with data. Thanks a lot to the ones who have been proofreading and to the ones helping me with git and tikz. Finally, I wish to express gratitude to all of my friends and fellow students for making my years in Trondheim unforgettable.

Abstract

Delusions are one of the main symptoms of schizophrenia, and delusion prone individuals have been linked to a 'jumping to conclusions' bias. That means drawing conclusions without having sufficient information. An information sampling task called the box task has been proposed to find if participants have this bias. In the box task, we have a grid of grey boxes that, when opened, either display the colour red or blue. Participants are informed that one colour is in the majority and that their task is to find out which one. We use two versions of the box task, one where the participants can open as many boxes as they want and another where the test terminates when they try to open a random box. These are called the unlimited and limited versions, respectively. In this report, we find an Ideal Observer solution of the box task, where an Ideal Observer is someone who would make the optimal choice each time a box is opened. We have data from 76 participants who have done both versions of the box task, and we define a model for how they make decisions using a softmax model. The model includes parameters, α, that is a minor loss or penalty a participant gets each time a box is opened, β, that is the loss we get if the test terminates in a limited trial and η, that is a measure of how good the decisions the participant make are. In the model, the probability that a box is red, Θ, has a prior distribution with hyperparameters γ and κ. We estimate the model parameters for each participant with maximum likelihood estimation and find confidence intervals using parametric bootstrapping. Finally, we look at the sensitivity to the hyperparameters in the prior distribution for Θ.

This model is a good fit for the participants who make good choices but not for those who make bad choices. Parametric bootstrapping makes the confidence intervals for the participants that make optimal, or close to optimal, choices have length zero, meaning that, for these participants, this is not the best choice of method for finding these intervals. Looking at the sensitivity in the unlimited case, we find that the values of $\hat{\eta}$ are not sensitive to the changes in the prior, whereas the values of $\hat{\alpha}$ tend to be smaller with one of the priors. However, in the limited version, the model is sensitive to the changes in the prior and tend to estimate smaller values for all three parameters for the smallest values of the hyperparameters we use here.

Sammendrag

Et av de viktigste symptomene på schizofreni er vrangforestillinger. Mange med dette symptomet har vist seg å trekke forhastede slutninger uten å ha nok informasjon, de har da har et 'jumping to conclusions' (JTC) bias. Bokstesten (the box task) har blitt foreslått å bruke for å finne ut om noen har et JTC bias. Da ser man tolv grå bokser i et rutenett. Man åpner en og en boks, og bak hver boks skjuler det seg en av to farger, rød eller blå. Deltakerne får beskjed om at en av fargene alltid er i majoritet og at de skal finne ut hvilken det er. I denne rapporten brukes to versjoner av bokstesten, én hvor man får åpne alle de tolv boksene og en annen hvor deltakerne får beskjed om at testen terminerer når en tilfeldig boks åpnes. Disse kalles henholdsvis ubegrenset og begrenset versjon. I denne rapporten finner vi en Ideell Observatør-løsning, hvor en Ideell Observatør er en deltaker som alltid tar optimale valg. Vi har data fra 76 personer som har gjort begge versjoner av bokstesten, og vi modellerer hvordan disse tar valg når de tar testen med en softmax-modell. Modellen inkluderer parameterne α, som representerer et lite tap man får hver gang en boks åpnes, β, som er det tapet man får hvis testen terminerer i den begrensete versjonen og η, som sier noe om hvor gode valg man tar. Sannsynligheten for at en boks er rød, Θ, har en apriorifordeling som inkluderer hyperparameterne γ og κ. Modellparameterne estimeres med sannsynlighetsmaksimering, og konfidensintervaller beregnes ved hjelp av parametrisk bootstrapping. Deretter ser vi på hvor sensitive resultatene er når vi forandrer på hyperparameterne i apriorifordelingen til Θ.

Modellen passer bra hvis deltakerne tar gunstige valg, men ikke fullt så bra hvis de tar dårlige valg. Lengden på konfidensintervallene til individene som tar optimale eller nesten optimale valg blir null. For disse deltakerne er derfor ikke parametrisk bootstrapping den beste måten å finne disse intervallene på. I den ubegrensede versjonen blir $\hat{\eta}$ påvirket lite når hyperparameterene forandres, mens $\hat{\alpha}$ tenderer til å bli mindre for de minste hyperparameterene brukt her. I den begrensede versjonen, derimot, får mange deltakere lavere estimater for alle tre parametere.

Contents

List of Figures

Chapter 1

Introduction

Schizophrenia is a psychotic disorder where at least two symptoms: delusions, hallucinations, disorganized speech, grossly disorganized or catatonic behaviour or negative symptoms such as reduced emotional expressions and lowered motivation, have to be present. Delusions are beliefs that will not change if contradicting evidence is presented. The most common type of delusions is persecutory delusions. People who have those kinds of delusions might think that they will be hurt, injured, tormented or so on by others. Referential delusions are also common. Then a person puts meaning into comments, gestures and actions, thinking that they are about themselves when they not necessarily are. Completely improbable beliefs are called bizarre delusions. These are delusions others find far-fetched, and they are things that cannot happen in real life. A bizarre delusion could, for example, be that a person believes that their organs have been removed and replaced by someone else's organs without there being any scars or other evidence of that happening. A delusion that is not bizarre could be that you think you are under police surveillance without there being any evidence supporting this. It might be hard to distinguish between delusions and strongly held ideas. The main distinction is about the degree of conviction and how much or little the beliefs can be amended when contradicting facts are presented (American Psychiatric Association, 2013).

Delusions are one of the main characteristics of schizophrenia as it appears in about three out of four of those diagnosed (Garety et al., 2011). Researchers have been trying to understand how the delusions are formed and maintained to improve treatment (Dudley et al., 2016). One important finding is that deluded individuals seem to make decisions based on less evidence than healthy and other psychiatric individuals. Making decisions based on little evidence is often referred to as a "jumping to conclusions" (JTC) bias. A person with this bias might reach decisions or form beliefs before reaching realistic conclusions and thus accept unrealistic ideas. They are therefore more prone to delusions. The hope is that if we can detect the JTC bias, we can reduce delusional thinking and prevent delusions.

The JTC bias is traditionally tested with a probabilistic reasoning task called the beads task. The participants are presented with two jars containing beads of two colours, for example, red and blue. The two jars have opposite ratios of each colour, meaning that if the first has 85% red beads and 15% blue, the second has 15% red and 85% blue. The participants are told that beads are drawn from one of the jars, and their task is to find out which one that is. We ask them to choose only when they are entirely sure, and they draw as many beads as they want. The beads are drawn sequentially, and after each draw, the participants are asked if they want to choose which jar we draw beads from or if they will continue to draw more beads. One is usually said to have a JTC bias if one decides after one or two beads (Moritz et al., 2017). However, the beads task has shown to pose some problems.

Some of the first to use the beads task were Huq et al. (1988). Already in the first article, they presented some of the problems with the beads task. They used an 85-15 ratio of the beads. When the two first beads that we draw are of the same colour, it is a 97% probability that the beads are from the jar with 85% of the beads in that colour. Therefore, one might argue that choosing a jar at that point is reasonable and does not show a JTC bias. Deluded individuals make decisions earlier than the control groups, but Huq et al. argue that non-deluded individuals are more conservative and that people with delusions cancel out that bias when gathering less information. In an article by Moritz et al. (2017), other problems with the beads task are discussed, for example, that many participants seem not to understand that we draw all the beads from the same jar. Thus, they might think that each time we draw a bead, they have to guess from which jar that single bead is coming. These participants are then classified to have a JTC bias. We can also see that it is common to make logical errors due to miscomprehension. In an article by Moritz and Woodward (2005), they found that 52% of the schizophrenic participants and 23% of the healthy controls had at least one response that was not logical. The participants that misunderstand are more likely to choose early. Moritz et al. (2017) further states that the beads task is correlated with intelligence. Lack of intelligence might be a reason for or a confound for misunderstanding the task. They also say that confidence influences decision-making. The participants are asked to choose when they are entirely sure which jar we draw beads from, which could make more confident participants decide earlier. We might conclude that participants make hasty decisions because they like to take risks or are not cautious. However, other tasks that account for confidence also display a JTC bias with the delusion-prone participants. Additionally, there is only a one-dimensional sequence of events in the beads task. Thus, it is harder to find different versions to test multiple times.

The box task has been suggested as an alternative to the beads task. Here, we present the participants with a grid of a fixed number of boxes. When we open a box, one out of two colours is displayed, for example, blue or red. The participants are told that one of the colours always is in the majority, and their task is to find out which one (Moritz et al., 2017). They can open as many boxes as they want before making a decision. We can change the number of boxes and the ratio of the two colours for each new trial.

In this report, we model how the participants make decisions in the box task. In the version of the box task used here, there are twelve boxes. The participants cannot choose which boxes they open, only if they open the next box or not. We use two different versions of the box task. The first is an unlimited one, where the participants can open as many boxes as they want, even until all twelve boxes are opened, before reaching a decision. In the second version, the participants are told that the test will terminate at a random point. If the test terminates before the participant has decided what the majority colour is, this counts as a failed trial. We call this the limited version. In Figure 1.1, we see a limited trial of the box task with red and blue boxes. The participant has opened two boxes and has to choose whether to open another box or decide whether blue or red is the dominant colour.

We have data from 76 participants that have done ten trials each of the box task. The first trial was a practice trial of the limited version, followed by three unlimited and six limited trials. We model the participants' decisions using a softmax model and fit this model to each participant with maximum likelihood estimation. Thus, we find the maximum likelihood estimates in the softmax model. We then find confidence intervals for each of these estimates using parametric bootstrapping and percentile intervals.

We also find an Ideal Observer solution of the box task. An Ideal Observer is a participant that always makes optimal, or ideal, choices and thus finds the best solution (He et al., 2013). Each time a box is opened, the participant has three options. The first is to choose that blue is the majority colour, the second that red is, and the third option is to open another box. We have defined loss functions for each of these alternatives, which represent the cost of choosing the different options. An Ideal Observer would always choose

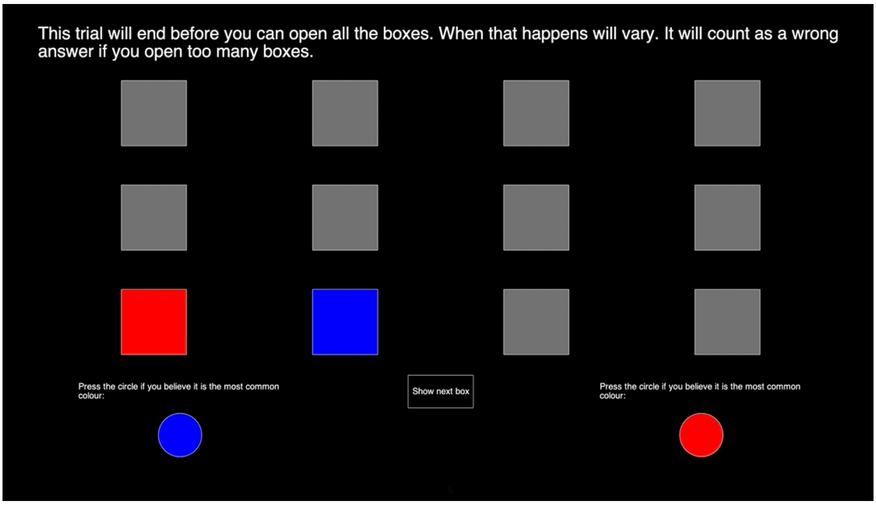

Figure 1.1: A limited trial of the box task with two opened boxes. The participants are to find out what the majority colour is given that one of them always is in the majority.

the decision with the least expected loss and end up with the overall optimal, or ideal, solution. In this report, we assume a binomial distribution with parameters 12 and some probability Θ for the total number of red boxes. Each time a box is opened, the probability that the box is red is Θ and $1-\Theta$ that the box is blue. We also assume a beta prior for Θ with parameters γ and κ. If we have any prior beliefs about the distribution of colours, we can incorporate this knowledge here. If both γ and κ are one, this is a uniform prior, meaning that Θ has the same probability of taking any value between zero and one.

In this report, we first go through some of the background theory used later on. Then, we will formulate the model for the decisions the participants make. That includes finding an Ideal Observer solution of the box task and describing how to find parameter estimates and confidence intervals. Further on, we will present some results, both the Ideal Observer solution and the parameter estimated with their respective confidence intervals. Additionally, we look at the sensitivity to the hyperparameters γ and κ. Lastly, we have some closing remarks.

Chapter 2

Background Theory

In this chapter, we go through some of the statistical theory used in this report. This includes the theorem of total probability, Bayes' theorem, the beta and gamma functions, Bayesian modelling, loss functions, the law of total expectation, the softmax function, maximum likelihood estimation and bootstrapping.

2.1 The Theorem of Total Probability

The theorem of total probability is often used when we want to find some probability, and this probability is hard to find. Then, sometimes it might be easier to find that probability if we condition on something, and use the theorem of total probability.

Theorem 1 (Theorem of Total Probability, Continuous Variables) *If we have a continuous variable, Θ, and a discrete variable, U, and both $P(U = u|\Theta = \theta)$ and $f_\Theta(\theta)$ are known for all θ, then we can find $P(U = u)$ from (Schay, 2016)*

$$P(U = u) = \int_{-\infty}^{\infty} P(U = u|\Theta = \theta) f_\Theta(\Theta = \theta)\, \mathrm{d}\theta. \tag{2.1}$$

Consider, for example, two discrete random variables U and V that are conditionally independent given the continuous stochastic variable Θ. To find the probability that $U + V$ is equal to some integer j, we can use the theorem of total probability to condition on theta. Thus,

$$P(U+V=j) = \int_{-\infty}^{\infty} P(U+V=j|\Theta=\theta) f_{\Theta}(\Theta=\theta)\, d\theta.$$

Later, we can exploit the conditional independence. If θ is a probability defined on the interval (0,1), this will be integrated on that interval, such that

$$P(U+V=j) = \int_{0}^{1} P(U+V=j|\Theta=\theta) f_{\Theta}(\Theta=\theta)\, d\theta.$$

2.2 Bayes' Rule

We can use Bayes' rule to find conditional probabilities and distributions.

Theorem 2 (Bayes' Rule) *Consider two events, A and B. We can find the probability of A given event B by the use of the probability of event B given A and the probabilities of the events A and B separately (Casella and Berger, 2002). Hence,*

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \tag{2.2}$$

As an example, consider a discrete random variable, U. We can find the probability that U is greater than or equal to 7, and condition on it being different from six by using (2.2). Then $U \geq 7$ is an event, and $U \neq 6$ is another event. Thus,

$$P(U \geq 7|U \neq 6) = \frac{P(U \neq 6|U \geq 7)P(U \geq 7)}{P(U \neq 6)}. \tag{2.3}$$

2.3 The Beta and Gamma Functions

Later, we will use the beta and gamma functions and some of their properties. These are therefore stated here. This theory can for example be found in Casella and Berger (2002). The gamma function for a parameter κ is

$$\Gamma(\kappa) = \int_{0}^{\infty} t^{\kappa-1} e^{-t} \mathrm{d}t.$$

A useful property of the gamma function is that it is recursive. Hence,

$$\Gamma(\kappa+1) = \kappa\Gamma(\kappa), \quad \kappa > 0. \tag{2.4}$$

Additionally, the beta function with parameters γ and κ is defined as

$$\mathrm{B}(\gamma, \kappa) = \int_0^1 \theta^{\gamma-1}(1-\theta)^{\kappa-1}\,\mathrm{d}\theta. \tag{2.5}$$

We can express the beta function as a product of gamma functions. This yields

$$\mathrm{B}(\gamma, \kappa) = \frac{\Gamma(\gamma)\Gamma(\kappa)}{\Gamma(\gamma+\kappa)}. \tag{2.6}$$

2.4 Bayesian Modelling

Consider a stochastic variable, U, that has a probability density function $f(u|\theta)$, where θ is a parameter upon which U depends. In classical statistics, θ is said to be a fixed but unknown value. The goal is to find this one true value. However, in Bayesian statistics we consider θ as a stochastic variable, such that θ has a density function. Here, the goal is to find the underlying density. To do so we propose a prior distribution for θ, $f(\theta)$. The prior distribution represents the prior knowledge we have about θ before observing any data. That could be our own subjective believes about the parameter or information based on other previously collected data or studies. One could also choose a prior distribution that does not say anything about the parameter at all. This is called a non-informative prior, and it is often used when we have none or little prior information about the parameter (Givens and Hoeting, 2012). If we have collected data, denoted u, we can update our prior beliefs with the information we get from that data. The resulting distribution is called the posterior distribution of θ, $f(\theta|u)$. We can find this using Bayes' theorem, and it includes both the prior information we have and the new information we get from the data.

Consider a discrete stochastic variable, U, that has a sampling distribution $P(U = u|\theta)$, and let $P(U = u)$ be the marginal distribution of U. Additionally, let $f(\theta)$ be the prior distribution of θ. Using Bayes' rule as it is stated in (2.2), we get that the posterior distribution of θ given u, $f(\theta|u)$, can be expressed as (Casella and Berger, 2002)

$$f(\theta|u) = \frac{P(U = u|\theta)f(\theta)}{P(U = u)}.$$

We can sometimes exploit the fact that the posterior distribution is proportional to the numerator in the above expression. This is because the denominator is a normalising constant. Hence,

$$f(\theta|u) \propto P(U = u|\theta)f(\theta). \tag{2.7}$$

If (2.7) has the form of a known distribution, then that known distribution is the posterior distribution.

As an example, consider a random variable, U, that is binomially distributed with parameters 12 and some probability, θ. Thus,

$$(U|\Theta = \theta) \sim \text{Binomial}(12, \theta).$$

Hence, the probability that we have u successes out of twelve, given θ, is

$$f(u|\theta) = \binom{12}{u} \theta^u (1-\theta)^{12-u}. \tag{2.8}$$

As θ is a probability, its value is on the interval $[0, 1]$. We know that the beta distribution is conjugate with the binomial distribution and has value between 0 and 1 (Casella and Berger, 2002), thus we choose a beta prior for θ with parameters γ and κ. Hence,

$$\Theta \sim \text{Beta}(\gamma, \kappa). \tag{2.9}$$

The prior density of Θ is then

$$f(\theta) = \frac{1}{\mathrm{B}(\gamma, \kappa)} \theta^{\gamma-1} (1-\theta)^{\kappa-1}, \tag{2.10}$$

where $B(\gamma, \kappa)$ is the beta function as defined in (2.5). We can find the posterior distribution of θ using (2.7), (2.8) and (2.10). Thus,

$$\begin{aligned} f(\theta|u) &\propto f(u|\theta) f(\theta) \\ &\propto \binom{12}{u} \theta^u (1-\theta)^{12-u} \frac{1}{\mathrm{B}(\gamma, \kappa)} \theta^{\gamma-1} (1-\theta)^{\kappa-1} \end{aligned}$$

All the factors that do not include θ are constants, and we collect them together as one constant, denoted C. Then

$$f(\theta|u) \propto C\, \theta^{u+\gamma-1} (1-\theta)^{12-u+\kappa-1}.$$

We can see that this is proportional to a beta distribution like the one in (2.10), but in this case with parameters $u + \gamma$ and $12 - u + \kappa$. Hence, the posterior distribution is a beta distribution with these parameters,

$$\Theta | U = u \sim \text{Beta}(u + \gamma, 12 - u + \kappa).$$

2.5 Loss Functions

A loss function typically says something about the cost, or loss, of an action related to a parameter. Let Ω_δ be the action space, consisting of all the actions that we can do, where δ is an action. Then,

$$\delta \in \Omega_\delta.$$

Additionally, let z be the true, but unknown state of nature, where

$$z \in \Omega_z.$$

We can define a loss function that depends on z and δ, which we denote $L(z, \delta)$. This is then the loss when making a decision, δ, regarding z (Liese and Miescke, 2008).

A loss function could for example be the 0-1-loss function. If for example $\Omega_\delta = \Omega_z = \{0, 1\}$, the loss function could be

$$L(z, \delta) = I(z \neq \delta), \tag{2.11}$$

where I is an indicator function such that

$$L(z, \delta) = \begin{cases} 0, & \text{if } z = \delta, \\ 1, & \text{if } z \neq \delta. \end{cases}$$

In some cases, we would like to find the expected value of the loss function. Taking the expected value of an indicator function gives the probability that the event is happening (Cormen et al., 2009). Hence, taking the expectation of (2.11) gives

$$E[L(z, \delta)] = E[I(z \neq \delta)] = P(z \neq \delta). \tag{2.12}$$

As an example, consider the box task with twelve boxes that could be either blue or red once they are opened. We define a stochastic variable, X_i, that represents the colour of the i-th opened box, such that

$$X_i = \begin{cases} 0, & \text{if box } i \text{ is blue,} \\ 1, & \text{if box } i \text{ is red.} \end{cases} \tag{2.13}$$

When i boxes are opened, let $X_{1:i}$ denote the colours of the i boxes, such that

$$X_{1:i} = (X_1, X_2, ..., X_i). \tag{2.14}$$

Additionally, let Z be the colour that is in the majority when all twelve boxes are opened, the true majority colour. This is also a stochastic variable as it depends on the colours of the twelve boxes, the X_i's. We define Z as

$$Z = I\left(\sum_{j=1}^{12} X_j > 6\right). \tag{2.15}$$

Then,

$$Z = \begin{cases} 0, & \text{if blue is the true majority colour,} \\ 1, & \text{if red is the true majority colour.} \end{cases} \tag{2.16}$$

Also, let δ be the choice the participant makes about which colour that is the dominant colour, such that

$$\delta = \begin{cases} 0, & \text{if the participant chooses blue as the majority colour,} \\ 1, & \text{if the participant chooses red as the majority colour.} \end{cases}$$

We can then define a loss function for the choice that the participant makes. This can be a 0-1 loss as in (2.11), and the loss function can therefore be defined as

$$L(Z, \delta) = I(Z \neq \delta), \tag{2.17}$$

Then, the loss is zero if the participant chooses the right colour as the majority colour and one if she chooses the wrong colour.

To find the expected loss, we take the expectation of the loss function. As Z depends on the colours of the twelve boxes, we condition on the colour of the already opened boxes, $X_{1:i} = x_{1:i}$. The expectation of the loss function is then

$$E[L(Z, \delta)|X_{1:i} = x_{1:i}] = E[I(Z \neq \delta)|X_{1:i} = x_{1:i}].$$

As in (2.12), this expectation is the probability that $\delta \neq Z$, but here the probability depends on X. Thus,

$$E[L(Z, \delta)|X_{1:i} = x_{1:i}] = P(Z \neq \delta|X_{1:i} = x_{1:i}). \tag{2.18}$$

2.6 The Law of Total Expectation

Let $\{A_1, A_2, ..., A_k\}$ be a partition of the sample space, S. Thus, there are k non-overlapping parts, such that $A_i \cap A_j = \emptyset, \forall\, i \neq j$. Then we also have that $S = A_1 \cup A_2 \cup ... \cup A_k$. If we want to find the expectation of an event, B, and we have the expectation of B on each of these partitions, we can use the law of total expectation. It states that

$$E[B] = \sum_i E[B|A_i]P(A_i). \tag{2.19}$$

This can also be used to find the expectation of functions (Schay, 2016). Let $g(B)$ be the function that we want to take the expectation of, then

$$E[g(B)] = \sum_i E[g(B)|A_i]P(A_i). \tag{2.20}$$

Later we will use the law of total expectation when we find the expectation of a loss function that says something about the loss of opening the next box in the box task. This expected loss is dependent on the colour of the box that will be opened. Thus, to find that expected loss, we use the law of total expectation and condition on the colour of the following box.

2.7 The Softmax Function

The softmax function is commonly used in classification problems with more that two classes (Bishop, 2013). Consider a decision, Δ, which now is a stochastic variable for which we want to construct a distribution. We find a probability mass function for $\Delta = \delta$ using a softmax function. Let there be D decisions, such that

$$\delta \in \{0, 1, 2, ..., D-1\}.$$

Additionally, let $\mathcal{E}_\delta(\varphi)$ be values tied to each decision that depends on some parameters, φ. The probability mass function for each decision, δ, could be found using a softmax function, such that

$$f(\delta|\varphi, \eta) = \frac{\exp(-\eta\mathcal{E}_\delta(\varphi))}{\sum_{d=0}^{D-1}\exp(-\eta\mathcal{E}_d(\varphi))}, \tag{2.21}$$

where η is some parameter.

These decisions could, for example, be the three choices we have each time we open a box in the box task. These choices are that blue is the majority colour, or that red is, denoted $\delta = 0$ and $\delta = 1$, respectively. The last choice is to open another box, which we denote $\delta = 2$. Then, we let $\mathcal{E}_0(\varphi)$ be the expected loss when choosing that blue is the majority colour, similarly to (2.18). Additionally, we let $\mathcal{E}_1(\varphi)$ and $\mathcal{E}_2(\varphi)$ be the expected loss of choosing red as the majority colour and of opening another box, respectively. Then, the probability mass function for δ could be as in (2.21). We then have a probability mass function for each of the three decisions that depend on the expected losses, parameters φ and some parameter η.

2.8 Maximum Likelihood Estimation

Maximum likelihood estimation is used to find estimates for parameters in a distribution. These are the estimates that, as the name implies, maximises the likelihood, and for short, we call them MLEs. Assume that we have probability distribution for a stochastic variable, Δ, and consider n samples, $\delta_1, \delta_2, ..., \delta_n$, of Δ. Denote the probability mass function for each of these δ's as $f(\delta|\varphi)$, where φ contains the parameters in the probability mass function. If the δ_i's are independent, the likelihood function is defined as

$$L(\varphi|\delta_1, \delta_2, ..., \delta_n) = \prod_{i=1}^{n} f(\delta_i|\varphi). \tag{2.22}$$

The MLEs are then the estimates of φ that maximises this function, and they are usually denoted $\hat{\varphi}$. It is often hard to maximize the likelihood function, then it might be easier to take the logarithm of the likelihood function and maximize that instead. This is called the log likelihood function, and is normally denoted as l. Thus,

$$\begin{aligned} l(\varphi|\delta_1, \delta_2, ..., \delta_n) =& \log\left(L(\varphi|\delta_1, \delta_2, ..., \delta_n)\right) \\ =& \log\left(\prod_{i=1}^{n} f(\delta_i|\varphi)\right). \end{aligned} \tag{2.23}$$

As the logarithm of products is the sum of the logarithms, we get that the log likelihood is

$$l(\varphi|\delta_1, \delta_2, ..., \delta_n) = \sum_{i=1}^{n} \log(f(\delta_i|\varphi)). \tag{2.24}$$

Maximizing this will give the same maximum point as if we maximize the likelihood function (Casella and Berger, 2002).

As an example, consider that the δ_i's have probability mass function as in (2.21). The parameters that we want to find estimates for are then φ and η. If we have n samples of Δ, denoted δ_i, where $i \in \{1, 2, ..., n\}$, the likelihood function would be

$$\begin{aligned} L(\varphi, \eta|\delta_1, \delta_2, ..., \delta_n) &= \prod_{i=1}^{n} f(\delta_i|\varphi, \eta) \\ &= \prod_{i=1}^{n} \frac{\exp(-\eta\mathcal{E}_{\delta_i}(\varphi))}{\sum_{d=0}^{2}\exp(-\eta\mathcal{E}_d(\varphi))}. \end{aligned}$$

The log likelihood would then be

$$\begin{aligned} l(\varphi, \eta|\delta_1, \delta_2, ..., \delta_n) &= \sum_{i=0}^{N} \log\left(\frac{\exp(-\eta\mathcal{E}_{\delta_i}(\varphi))}{\sum_{d=0}^{2}\exp(-\eta\mathcal{E}_d(\varphi))}\right) \\ &= \sum_{i=0}^{N}\left(-\eta\mathcal{E}_{\delta_i} - \log\left(\sum_{d=0}^{2}\exp\left(-\eta\mathcal{E}_d(\varphi)\right)\right)\right). \end{aligned}$$

The maximum likelihood estimators of φ and η would then be the values that maximises this log likelihood function. We denote them as $\hat{\varphi}$ and $\hat{\eta}$.

2.9 Bootstrapping

Consider a sample, $(\delta_1, \delta_2, ..., \delta_n)$, where the δ_i's are identically and independently distributed from an unknown distribution, F. We can use this sample to estimate this distribution, denoted by $\hat{F}$. To get some ideas about the properties of F, we can find the properties of $\hat{F}$. Sometimes it is challenging to do this analytically. Instead, we can use simulations, and this is where bootstrapping is useful. Bootstrapping is a way of finding new samples, either from the original sample, $(\delta_1, \delta_2, ..., \delta_n)$, or from the estimated distribution, $\hat{F}$. We can then use those samples to find, for example, standard error, bias, variance, or perhaps the most common; confidence intervals (Efron and Tibshirani, 1993).

There are two types of bootstrapping, nonparametric and parametric. In the nonparametric bootstrap, $\hat{F}$ is the empirical distribution of the data, and we take samples from our original sample. Consider for example that you have a dataset, $\boldsymbol{\delta} = (\delta_1, \delta_2, \delta_3, \delta_4, \delta_5)$. A bootstrap sample of this might then be $(\delta_5, \delta_5, \delta_2, \delta_3, \delta_1)$ and another might be $(\delta_2, \delta_4, \delta_2, \delta_2, \delta_1)$. These are resampled versions of $\boldsymbol{\delta}$. Thus, the bootstrap samples consists of elements from the original dataset, but some of them might not appear at all in a bootstrap sample while others might appear more than once. Drawing B of these samples, we can do inference about the population the original data is from.

In the parametric bootstrap, we make assumptions about the population, and $\hat{F}$ is the parametric distribution. Consider a sample, $(\delta_1, \delta_2, ..., \delta_n)$, from a distribution that has a probability mass function $f(\delta|\varphi)$, where φ might be a vector of parameters (Casella and Berger, 2002). We can for example find an estimate, $\hat{\varphi}$, of φ, using maximum likelihood estimation as in Chapter 2.8. When we have done that, we can draw new samples, denoted δ_i^* from $f(\delta|\hat{\varphi})$, such that

$$\delta_1^*, \delta_2^*, ..., \delta_n^* \sim f(\delta|\hat{\varphi}).$$

If we draw B samples, we can, as for the nonparametric bootstrap, do inference.

2.9.1 Confidence Intervals with Bootstrap Samples

One way of doing inference is to find confidence intervals. When we have B bootstrap samples, there are multiple methods for finding these. A confidence interval (CI) for a parameter is an interval that will contain the true value of the parameter a given proportion of the times an interval is constructed. If we, for example, have a 90% CI, then the true value of the parameter will be in the interval 90% of the times we construct a new one (Efron and Tibshirani, 1993).

One method for finding CIs with bootstrap samples is the percentile method. The percentile method is simple to both understand and implement. However, these confidence intervals might be biased. Then, one could instead

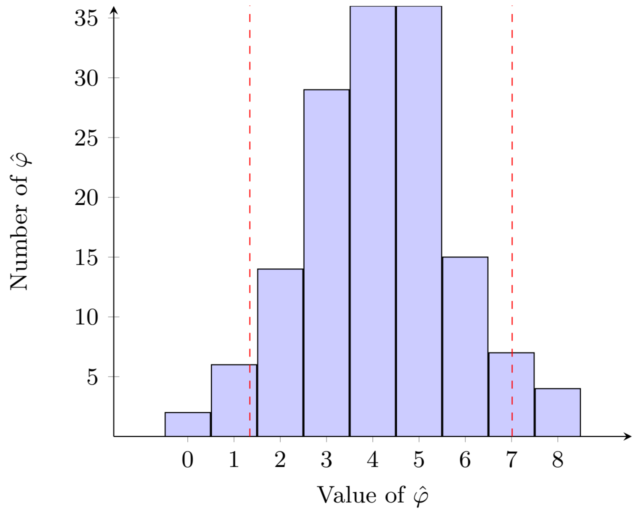

Figure 2.1: Here we have plotted the MLEs of 150 bootstrap samples in a histogram. The red dashed lines represent the 5-th and 95-th percentiles.

use approaches such as *bias corrected and accelerated* intervals or *approximate bootstrap confidence* intervals. In this report, we use the percentile method to find confidence intervals.

Consider a situation with B bootstrap samples. Let the vector φ be a parameter, for which we want to find a confidence interval. Then, we find the MLE of φ for each of the B samples. If we want to find a 90% confidence interval using the percentile method, we find the 5-th and 95-th percentiles. Plotting the MLEs of φ is a histogram, the 5-th percentile is the value of $\hat{\varphi}$ in the histogram where 5% of the samples are below. The 95-th is where 5% of the values are above. This is visualised in Figure 2.1. Here we have 150 bootstrap samples, and we have found the MLE of φ for each sample. These values are plotted in a histogram, where the red dashed lines represent the 5-th and 95-th percentiles. Then 5% of the MLEs lie to the left of the left red line, and 5% lie to the right of the right red line. The 90% CI for φ is around (1.4,7) when using the percentile method.

Chapter 3

Model formulation

The box task is an information sampling task used to assess a 'jumping to conclusions' (JTC) bias (Balzan et al., 2017). In the box task used in this report, the participants are shown a grid of twelve boxes, and each time a box is opened, one out of two colours, for example, blue or red, is displayed. Participants are told that one colour is always in the majority and that their task is to find out which one. We use two different versions of the box task. In the first one, the participants can open as many of the twelve boxes as they want before deciding which of the two colours is in the majority. We call this the unlimited version. In the second one, which we call the limited version, the participants are informed that the test will terminate at one point when a random box is opened. If the participant has not decided what the majority colour is when the test terminates, this counts as a failed trial. The participant could, for example, try to open the fourth box when the test terminates. Then, she does not get to see what colour that fourth box has. She cannot choose what she thinks is the majority colour, and this is a failed trial.

We have data from 76 participants that have done multiple trials of both versions of the box task. The experiment where the data was collected was carried out by Professor Gerit Pfuhl and Doctoral Research Fellow Kristoffer Klevjer at UiT The Arctic University of Norway in February 2020. They recruited participants from an undergraduate psychology course. First, they did a practice trial that was a limited trial that terminated after opening three boxes. That means that if they tried to open the fourth, the test terminated, and they could not make a decision. That trial is not analyzed here. Following the practice trial were three unlimited trials. The

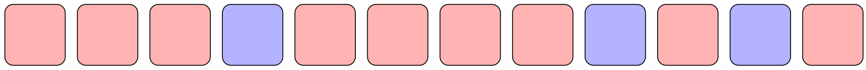

Figure 3.1: The order of the boxes in Trial 2. This is an unlimited trial.

participants could, in these trials, open as many of the twelve boxes as they wanted before deciding on what they think is the dominant colour—lastly, there were six limited trials. Three of them terminated after the participants had opened six boxes, and the other three terminated after they had opened nine boxes. These nine trials are the ones we analyze here, meaning that we analyze Trials 2 to 10. We have data for how many boxes each participant has opened in each of the nine trials. We call this 'draws to decision'. The participants have either opened boxes until they have decided what they think is the majority colour or until the test terminates. We have data for what they chose or whether the test terminated before they were able to choose. To compensate for possible biases towards one colour, the two colours were changed for each new trial. They could, for example, be green and pink in the first trial and blue and yellow in the second trial. For simplicity, we are in this report referring to these colour as blue and red for all trials.

For each trial, there is a fixed sequence of boxes. The participants can only choose whether to open the next box or not; they cannot choose which box they open. Thus, we know how many of the boxes that were blue and how many that were red for each step in the different trials. In Trial 2, which is an unlimited trial, the boxes were opened in the order that is shown in Figure 3.1. In Figure 3.2, the draws to decisions for all participants are shown in a histogram. Here, the number of boxes that are opened when the participant chooses what she thinks is the majority colour is on the horizontal axis. On the vertical axis are the number of participants that have decided on that particular box. We see that many participants have chosen the majority colour after they have opened three boxes. All of these three boxes are red. Thus, there is a high probability that red is the dominant colour. As the participants are told that one of the colours is always in the majority, we can be completely sure if six of the opened boxes are red, that red is the dominant colour. This is because there cannot be six of each box if one of them is in the majority. When seven boxes are opened in Trial 2, six of them are red, and we know then that red is the dominant colour. Seven participants have chosen colour after seven boxes are opened. Some participants wait longer, even though they can be completely sure after seven boxes are opened.

In Trial 3, which is also an unlimited trial, it takes more boxes to be completely sure what the majority colour is. As shown in Figure 3.3, there are

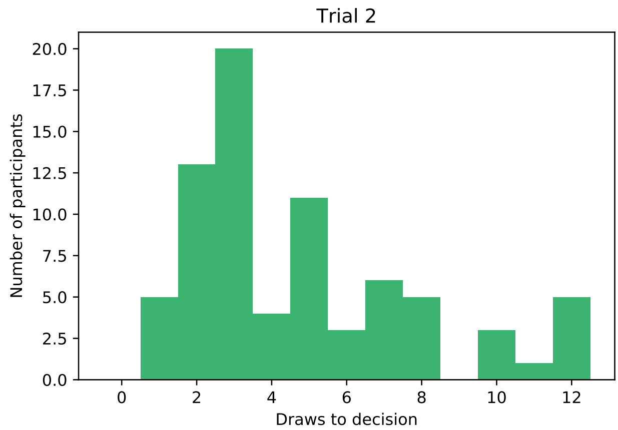

Figure 3.2: Histogram of the draws to decisions for all participants in Trial 2.

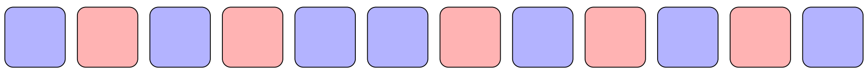

Figure 3.3: The order of the boxes in Trial 3. This is an unlimited trial.

six blue boxes when ten of the boxes are opened. We see in Figure 3.4 that the participants in general open more boxes before choosing the majority colour in this trial than in Trial 2.

Both Trial 5 and Trial 8 are limited trials that terminate after nine boxes are opened. The order of the boxes in Trial 5 is shown in Figure 3.5. When seven boxes are opened, six of them are blue, and we can therefore conclude when seven boxes are opened that blue is the majority colour. We see in Figure 3.7 that many participants choose the majority colour after three boxes are opened. All of these three are blue boxes. In Trial 8, there are never two boxes of the same colour following each other, as shown in Figure 3.6. There are at no point in this trial six of one of the colours, meaning that we never can be completely sure which one is the majority colour. This is reflected in the draws to decision for the participants, as shown in Figure 3.8. We see that the test terminates for many of the participants before choosing what they think is the majority colour.

The order of the boxes for all trials can be found in Appendix A. Here we also have histograms of the draws to decision for all of the trials.

In the following, we formulate a model for how the participants make decisions in the box task, and we estimate parameters such that we can fit

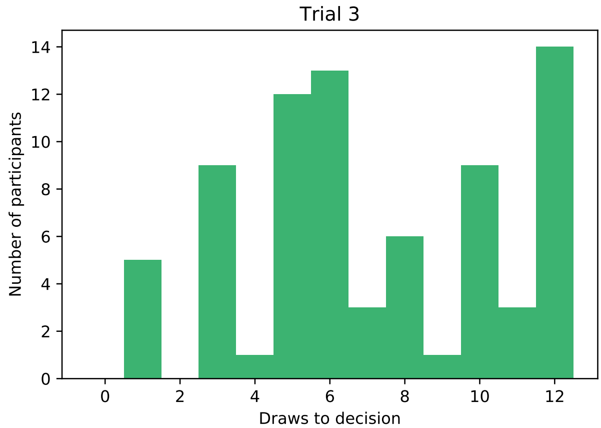

Figure 3.4: Histogram of the draws to decisions for all participants in Trial 3.

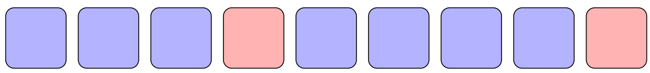

Figure 3.5: The order of the boxes in Trial 5. This is a limited trial that terminates after nine boxes are opened.

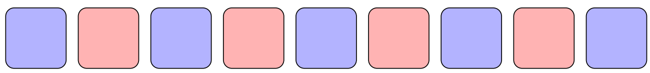

Figure 3.6: The order of the boxes in Trial 8. This is a limited trial that terminates after nine boxes are opened.

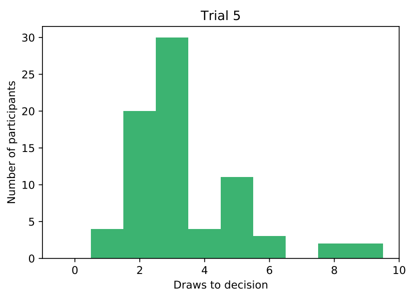

Figure 3.7: The draws to decisions for all participants in Trial 5. That is, how many boxes they open before they choose what they think is the majority colour, or before the test terminates.

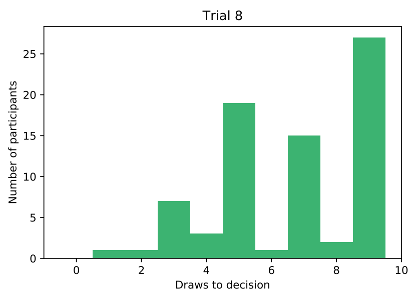

Figure 3.8: The draws to decisions for all participants in Trial 8. That is, how many boxes they open before they choose what they think is the majority colour, or before the test terminates.

the model to each person. We also find a so-called Ideal Observer solution of the box task. An Ideal Observer would always make optimal decisions (He et al., 2013). Thus, an Ideal Observer solution is close to an optimal solution of the box task.

3.1 Modelling Framework

Before we start with the formulation of the model, we will introduce some notation and present some assumptions.

We let X_i be the colour of the i-th opened box as in (2.13). Then, if the box is blue, X_i is zero, and if the box is red, X_i is one. We assume that each X_i has a Bernoulli distribution with success probability Θ, where we later condition on there not being six blue and six red boxes. Then,

$$X_i \sim \text{Bernoulli}(\Theta).$$

We also define a vector, $X_{1:i}$, that contains the colours of the first i boxes that are or will be opened, just as in (2.14). In the same way, we let $x_{1:i} = (x_1, x_2, ..., x_i)$.

Additionally, let U_i be the number of the first i opened boxes that are red. Thus, U_i is a stochastic variable defined as

$$U_i = \sum_{j=1}^{i} X_j. \tag{3.1}$$

The sum of Bernoulli distributed variables is binomially distributed (Casella and Berger, 2002). Thus, U_i is binomially distributed with parameters i and Θ. We define another stochastic variable, V_i, that is the number of red boxes that are not opened when i boxes are opened. Thus, V_i is the number of red boxes out of the $12 - i$ boxes that are not opened, which yields,

$$V_i = \sum_{j=i+1}^{12} X_j. \tag{3.2}$$

This variable is also binomially distributed, but with parameters $12 - i$ and Θ. Thus,

$$\begin{aligned} U_i &\sim \text{Binomial}(i, \Theta) \\ V_i &\sim \text{Binomial}(12 - i, \Theta). \end{aligned} \tag{3.3}$$

Then, we have that

$$P(U_i = u_i | \Theta = \theta) = \binom{12}{u_i} \theta^{u_i} (1 - \theta)^{12 - u_i}, \tag{3.4}$$

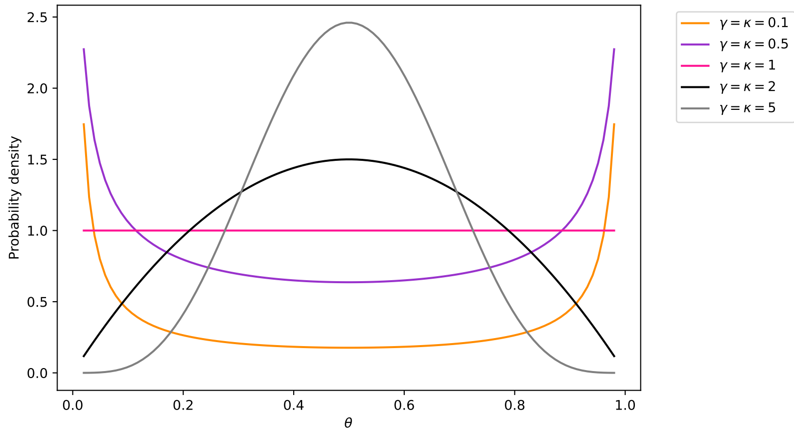

Figure 3.9: The probability density function for the beta distribution plotted for different values of the hyperparameters γ and κ.

and

$$P(V_i = v_i|\Theta = \theta) = \binom{12-i}{v_i}\theta^{v_i}(1-\theta)^{12-i-v_i}. \tag{3.5}$$

Just as in Chapter 2.4, we let Θ have a conjugate beta prior with parameters γ and κ, as shown in (2.9). The prior distribution of Θ is then as given in (2.10).

Figure 3.9 shows the probability density function of the beta distribution for different values of γ and κ. The pink line represents the situation where $\gamma = \kappa = 1$. This is the same as having a uniform prior for Θ. That means that the probability of Θ being anywhere on the interval between zero and one is constant. As the participants are told that one of the colours will be in the majority but get no information about which one, this might be a suitable prior. However, one might argue that our prior beliefs resemble the purple or orange lines as we know that one of the colours will definitively be in the majority. Thus it might not be reasonable to assume that Θ is 0.5. For this reason, we exclude all priors that have γ and κ larger than 1, which is the situation for the black and grey lines.

Of all the 12 boxes, $U_i + V_i$ is the total number of red boxes. Consequently, if $U_i + V_i$ is bigger than 6, it is a red majority in the box task, and if it is smaller than 6, the true majority colour is blue. We denote this true majority colour as Z, such that

$$Z = I(U_i + V_i > 6). \tag{3.6}$$

This is the same as defining Z as in (2.15), as $U_i + V_i = \sum_{j=1}^{12} X_j$, and the order that the boxes are opened in does not affect the majority. Then, as in (2.16), Z is zero if the true majority colour is blue and one if the true majority colour is red.

Each time a box is opened, the participants have three choices. The first is to choose blue as the dominant colour, the second that red is, and the third is to choose to open another box. We denote these decisions as δ_i, where i is the number of opened boxes. If $\delta_i = 0$, the participant chooses that blue is the more prominent colour, thus that there are in total, of all twelve boxes, more blue boxes than red. Moreover, $\delta_i = 1$ means that the participant has chosen that red is the dominant colour, and $\delta_i = 2$ represents the situation where the participant chooses to open the next box. Thus,

$$\delta_i = \begin{cases} 0, & \text{if blue is chosen as majority colour,} \\ 1, & \text{if red is chosen as majority colour,} \\ 2, & \text{if the participant chooses to open the next box.} \end{cases} \tag{3.7}$$

We can define loss functions for each of these decisions. These loss functions depend on the true majority colour, Z, and the decision, δ_i. Similarly to Chapter 2.5, we denote the loss function when i boxes are opened as $L_i[Z, \delta_i; \varphi]$. In our case, we have that

$$\Omega_Z = \{0, 1\}$$

and

$$\Omega_{\delta_i} = \{0, 1, 2\}.$$

If we take the expectation of this loss function, we get the expected loss for each of these decisions when i boxes are opened, which we denote

$$\mathcal{E}^i_{\delta_i}(\varphi) = E\left[L_i[Z, \delta_i; \varphi] | X_{1:i} = x_{1:i}\right],$$

as it depends on some parameters φ. This expected loss also depends on the colours of the i boxes that already are opened, $x_{1:i}$.

In the limited version of the box task, the participants are told that the test will terminate when a random box is opened. Thus, we need a random variable representing how many boxes that are open when the test terminates. We call this variable T. If $T = 3$, the participant has opened three boxes and wants to open the fourth when the test terminates. Then, instead of seeing the colour of the fourth box, the test terminates, and this is a failed trial. The information given to the participants regarding this is that the test will terminate when a random box is opened. We assume that the first box can always be opened, but the probabilities that the test terminates at the subsequent boxes are the same. When 12 boxes are opened, there are no more boxes to open and, therefore, no more chances for the test to terminate. Thus, T is uniformly distributed with on $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$,

$$T \sim \text{Uniform}(\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}). \tag{3.8}$$

Now we have all the notation needed to define the model for how the participants make decisions.

3.2 The Model for the Decisions

Having the notation for the expected losses and decisions, we can define the probability mass function for the decisions using a softmax function similar to the one in (2.21).

For each participant we have observed decisions, $\boldsymbol{\delta} = (\delta_1, \delta_2, ..., \delta_n)$, where $\delta_j \in \{0, 1, 2\}$ as in (3.7), and n is the total amount of decisions we have for each participant. Thus, $j \in \{1, 2, ..., n\}$. As the participants have opened a different amount of boxes each time, n varies form participant to participant. Recall that i is the number of boxes that are opened, and that i is reset for each new trial. Thus, the probability mass function for the decisions can be expressed as

$$f(\delta_j | \varphi, \eta; x_{1:i}) = \frac{\exp(-\eta \mathcal{E}^i_{\delta_j}(\varphi))}{\sum_{d=0}^{2} \exp(-\eta \mathcal{E}^i_d(\varphi))}, \tag{3.9}$$

where η is some parameter. η can be interpreted as a measure of how far the choices the participant makes are away from the decision with the least expected loss. If η is infinity, they always make the decision with the lowest expected loss, and if η is zero, they choose arbitrarily. A negative value of η indicates that the participant tends to choose the decisions with higher expected losses.

When we have this model, we can find estimates of the parameters, φ and η, for each participant such that the model is adapted to each one of the participants. This estimation is done by finding the maximum likelihood estimates (MLEs) as described in Chapter 2.8. We can also find confidence intervals tied to each of the parameters for all of the participants using the bootstrap as described in Chapter 2.9. This will be done in the subsequent sections, but firstly we find an Ideal Observer solution of the box task and use this to find expressions for the loss functions and expected losses.

3.3 Loss Functions

Before we start finding Ideal Observer solutions, we formulate loss functions in the unlimited and limited cases.

3.3.1 Loss Functions in the Unlimited Case

Starting with the unlimited case, we define loss functions for each of the three choices we have when i boxes are opened and put them together as one function. If the participant chooses blue as the majority colour, $\delta_i = 0$, we say that the loss is zero if blue is the true majority colour and one if it is not. Thus, this can be expressed as an indicator function as in (2.11). Recall that the true majority colour is denoted Z. Then, we can express the loss of choosing blue as the majority colour when i boxes are opened as

$$L_i[Z, \delta_i = 0; \varphi] = I(Z \neq 0) = I(Z = 1). \tag{3.10}$$

We define the loss function for when the participant chooses that red is the majority colour, $\delta_i = 1$, similarly to (3.10). This time the loss is zero if the true majority colour is red and one if the blue is the true majority colour. Thus,

$$L_i[Z, \delta_i = 1; \varphi] = I(Z \neq 1) = I(Z = 0). \tag{3.11}$$

We imagine that some participants have some minor penalty or loss of opening another box. That might be because it is tiresome for them to sit through a full trial and they want to finish fast, or that they get some inner reward or feeling of victory when they finish early. A parameter, α, represent this. The loss function for the choice of opening the next box depends on the successive losses. As we do not know the choices that will be made later, we do not know what these losses are. However, we can model these choices as the choices that an Ideal Observer would make. These choices depend on the colour of the next box, X_{i+1}, and the colours of the already opened boxes, $x_{1:i}$. We denote these choices as $IO(x_{1:i}, X_{i+1})$, where $X_{i+1} \in \{0, 1\}$. We define the loss for the decision to open the next box as α plus the loss in the next step. Thus,

$$L_i[Z, \delta_i = 2; \varphi] = \alpha + L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi], \tag{3.12}$$

where $L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi]$ is the loss when the next box is opened.

Putting (3.10), (3.11) and (3.12) together, we get that the total loss function in the unlimited case can be expressed as

$$\begin{aligned} L_i[Z, \delta_i; \varphi] = & I(Z = 0)\, I(\delta_i = 0) \\ & + I(Z = 1)\, I(\delta_i = 1) \\ & + \big(\alpha + L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi]\big) I(\delta_i = 2). \end{aligned} \tag{3.13}$$

Having the loss functions for the unlimited case, we proceed with formulating the loss functions for the limited trials of the box task.

3.3.2 Loss Functions in the Limited Case

The loss functions in the limited case are highly comparable to the ones in the unlimited case. Recall that in a limited trial, the participants might be stopped when a random box opens and that this counts as a failed trial.

Firstly, we have a look at the loss function for choosing blue as the majority colour. We see that this is not dependent on any of the boxes that are not opened in the unlimited case. When i boxes are opened in a limited trial, and the participant chooses that blue is the majority colour, this is, as in the unlimited trial, not affected by the colours of the unopened boxes. If i boxes are opened, and one chooses what the majority colour is here, we know that the test will not terminate, as the participant will not open more boxes. Thus, we can put the loss function for choosing blue as the majority colour in a limited trial as the same as the loss function for choosing blue in an unlimited trial. The loss function is then as in (3.10).

The same argument holds for the loss function for choosing red as the majority colour in a limited trial. Thus, that loss function is the same as in (3.11).

For the choice of opening the next box, we have to consider that the test might terminate. We define a parameter, β, that only appears in the loss function for opening the next box in limited trials. We let it be the loss the participant gets when the test terminates before choosing what the majority colour is. Recall that T is the number of boxes that already are opened when the test terminates and that it is uniformly distributed as in (3.8). The loss when the test does not terminate will be the loss for when the next box is open, in the same way as for the unlimited trials. We can include the event of the test terminating as an indicator function, where an indicator function is as seen in (2.11). Thus, the loss function for opening the next box in an unlimited trial is the loss you get when the next box is opened plus α, times an indicator function that is one if the test does not terminate. I addition to this, we have the loss for when the test terminates, β, times an indicator for that the test terminates. Hence,

$$\begin{aligned} L_i[Z, \delta_i = 2; \varphi] =& \big(\alpha + L_i[Z, IO(x_{1:i}, X_{i+1}); \varphi]\big)\, I(T \neq i) \\ &+ \beta\, I(T = i), \end{aligned} \tag{3.14}$$

where $IO(x_{1:i}, X_{i+1})$ are the choices that an Ideal Observer would do in the next steps.

We get the total loss function in the limited case using (3.10), (3.11) and (3.14), such that

$$\begin{aligned} L_i[Z, \delta_i; \varphi] =& I(Z = 0)\, I(\delta_i = 0) \\ &+ I(Z = 1)\, I(\delta_i = 1) \\ &+ \Big(\big(\alpha + L_i[Z, IO(x_{1:i}, X_{i+1}); \varphi]\big)\, I(T \neq i) \\ &\quad + \beta\, I(T = i)\Big) I(\delta_i = 2) \end{aligned} \tag{3.15}$$

As we have the loss functions in the unlimited and the limited cases, we now continue with finding an Ideal Observer solution of the box task. This solution depends on the expectation of the loss functions, the expected losses.

3.4 Ideal Observer Solution

Having defined loss functions for both the unlimited and the limited versions of the box task, we now want to find an Ideal Observer (IO) solution. We find one for the unlimited case and another for the limited. As stated above, an Ideal Observer acts like a participant that always makes optimal decisions. In the box task, the optimal decision for each opened box is the decision that gives the least expected loss. If one makes the decision with the lowest expected loss each time a box is opened, the total solution is the Ideal Observer solution. Thus, we need to find these expected losses. They depend on the parameter α in the unlimited version of the box task and both α and β in the limited version. Therefore, we get many different IO solutions depending on the values of the parameters. Recall that when i boxes are opened, the participants have three choices as stated in (3.7). We want to find expected losses for all three decisions in both the unlimited and limited versions of the box task.

3.4.1 Expected Losses

When we find the expected losses, we take the expectation of the loss functions like in (2.12). Recall that taking the expectation of an indicator function gives the probability of the event and that $x_{1:i}$ is a vector containing the colours of the i opened boxes.

As stated in Chapter 3.1, the expected losses when i boxes are opened, are denoted $\mathcal{E}^i_{\delta_i}(\varphi)$, with $\delta_i \in (0,1,2)$. We start with the expected loss for choosing blue as the majority colour, $\mathcal{E}^i_0(\varphi)$. This expected loss is the same for both the unlimited and limited trials as the loss functions, stated in (3.10), are identical. We condition on the colours of the opened boxes, as the true majority colour, Z, depends on the colours of all the twelve boxes. Thus,

$$\begin{aligned}\mathcal{E}^i_0(\varphi) =& E\big[L_i[Z,\delta_i=0;\varphi]\big|X_{1:i}=x_{1:i}]\\ =& E\big[I(Z=1)\big|X_{1:i}=x_{1:i}]\\ =& P(Z=1|X_{1:i}=x_{1:i}).\end{aligned} \tag{3.16}$$

We see that the expected loss of choosing blue as the majority colour is equal to the probability that red is the majority colour given the colours of the opened boxes, for both the unlimited and limited versions. The only thing that this expected loss depends on is the colours of the first i boxes, $x_{1:i}$.

We find the expected loss of choosing red as the majority colour when i boxes are opened, $\mathcal{E}^i_1(\varphi)$, similarly to (3.16). Again, conditioning on $X_{1:i}=x_{1:i}$, we get

$$\begin{aligned}\mathcal{E}^i_1(\varphi) =& E\big[L_i[Z,\delta_i=1;\varphi]\big|X_{1:i}=x_{1:i}]\\ =& E\big[I(Z=0)\big|X_{1:i}=x_{1:i}]\\ =& P(Z=0|X_{1:i}=x_{1:i}).\end{aligned} \tag{3.17}$$

The expected loss for choosing red as the majority colour is then the probability that blue is the majority colour, conditioned on $X_{1:i}=x_{1:i}$, for both the unlimited and the limited case.

When we find the expected losses for opening the next box, we have to distinguish between the unlimited and limited cases. Starting with the unlimited case, we continue in the same way as for choosing blue or red as the majority colour, by taking the expectation of the loss function, as it is stated in (3.12), and conditioning on the colours of the i opened boxes. Recall that $IO(x_{1:i},X_{i+1})$ are the choices that an Ideal Observer would make in the next steps. We then get that

$$\mathcal{E}^i_2(\varphi) = E\big[\alpha + L_i[Z, IO(x_{1:i},X_{i+1});\varphi]|X_{1:i}=x_{1:i}\big].$$

Taking the expectation of a constant gives the constant. Thus

$$E[\alpha|X_{1:i}=x_{1:i}]=\alpha,$$

as α is not dependent on the colours of the boxes. Then,

$$\mathcal{E}_2^i(\varphi) = \alpha + E\big[L_i[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}\big]. \tag{3.18}$$

We see that $E\big[L_i[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}\big]$ is the expected loss in the next step, and it depends on the colour of the box that opens, X_{i+1}. We find this expectation using the law of total expectation as in (2.20). Then,

$$\begin{aligned} &E\big[L_i[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}\big] \\ &\quad = E\big[L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}, X_{i+1} = 0\big] \\ &\quad \times P(X_{i+1} = 0 | X_{1:i} = x_{1:i}) \\ &\quad + E\big[L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}, X_{i+1} = 1\big] \\ &\quad \times P(X_{i+1} = 1 | X_{1:i} = x_{1:i}), \end{aligned} \tag{3.19}$$

where $P(X_{i+1} = 0 | X_{1:i} = x_{1:i})$ is the probability that box $i+1$ is blue given the colours of the first i boxes and $P(X_{i+1} = 1 | X_{1:i} = x_{1:i})$ is the probability that it is red.

Inserting (3.19) into (3.18), we get that the expected loss for opening the next box in the unlimited case is

$$\begin{aligned} \mathcal{E}_2^i(\varphi) = \alpha + &E\big[L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}, X_{i+1} = 0\big] \\ &\times P(X_{i+1} = 0 | X_{1:i} = x_{1:i}) \\ &+ E\big[L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}, X_{i+1} = 1\big] \\ &\times P(X_{i+1} = 1 | X_{1:i} = x_{1:i}). \end{aligned} \tag{3.20}$$

Note that

$$E\big[L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}, X_{i+1} = 0\big] = \mathcal{E}^{i+1}_{IO(x_{1:i},0)}(\varphi)$$

and

$$E\big[L_{i+1}[Z, IO(x_{1:i}, X_{i+1}); \varphi] | X_{1:i} = x_{1:i}, X_{i+1} = 1\big] = \mathcal{E}^{i+1}_{IO(x_{1:i},1)}(\varphi).$$

The expression for the expected loss of opening the next box in the unlimited case, (3.20), is then

$$\begin{aligned} \mathcal{E}_2^i(\varphi) = \alpha + &\mathcal{E}^{i+1}_{IO(x_{1:i},0)}(\varphi) P(X_{i+1} = 0 | X_{1:i} = x_{1:i}) \\ &+ \mathcal{E}^{i+1}_{IO(x_{1:i},1)}(\varphi) P(X_{i+1} = 1 | X_{1:i} = x_{1:i}). \end{aligned} \tag{3.21}$$

In the unlimited case, the expected loss depends on the parameter α. Thus, $\varphi = \alpha$.

We proceed in a similar manner when we find the expected loss of opening the next box in the limited case. Taking the expectation of the loss function in (3.14), we get the expected loss when i boxes are opened given that the test has not terminated yet, $\mathcal{E}_2^i(\varphi)$. Recall that if the box terminates when i boxes already are open, then the parameter T is equal to i. We have to condition on T being greater than or equal to i when we find the expected loss, meaning that the test has not terminated yet when i boxes are open. Using (3.14) we get that

$$\begin{aligned}\mathcal{E}_2^i(\varphi) =& E\big[L_i[Z,\delta_i=2;\varphi] \mid X_{1:i}=x_{1:i}, T\geq i\big]\\ =& E\big[\big(\alpha + L_{i+1}[Z, IO(x_{1:i}, X_{i+1});\varphi]\big)\, I(T\neq i)\\ &+\beta\, I(T=i) \mid X_{1:i}=x_{1:i}, T\geq i\big].\end{aligned} \tag{3.22}$$

We start with the first term in (3.22). When the test terminates is independent of the colours of the boxes, such that T is independent of $x_{1:i}$. The indicator function will then be the probability of $T\neq i$, whereas, for the expectation of the loss function when $i+1$ boxes are opened, we use the law of total expectation as in (2.20), and condition on the colour of the next box, X_{i+1}.

We also have that $\mathcal{E}_2^{i+1}(\varphi, X_{i+1}=j)$ is the expected loss in the next step given the colours of the i opened boxes, the colour of box $i+1$ and given that the test has not terminated yet. Thus,

$$\begin{aligned}&E\big[\big(\alpha + L_{i+1}[Z, IO(x_{1:i}, X_{i+1});\varphi]\big)\, I(T\neq i)|X_{1:i}=x_{1:i}, T\geq i\big]\\ &=\Big(\alpha+\sum_{j=0}^{1}\mathcal{E}_2^{i+1}(\varphi, X_{i+1}=j)P(X_{i+1}=j|X_{1:i}=x_{1:i}, T\geq i)\Big)\\ &\quad\times P(T\neq i|T\geq i)\end{aligned} \tag{3.23}$$

As the x's and T are independent, we have that

$$P(X_{i+1}=j|X_{1:i}=x_{1:i}, T\geq i) = P(X_{i+1}=j|X_{1:i}=x_{1:i}). \tag{3.24}$$

Putting (3.24) into (3.23), we get

$$\begin{aligned}&E\big[\big(\alpha + L_{i+1}[Z, IO(x_{1:i}, X_{i+1});\varphi]\big)\, I(T\neq i)|X_{1:i}=x_{1:i}, T\geq i\big]\\ &=\Big(\alpha+\sum_{j=0}^{1}\mathcal{E}_2^{i+1}(\varphi, X_{i+1}=j)P(X_{i+1}=j|X_{1:i}=x_{1:i})\Big)\\ &\quad\times P(T\neq i|T\geq i).\end{aligned} \tag{3.25}$$

The last term in (3.22) becomes

$$E\big[\beta\, I(T=i)|X_{1:i}=x_{1:i}, T\geq i\big] = \beta\, P(T=i|T\geq i) \tag{3.26}$$

as it does not depend on the colours of the boxes, $x_{1:i}$.

Putting (3.25) and (3.26) into (3.22), we get that the expected loss for opening another box in the limited case is

$$\begin{aligned}\mathcal{E}_2^i(\varphi) = &\Big(\alpha + \sum_{j=0}^{1} \mathcal{E}_2^{i+1}(\varphi, X_{i+1} = j) P(X_{i+1} = j | X_{1:i} = x_{1:i})\Big) \\ &\times P(T \neq i | T \geq i) \\ &+ \beta\, P(T = i | T \geq i).\end{aligned} \tag{3.27}$$

The expected losses in the limited cases depend on α and β, thus in this case we have that $\varphi = (\alpha, \beta)$.

Now that we have expressions for the expected losses, we have to find the probabilities in these expressions.

3.4.2 Probabilities

As we now have expressions for the expected losses, we find the probabilities needed for finding the expected losses. That is $P(Z = 1 | X_{1:i} = x_{1:i})$, $P(Z = 0 | X_{1:i} = x_{1:i})$, $P(X_{i+1} = 1 | X_{1:i} = x_{1:i})$, $P(X_{i+1} = 0 | X_{1:i} = x_{1:i})$, $P(T \neq i | T \geq i)$ and $P(T = i | T \geq i)$.

The Majority Colour

When we find the probabilities used in the expressions for the expected losses, we start with the expected loss for choosing blue as the majority colour, as given in (3.16). Then we need the probability

$$P(Z = 1 | X_{1:i} = x_{1:i}). \tag{3.28}$$

This is the probability that red is the majority colour, given the colours of the boxes that already are observed. Using the definition of Z as it is in (3.6), we can express (3.28) using U_i and V_i. Recall that they are defined as in (3.1) and (3.2), respectively. (3.28) can then be expressed as the probability that $U_i + V_i > 6$, or that $U_i + V_i \geq 7$, given the colours of the i first boxes. However, we also need to condition on $U_i + V_i \neq 6$, as we know that one of the colours always is in majority, such that there will never be six blue and six red boxes all together. Thus, we find $P(U_i + V_i \geq 7 | X_{1:i} = x_{1:i}, U_i + V_i \neq 6)$. As the order the boxes have been opened in is irrelevant here, and $U_i = \sum_{j=1}^{i} X_j$, we use $U_i = u_i$ instead of $X_{1:i} = x_{1:i}$, to be consistent with the other notation. Thus, we have that

$$P(Z = 1 | X_{1:i} = x_{1:i}) = P(U_i + V_i \geq 7 | U_i = u_i, U_i + V_i \neq 6). \tag{3.29}$$

Using Bayes rule as described in (2.2), we get that

$$
\begin{aligned}
&P(U_i + V_i \geq 7 | U_i = u_i, U_i + V_i \neq 6) \\
&\quad = \frac{P(U_i + V_i \neq 6 | U_i = u_i, U_i + V_i \geq 7) P(U_i + V_i \geq 7 | U_i = u_i)}{P(U_i + V_i \neq 6 | U_i = u_i)}.
\end{aligned} \tag{3.30}
$$

As

$$
P(U_i + V_i \neq 6 | U_i = u_i, U_i + V_i \geq 7) = 1,
$$

we get that

$$
P(U_i + V_i \geq 7 | U_i = u_i, U_i + V_i \neq 6) = \frac{P(U_i + V_i \geq 7 | U_i = u_i)}{P(U_i + V_i \neq 6 | U_i = u_i)}. \tag{3.31}
$$

We have that

$$
P(U_i + V_i \geq 7 | U_i = u_i) = \sum_{j=7}^{12} P(U_i + V_i = j | U_i = u_i) \tag{3.32}
$$

Thus, to be able to find $P(U_i + V_i \geq 7 | U_i = u_i)$, we start with finding $P(U_i + V_i = j | U_i = u_i)$. Using the law of total probability as in (2.1), and conditioning on θ, we get

$$
\begin{aligned}
&P(U_i + V_i = j | U_i = u_i) \\
&\quad = \int_0^1 P(U_i + V_i = j | U_i = u_i, \Theta = \theta) f(\theta | U_i = u_i) \mathrm{d}\theta \\
&\quad = \int_0^1 P(V_i = j - u_i | \Theta = \theta) f(\theta | U_i = u_i) \mathrm{d}\theta
\end{aligned} \tag{3.33}
$$

Thus, we need to find $P(V_i = j - u_i | \Theta = \theta)$ and $f(\theta | U_i = u_i)$.

Since V_i has a binomial distribution as in (3.3), we get that

$$
P(V_i = j - u_i | \Theta = \theta) = \binom{12 - i}{j - u_i} \theta^{j - u_i} (1 - \theta)^{12 - i - (j - u_i)} \tag{3.34}
$$

We can find $f(\theta | U_i = u_i)$ using Bayes rule as given in (2.2). Hence,

$$
f(\theta | U_i = u_i) = \frac{P(U_i = u_i | \Theta = \theta) f(\theta)}{P(U_i = u_i)},
$$

which is proportional to the numerator of the right-hand side as in (2.7). Using that $U_i | \Theta$ has a binomial distribution with probability mass function as in (3.4), and that Θ has a beta prior, with density function as in (2.10), we get that

$$\begin{aligned} f(\theta|U_i = u_i) &\propto P(U_i = u_i|\Theta = \theta) f(\theta) \\ &\propto \theta^{u_i}(1-\theta)^{i-u_i}\theta^{\gamma-1}(1-\theta)^{\kappa-1} \\ &= \theta^{u_i+\gamma-1}(1-\theta)^{i-u_i+\kappa-1}. \end{aligned}$$

This is proportional to the density of a beta distribution with parameters $u_i + \gamma$ and $i - u_i + \kappa$. Hence, we can conclude that

$$\Theta|U_i \sim \text{Beta}(u_i + \gamma, i - u_i + \kappa),$$

and therefore that

$$f(\theta|U_i = u_i) = \frac{1}{\text{B}(u_i + \gamma, i - u_i + \kappa)}\theta^{u_i+\gamma-1}(1-\theta)^{i-u_i+\kappa-1}. \tag{3.35}$$

We now have expressions for $P(V_i = j - u_i|\Theta = \theta)$ and $f(\theta|U_i = u_i)$, as given in (3.34) and (3.35), respectively. We put these into (3.33), and get

$$\begin{aligned} &P(U_i + V_i = j|U_i = u_i) \\ &= \int_0^1 P(V_i = j - u_i|\Theta = \theta)P(\Theta = \theta|U_i = u_i)\mathrm{d}\theta \\ &= \int_0^1 \binom{12-i}{j-u_i}\theta^{j-u_i}(1-\theta)^{12-i-(j-u_i)}\frac{\theta^{u_i+\gamma-1}(1-\theta)^{i-u_i+\kappa-1}}{\text{B}(u_i + \gamma, i - u_i + \kappa)}\mathrm{d}\theta. \end{aligned} \tag{3.36}$$

Taking the parts that do not depend on θ outside of the integral and summing the exponents of θ and $(1 - \theta)$, we get that (3.36) is

$$\begin{aligned} &P(U_i + V_i = j|U_i = u_i) \\ &= \frac{\binom{12-i}{j-u_i}}{\text{B}(u_i + \gamma, i - u_i + \kappa)} \\ &\quad \times \int_0^1 \theta^{j-u_i+u_i+\gamma-1}(1-\theta)^{12-i-(j-u_i)+i-u_i+\kappa-1}\mathrm{d}\theta \\ &= \frac{\binom{12-i}{j-u_i}}{\text{B}(u_i + \gamma, i - u_i + \kappa)}\int_0^1 \theta^{j+\gamma-1}(1-\theta)^{12-j+\kappa-1}\mathrm{d}\theta. \end{aligned} \tag{3.37}$$

The part inside the integral is proportional to the density of a beta distribution with parameters $j + \gamma$ and $12 - j + \kappa$. The integral of a density over the parameter space is one, hence

$$\int_0^1 \frac{1}{\mathrm{B}(j+\gamma, 12-j+\kappa)} \theta^{j+\gamma-1}(1-\theta)^{12-j+\kappa} \mathrm{d}\theta = 1.$$

Therefore,

$$\int_0^1 \theta^{j+\gamma-1}(1-\theta)^{12-j+\kappa} \mathrm{d}\theta = \mathrm{B}(j+\gamma, 12-j+\kappa). \tag{3.38}$$

Putting (3.38) into (3.37), we get

$$P(U_i+V_i = j | U_i = u_i) = \binom{12-i}{j-u_i} \frac{\mathrm{B}(j+\gamma, 12-j+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}. \tag{3.39}$$

Putting (3.39) into (3.32), we get that

$$P(U_i+V_i \geq 7 | U_i = u_i) = \sum_{j=7}^{12} \binom{12-i}{j-u_i} \frac{\mathrm{B}(j+\gamma, 12-j+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}. \tag{3.40}$$

We have that

$$P(U_i+V_i \neq 6 | U_i = u_i) = 1 - P(U_i+V_i = 6 | U_i = u_i)$$

and, using (3.39), we get

$$\begin{aligned} P(U_i+V_i \neq 6 | U_i = u_i) &= 1 - \binom{12-i}{6-u_i} \frac{\mathrm{B}(6+\gamma, 12-6+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)} \\ &= 1 - \binom{12-i}{6-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}. \end{aligned} \tag{3.41}$$

Putting (3.41) and (3.40) into (3.31), we get

$$\begin{aligned} & P(U_i+V_i \geq 7 | U_i = u_i, U_i+V_i \neq 6) \\ &= \frac{\sum_{j=7}^{12} \binom{12-i}{j-u_i} \frac{\mathrm{B}(j+\gamma, 12-j+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}}{1 - \binom{12-i}{6-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}}. \end{aligned} \tag{3.42}$$

This is the probability that there is a red majority in total, given the colour of the first i boxes that are opened, and given that one of the colours is in the majority. This is also the expected loss of choosing blue as the majority colour.

In the expected loss for choosing red as the majority colour, we have $P(Z = 0|X_{1:i} = x_{1:i})$, as in (3.17). The same argument holds here as in (3.29). Thus, we have that

$$P(Z = 0|X_{1:i} = x_{1:i}) = P(U_i + V_i \leq 5|U_i = u_i, U_i + V_i \neq 6). \tag{3.43}$$

This is the probability that blue is the majority colour, which is the complementing probability to the probability that red is the majority colour. Therefore,

$$\begin{aligned} &P(U_i + V_i \leq 5|U_i = u_i, U_i + V_i \neq 6) \\ &\qquad = 1 - P(U_i + V_i \geq 7|U_i = u_i, U_i + V_i \neq 6). \end{aligned} \tag{3.44}$$

Putting the expression in (3.42) into (3.44), we get that the probability of blue being the dominant colour is

$$\begin{aligned} &P(U_i + V_i \leq 5|U_i = u_i, U_i + V_i \neq 6) \\ &\qquad = 1 - \frac{\sum_{j=7}^{12} \binom{12-i}{j-u_i} \frac{\mathrm{B}(j+\gamma, 12-j+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}}{1 - \binom{12-i}{6-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}}, \end{aligned} \tag{3.45}$$

which also it the expected loss of choosing red as the majority colour.

The Colour of the Next Box

We now have a look at the expected losses for opening the next box, both in the unlimited and the limited cases, as given in (3.21) and (3.27), respectively. In both expressions, we have the probability that the next box is either red or blue, given the colours of the opened boxes. These probabilities are $P(X_{i+1} = 1|X_{1:i} = x_{1:i})$ and $P(X_{i+1} = 0|X_{1:i} = x_{1:i})$, where

$$P(X_{i+1} = 0|X_{1:i} = x_{1:i}) = 1 - P(X_{i+1} = 1|X_{1:i} = x_{1:i}), \tag{3.46}$$

as there are only two possible colours the box could have. Thus, we find the probability that that the next box is red, and can then easily find the probability of it being blue using (3.46).

Again, we change the notation from $X_{1:i} = x_{1:i}$ to $U_i = u_i$ and $V_i = v_i$, with the same argument as for (3.29). Thus,

$$P(X_{i+1} = 1|X_{1:i} = x_{1:i}) = P(X_{i+1} = 1|U_i = u_i, U_i + V_i \neq 6) \tag{3.47}$$

Using Bayes' rule, stated in (2.2), we get that this is

$$\begin{aligned} &P(X_{i+1} = 1|U_i = u_i, U_i + V_i \neq 6) \\ &\qquad = \frac{P(U_i + V_i \neq 6|U_i = u_i, X_{i+1} = 1)P(X_{i+1} = 1|U_i = u_i)}{P(U_i + V_i \neq 6|U_i = u_i)}, \end{aligned} \tag{3.48}$$

where the expression in the denominator, $P(U_i + V_i \neq 6|U_i = u_i)$, is as given in (3.41).

We start by finding $P(X_{i+1} = 1|U_i = u_i)$. Using the law of total probability that is stated in (2.1), and conditioning on θ, we get

$$P(X_{i+1} = 1|U_i = u_i) = \int_0^1 P(X_{i+1} = 1|U_i = u_i, \Theta = \theta) \times f(\theta|U_i = u_i)\,\mathrm{d}\theta. \tag{3.49}$$

The expression for $f(\theta|U_i = u_i)$ is given in (3.35). All of the x's are Bernoulli distributed with probability θ, and they are conditionally independent of each other, given θ. Therefore, the probability that X_{i+1} is one, or red, is independent of the colour of the already opened boxes. The probability that a box that is opened is one is also equal to θ. Hence,

$$P(X_{i+1} = 1|U_i = u_i, \Theta = \theta) = P(X_{i+1} = 1|\Theta = \theta) = \theta. \tag{3.50}$$

Putting (3.50) and (3.35) into (3.49) gives

$$\begin{aligned} P(X_{i+1} &= 1|U_i = u_i) \\ &= \int_0^1 \theta \frac{1}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)} \theta^{u_i+\gamma-1}(1-\theta)^{i-u_i+\kappa-1}\mathrm{d}\theta \\ &= \frac{1}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)} \int_0^1 \theta^{u_i+\gamma}(1-\theta)^{i-u_i+\kappa-1}\mathrm{d}\theta. \end{aligned} \tag{3.51}$$

Again, the part inside the integral is proportional to the density of a beta distribution, here with parameters $u_i + \gamma + 1$ and $i - u_i + \kappa$. Integrating a distribution over the parameter space gives one, which in this case gives

$$\int_0^1 \frac{1}{\mathrm{B}(u_i+\gamma+1, i-u_i+\kappa)} \theta^{u_i+\gamma}(1-\theta)^{i-u_i+\kappa-1}\mathrm{d}\theta = 1.$$

Hence,

$$\int_0^1 \theta^{u_i+\gamma}(1-\theta)^{i-u_i+\kappa-1}\mathrm{d}\theta = \mathrm{B}(u_i+\gamma+1, i-u_i+\kappa). \tag{3.52}$$

Inserting (3.52) into (3.51) gives

$$P(X_{i+1} = 1|U_i = u_i) = \frac{\mathrm{B}(u_i+\gamma+1, i-u_i+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}. \tag{3.53}$$

Using the property of the beta function as stated in (2.6), we get that the numerator in (3.53) is

$$\mathrm{B}(u_i+\gamma+1, i-u_i+\kappa) = \frac{\Gamma(u_i+\gamma+1)\Gamma(i-u_i+\kappa)}{\Gamma(u_i+\gamma+1+i-u_i+\kappa)}, \tag{3.54}$$

and that the denominator is

$$\mathrm{B}(u_i+\gamma, i-u_i+\kappa) = \frac{\Gamma(u_i+\gamma)\Gamma(i-u_i+\kappa)}{\Gamma(u_i+\gamma+i-u_i+\kappa)}. \tag{3.55}$$

Inserting (3.54) and (3.55) into (3.53), gives

$$\begin{aligned} P(X_{i+1}=1|U_i=u_i) &= \frac{\frac{\Gamma(u_i+\gamma+1)\Gamma(i-u_i+\kappa)}{\Gamma(u_i+\gamma+1+i-u_i+\kappa)}}{\frac{\Gamma(u_i+\gamma)\Gamma(i-u_i+\kappa)}{\Gamma(u_i+\gamma+i-u_i+\kappa)}} \\ &= \frac{\frac{\Gamma(u_i+\gamma+1)}{\Gamma(\gamma+1+i+\kappa)}}{\frac{\Gamma(u_i+\gamma)}{\Gamma(\gamma+i+\kappa)}}. \end{aligned} \tag{3.56}$$

Using the recursive property of the gamma function as seen in (2.4), we get that the nominator in (3.56) is

$$\frac{\Gamma(u_i+\gamma+1)}{\Gamma(\gamma+1+i+\kappa)} = \frac{(\gamma+u_i)\Gamma(u_i+\gamma)}{(\gamma+\kappa+i)\Gamma(\gamma+i+\kappa)}. \tag{3.57}$$

Inserting (3.57) into (3.56), we get

$$\begin{aligned} P(X_{i+1}=1|U_i=u_i) &= \frac{\frac{(\gamma+u_i)\Gamma(u_i+\gamma)}{(\gamma+\kappa+i)\Gamma(\gamma+i+\kappa)}}{\frac{\Gamma(u_i+\gamma)}{\Gamma(\gamma+i+\kappa)}} \\ &= \frac{\gamma+u_i}{\gamma+\kappa+i}. \end{aligned} \tag{3.58}$$

In the expression in (3.48), it remains to find $P(U_i+V_i \neq 6|U_i=u_i, X_{i+1}=1)$. Firstly,

$$\begin{aligned} &P(U_i+V_i \neq 6|U_i=u_i, X_{i+1}=1) \\ &\quad = P(U_i+V_i \neq 6|U_{i+1}=u_i+1) \\ &\quad = P(U_{i+1}+V_{i+1} \neq 6|U_{i+1}=u_i+1). \end{aligned} \tag{3.59}$$

Using (3.41), we get that

$$\begin{aligned} &P(U_{i+1}+V_{i+1} \neq 6|U_{i+1}=u_i+1) \\ &= 1-\binom{12-(i+1)}{6-(u_i+1)} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+1+\gamma, i-(u_i+1)+\kappa)} \\ &= 1-\binom{11-i}{5-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+1+\gamma, i-u_i-1+\kappa)}. \end{aligned} \tag{3.60}$$

Then, using (3.59) and (3.60), we get that

$$\begin{aligned} &P(U_i + V_i \neq 6 | U_i = u_i, X_{i+1} = 1) \\ &\quad = 1 - \binom{11-i}{5-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+\gamma+1, i-u_i+\kappa)} \end{aligned} \tag{3.61}$$

Inserting (3.61), (3.58) and (3.41) into (3.48), we get

$$\begin{aligned} &P(X_{i+1} = 1 | U_i = u_i, U_i + V_i \neq 6) \\ &\quad = \frac{P(U_i + V_i \neq 6 | U_i = u_i, X_{i+1} = 1) P(X_{i+1} = 1 | U_i = u_i)}{P(U_i + V_i \neq 6 | U_i = u_i)} \\ &\quad = \frac{\left[1 - \binom{11-i}{5-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(\gamma+u_i+1, \kappa+i-u_i)}\right] \frac{\gamma+u_i}{\gamma+\kappa+i}}{1 - \binom{12-i}{6-u_i} \frac{\mathrm{B}(6+\gamma, 6+\kappa)}{\mathrm{B}(u_i+\gamma, i-u_i+\kappa)}}. \end{aligned} \tag{3.62}$$

As we now have the probability that the next box that is opened is red, we can find the probability that the next box that opens is blue using (3.46). Thus, we have both the probabilities for which colour the next box is.

When the Test Terminates

In (3.27) we see the expected loss for opening another box in the limited case. Here we have the probability that the test terminates when i boxes are opened given that the test has not terminated yet and the probability that it does not terminate when i boxes are opened given that it has not terminated yet. These are $P(T = i | T \geq i)$ and $P(T \neq i | T \geq i)$, which are complementary probabilities, such that

$$P(T \neq i | T \geq i) = 1 - P(T = i | T \geq i). \tag{3.63}$$

Thus, if we find $P(T = i | T \geq i)$, we can easily find $P(T \neq i | T \geq i)$ using (3.63).

Using Bayes' rule as it is given in (2.2), we get

$$P(T = i | T \geq i) = \frac{P(T \geq i | T = i) P(T = i)}{P(T \geq i)}. \tag{3.64}$$

We see that

$$P(T \geq i | T = i) = 1. \tag{3.65}$$

As T is uniformly distributed with 11 possible values, as in (3.8), we have that

$$P(T = i) = \frac{1}{11}. \tag{3.66}$$

It then remains to find $P(T \geq i)$. We have that

$$P(T \geq i) = 1 - P(T < i) = 1 - \sum_{j=1}^{i-1} P(T = j). \tag{3.67}$$

As in (3.66), we have that $P(T = j)$ is $\frac{1}{11}$. Thus, (3.67) becomes

$$\begin{aligned} P(T \geq i) =& 1 - \sum_{j=1}^{i-1} \frac{1}{11} = 1 - (i-1)\frac{1}{11} \\ =& \frac{11 - (i-1)}{1} = \frac{12 - i}{11}. \end{aligned} \tag{3.68}$$

Inserting (3.65), (3.66) and (3.68) into (3.64), we get

$$P(T = i | T \geq i) = \frac{\frac{1}{11}}{\frac{12-i}{11}} = \frac{1}{12 - i}. \tag{3.69}$$

We now have all that we need to find the three expected losses each time a box is opened in both unlimited and limited trials. Then we find the Ideal Observer solutions by always choosing the decision with the least expected loss each time a box is opened.

3.5 Maximum Likelihood Estimators

As we have defined a model for the participants' decisions and found expressions for the expected losses for each of the three decisions, we can now fit the model to each participant. We do this by finding maximum likelihood estimates of α and η in the unlimited case and α, β and η in the limited case, based on the decisions the participants have made. We generalise the situation to fit both the unlimited and limited case and denote α and β as φ.

We can find the likelihood, $L(\varphi, \eta | \boldsymbol{\delta})$ as in (2.22). If we have n decisions for each participant, denoted $\boldsymbol{\delta}$, we get that the likelihood is

$$L(\varphi, \eta | \boldsymbol{\delta}) = \prod_{j=1}^{n} f(\delta_j | \varphi, \eta) \tag{3.70}$$

Using the model as it is defined in (3.9) and the expected losses as they are formulated in Chapter 3.4.1, (3.70) becomes

$$L(\varphi, \eta|\boldsymbol{\delta}) = \prod_{j=1}^{n} \frac{\exp(-\eta \mathcal{E}^i_{\delta_j}(\varphi))}{\sum_{d=0}^{2} \exp(-\eta \mathcal{E}^i_d(\varphi))}. \tag{3.71}$$

We then find the log likelihood function, $l(\varphi, \eta|\boldsymbol{\delta})$, by taking the logarithm of (3.71) as in (2.23) and (2.24). Then,

$$\begin{aligned} l(\varphi, \eta|\boldsymbol{\delta}) &= \sum_{j=1}^{n} \log(f(\delta_j|\varphi, \eta)) \\ &= \sum_{j=1}^{n} \log\Big(\frac{\exp(-\eta \mathcal{E}^i_{\delta_j}(\varphi))}{\sum_{d=0}^{2} \exp(-\eta \mathcal{E}^i_d(\varphi))}\Big). \end{aligned} \tag{3.72}$$

Using the property of the logarithm that

$$\log\Big(\frac{a}{b}\Big) = \log(a) - \log(b),$$

we get that the log likelihood is

$$\begin{aligned} l(\varphi, \eta|\boldsymbol{\delta}) = \sum_{j=1}^{n} \Bigg(& \log\big(\exp(-\eta \mathcal{E}^i_{\delta_j}(\varphi))\big) \\ & -\log\big(\sum_{d=0}^{2} \exp(-\eta \mathcal{E}^i_d(\varphi))\big)\Bigg). \end{aligned} \tag{3.73}$$

Using that the last term inside the sum does not depend on j, and that

$$\log\big(\exp(-\eta \mathcal{E}^i_{\delta_j}(\varphi))\big) = -\eta \mathcal{E}^i_{\delta_j}(\varphi),$$

we can write (3.73) as

$$l(\varphi, \eta|\boldsymbol{\delta}) = \sum_{j=1}^{n} \Big(-\eta \mathcal{E}^i_{\delta_j}(\varphi)\Big) - n \log\big(\sum_{d=0}^{2} \exp(-\eta \mathcal{E}^i_d(\varphi))\big). \tag{3.74}$$

It does not matter if we maximise the likelihood or the log-likelihood because maximising (3.71) and (3.74) with respect to φ and η, gives the same estimates for the parameters. These are the maximum likelihood estimates, and we denote them as $\hat{\varphi}$ and $\hat{\eta}$, respectively. We maximize (3.74) here. Thus, we find

$$\hat{\varphi}, \hat{\eta} = \arg\max_{\varphi, \eta} \; l(\varphi, \eta|\boldsymbol{\delta}). \tag{3.75}$$

Recall that in the unlimited trials we have $\varphi = \alpha$, such that we find $\hat{\alpha}$ and $\hat{\eta}$ for all participants. In the limited trials we have that $\varphi = \alpha, \beta$, meaning that we find $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\eta}$ for all participants.

To maximise the log-likelihood, we use a version of the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm called L-BFGS-B. This algorithm uses less memory than the original BFGS and allows bounds on the parameters (Zhu et al., 1997). It is an extension of the limited memory BFGS (L-BFGS) algorithm that does not allow bounds on the parameters. L-BFGS-B is mostly used for nonlinear optimisation problems with bounded variables where it might be hard to find the Hessian matrix. As we want to maximise (3.74) and this depends on the expected losses, it is hard to find the Hessian matrix. Especially the expected loss for opening the next box, $\mathcal{E}_2^i(\varphi)$, makes this challenging as it depends on the expected losses in the next steps. We also have bounds on α and β. Recall that they are defined as the loss one gets when a new box is opened, and the loss one gets when the test terminates, respectively. If α or β were negative, this indicates some reward of opening the next box or the test terminating, which we believe is not likely. Therefore, the L-BFGS-B algorithm is a natural choice of optimisation algorithm. The L-BFGS-B is used for minimising and not maximising. Thus we minimise the negative of the log-likelihood instead of maximising the log-likelihood. We use Python to do this. To avoid finding local minimum points, we use several starting values and choose the results tied to the lowest value of the negative log-likelihood.

When we have the MLEs for α, β and η, it remains to find confidence intervals for the parameters.

3.6 Confidence Intervals for the Parameters

To say something about the uncertainty of the parameter estimates we have found, we now find confidence intervals for each of them. We then use parametric bootstrapping, which is described in Chapter 2.9.

Consider a person with parameter estimates $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\eta}$. Recall that $\hat{\varphi} = \hat{\alpha}$ in the unlimited case and $\hat{\varphi} = \hat{\alpha}, \hat{\beta}$ in the limited case. Using the softmax model, we can find the probabilities for each of the three choices in all the steps in the nine trials. That is, we find the probability that the participant chooses blue as the majority colour, that she chooses red and that she chooses to open the next box. These probabilities are then

$$P(\delta_j | \hat{\varphi}, \hat{\eta}, x_{1:i}) = \frac{\exp(-\hat{\eta}\mathcal{E}_{\delta_j}^i(\hat{\varphi}))}{\sum_{d=0}^{2} \exp(-\hat{\eta}\mathcal{E}_d^i(\hat{\varphi}))}, \tag{3.76}$$

where $\delta_j \in \{0, 1, 2\}$. For each opened box, we can simulate the decisions the participant would make using these probabilities. Then we get whole new sequences of decisions. Consider, for example, Trial 2, where the order of the boxes is as shown in Figure 3.1. Before any boxes are opened, we can imagine that the probability that the participant chooses to open the next box, that $\delta_0 = 2$, is relatively high, and the other two probabilities quite small and of equal size. The first box that opens is a red one, $X_1 = 1$. The probability that she chooses red as the majority colour is higher than the probability that she chooses blue, but whether it is higher than the probability of $\delta_1 = 2$ or not depends on the parameter estimates. We find these probabilities and draw decisions until the participant in the simulated trial has chosen what the majority colour is. We do that for each of the nine trials and then end up with a new set of simulated decisions. In the limited trials, we either stop when $\delta_i = 0$ or $\delta_i = 1$, or when the test terminates. If the test terminates, this is a failed trial, and the loss is β.

For those simulated decisions, we find new estimates for the parameters, again using maximum likelihood estimation as in Chapter 3.5. We simulate these decisions and find the new MLEs 1000 times. We then have 1000 bootstrap samples for all the parameters for each participant. These MLEs of the decisions in simulated trials are denoted

$$\{\varphi^{*(b)}, b = 1, 2, ..., 1000\}$$

and

$$\{\eta^{*(b)}, b = 1, 2, ..., 1000\}.$$

For each of these parameters, we use the percentile method to construct confidence intervals as described in Chapter 2.9.1. We find 90% CIs for each of the parameters, and we denote them as

$$[\hat{\varphi}^{*(5)}_{1000}, \hat{\varphi}^{*(95)}_{1000}]$$

and

$$[\hat{\eta}^{*(5)}_{1000}, \hat{\eta}^{*(95)}_{1000}].$$

Thus, we are finding the 5-th and 95-th percentiles.

We do this for each participant and get confidence intervals for α and η in the unlimited case and α, β and η in the limited case for all 76 participants.

We now have an Ideal Observer solution of the box task and parameter estimates and confidence intervals for all three parameters for all of the 76 participants. Next, we will show some results.

Chapter 4

Results

As we have found the maximum likelihood estimates and their respective confidence intervals, the next step is to show some of these results. However, we first present the Ideal Observer (IO) solutions for different values of the parameters, which depends on the expected losses. We first look at the situations where $\gamma = \kappa = 1$, meaning that the prior distribution for Θ is uniform. Secondly, we have a look at how sensitive the results are to the hyperparameters γ and κ in the beta prior we have for Θ.

4.1 Uniform Prior for Θ

When we present the results, we start with the results where we have a uniform prior for Θ. Recall that this means that it is equally likely that Θ takes any value between zero and one. We start with having a look at the probabilities that either blue or red are the majority colours.

4.1.1 Conditional Probabilities

We will here have a look at the probability that red is the dominant colour and the probability that blue is the dominant colour, as shown in (3.42) and (3.45), respectively. These probabilities can be represented for each possible combination of red and blue boxes. Thus, we can find those probabilities for all the trials the participants have done.

We present the probabilities in Figure 4.1 as a tree. The top node represents the situation where no boxes are opened. The probability that blue is the majority colour is then equal to the probability of red being the majority

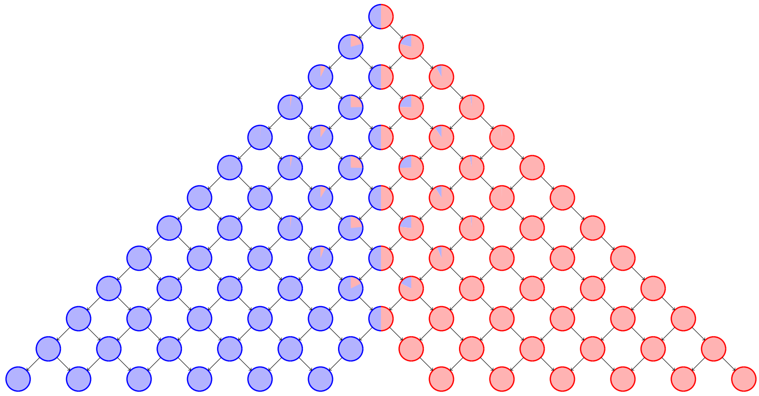

Figure 4.1: A tree representing the probabilities that either red or blue are in the majority in the box task with a uniform prior. The top node is the situation where no boxes are opened, where we do not have any information. Hence, the probabilities are equal. The fraction that is blue inside the node is the probability that a box is blue. The circle around the nodes represents which colour that is most likely to be in the majority. Hence the circle is split between red and blue in the top node. The node down to the left is the situation when a blue box is opened, the node down to the right is when one red box is opened, and so forth.

colour. This is represented as the proportion of blue and red inside the nodes, which in the top node are equal quantities. The circle around the node represent which of the colours that have the highest probability of being in the majority. In the top node, the probabilities are equal, and therefore, this circle is split in two. The node down to the left of the top node represents the situation where one box is opened, and that box is blue. One node down to the left of that one represents the situation where two boxes are opened, and both are blue, and so on. Similarly, the node down to the right of the top node represents the situation where one box that is red is opened and so forth down the tree. We see that in the last row of the tree, there is no middle node. This is because the last row represents the situations where twelve boxes are opened, and as there cannot be six of each colour, that node is not included in the tree. In the row above, we see that all the nodes are completely red or completely blue. That means that we can be sure what the majority colour is after eleven boxes are opened. This is because there cannot be six boxes of each of the colours. Then, if there are six of one colour and five of the other, we know that the colour with six boxes is the majority colour.

When we have these probabilities, the next step is to look at the Ideal Observer solution we have found.

4.1.2 An Ideal Observer Solution in the Unlimited Case

As we have all the expected losses for each possible combination of opened boxes, we can present these similarly to the probabilities in Figure 4.1. The expected losses in the unlimited case are given in (3.16), (3.17) and (3.21). In the unlimited case, we get different solutions for different values of α. We have looked at numerous solutions, but only a handful of them will be presented here.

Recall that an Ideal Observer (IO) solution is the solution we get when the decisions tied to the least expected loss are chosen each time a box is opened. In Figure 4.2, we have visualised the expected losses and the decisions an Ideal Observer would do if $\alpha = 0.0001$. As in Figure 4.1, the top node represents the situation where no boxes are opened, and the node down to the left represents the situation where one box is opened and that box is blue and so forth. The circles around the nodes represent the decision with the least expected loss, thus the decision that an Ideal Observer would make. A blue circle indicates that choosing blue as the majority colour has the least expected loss, and a red circle indicates that choosing red has. A green circle means that the decision to open the next box has the least expected loss. The colours inside the nodes represent what we could call the inverse of the expected losses. That means that if the decision of choosing red as the majority colour has a low expected loss, the amount of red inside that node is big. That means that the colour that represents the Ideal Observer solution is the most prominent in the different nodes. The inverse expected losses are found by adding together all the expected losses and then subtracting the expected loss in question. Later we are normalising these inverse expected losses as they do not sum to one. Let $\tau_i(\delta_i)$ be the inverse expected losses when i boxes are opened. Then, the inverse expected loss for choosing blue as the majority colour is

$$\tau_i(0) = \sum_{j=0}^{2} \mathcal{E}^i_{\delta_j}(\alpha) - \mathcal{E}^i_0(\alpha).$$

Similarly, for red it would be

$$\tau_i(1) = \sum_{j=0}^{2} \mathcal{E}^i_{\delta_j}(\alpha) - \mathcal{E}^i_1(\alpha),$$

and for opening the next box it will be

$$\tau_i(2) = \sum_{j=0}^{2} \mathcal{E}^i_{\delta_j}(\alpha) - \mathcal{E}^i_2(\alpha).$$

We need to normalise these. The proportion of blue in each node would

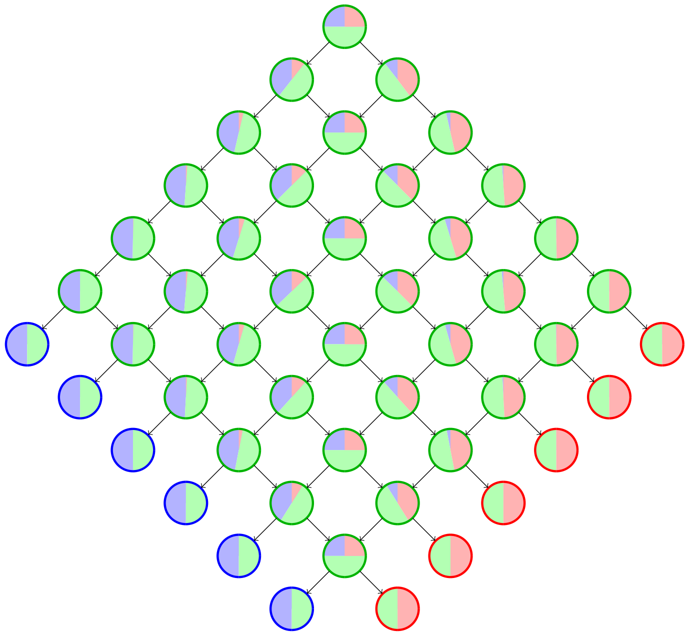

Figure 4.2: A decision tree for the unlimited version of the box task with $\alpha = 0.0001$ and $\gamma = \kappa = 1$. Green circles around each node show that opening another box has the least expected loss, and blue and red circles show that the least expected loss is for choosing the majority colour to be blue and red, respectively. The colours inside the nodes represent the inverse expected losses. As in Figure 4.1, the top node is the situation before any boxes are opened, and the node down to the left of the top node is when one blue box is opened and so on.

then be

$$\frac{\tau_i(0)}{\tau_i(0) + \tau_i(1) + \tau_i(2)}.$$

This tree stops when there is a red or blue circle around the node. That is when the Ideal Observer would choose what the majority colour is, and since this is the Ideal Observer solution, we stop the decision tree there.

As we can see in Figures 4.2, 4.3 and 4.4, where the α values are 0.0001, 0.01 and 0.05, respectively, the trees are slimmer for bigger values of α. Recall that α is the penalty of opening a box and that the expected loss for choosing to open another box increases with it. Hence the threshold for when that expected loss surpasses the expected loss for choosing either blue or red as the majority colour decreases with increasing α. Therefore, we decide at an earlier point when α is high, making the trees slimmer.

At some point, α could get so big that the Ideal Observer would decide

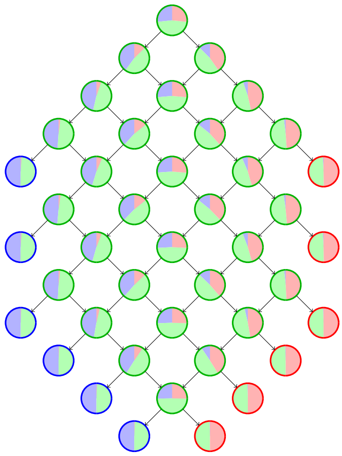

Figure 4.3: A decision tree for an unlimited trial with $\alpha = 0.01$ and $\gamma = \kappa = 1$. We can interpret this tree in the same way as Figure 4.2 where the circles around each node shows which decision the Ideal Observer would do.

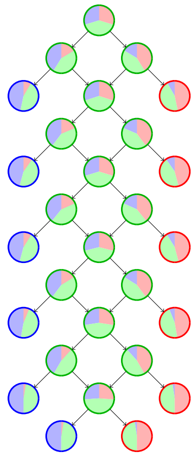

Figure 4.4: A decision tree in the unlimited case with $\alpha = 0.05$ and $\gamma = \kappa = 1$ that can be understood in the same way as Figure 4.2 where the circles around each node shows which decision the Ideal Observer would do.

what the majority colour is after only one box is opened. The expected loss for opening another box is dependent on both α and the expected losses for continuing to open boxes. In contrast, the expected losses for choosing the majority colour depend only on the probabilities that one of the colours is in the majority. For example, if the first box is red, the probability that red is in the majority increases. Therefore, if α is big enough, the expected loss for opening another box could be higher than the probabilities that one of the colours is in the majority; thus, the Ideal Observer would decide after one box is opened. This is the situation when α is 0.1 as in Figure 4.5. The expected loss before any boxes are opened is 0.5, both for choosing blue and red as the majority colour. For the choice of opening a box, the expected loss is 0.308. This is then the choice an Ideal Observer would make before any boxes are opened. If the box that opens is blue, the expected loss for choosing that blue is the majority colour is 0.208, 0.792 for choosing red, and 0.260 for opening another box. Hence, the Ideal Observer would decide that blue is the majority colour. This problem is symmetric. If the opened box is red, the expected loss for choosing that red is the majority colour is 0.208, 0.792 for choosing blue and 0.260 for opening another box. In that case, the expected loss is smallest when we choose red as the majority colour.

As α is the loss we get when we open a box, we can imagine that it also could be zero. In Figure 4.6 we see the decision tree for $\alpha = 0$. When we have six

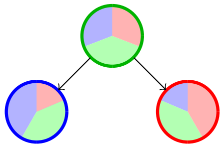

Figure 4.5: A decision tree for an unlimited trial with $\alpha = 0.1$, that can be interpreted in the same way as Figure 4.2

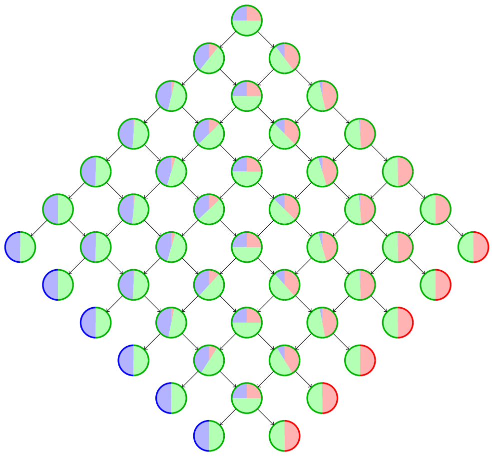

Figure 4.6: An Ideal Observer solution of the unlimited version of the box task with $\alpha = 0$ and $\gamma = \kappa = 1$. This tree can be interpreted the same way as the tree in Figure 4.2. Here we choose that the IO chooses the majority colour if the expected loss of opening the next box is the same as for choosing the majority colour. These are both zero if there are six boxes of one of the colour that is displayed, which is the situation in all the nodes that have circles that are split between two colours.

boxes of one of the colours, for example, red, the expected loss of choosing red as the majority colour is zero, but so is the expected loss of opening the next box. Thus, an Ideal Observer would choose arbitrarily between those. However, we have decided here that the IO would rather choose the majority colour than open the next box if both of these expected losses are zero, such that one does not open any more boxes than necessary. We see that this tree resembles the tree where $\alpha = 0.0001$ as shown in Figure 4.2. This indicated that such a small value of α is very close to having α equal to zero.

In Figure 4.7 we see the Ideal Observer solution in Trial 2 for an individual with $\alpha = 0.01$. Trial 2 is as shown in Figure 3.1. We see that the Ideal Observer would choose after six boxes are opened, hence before we can be completely sure that red is the majority colour, but the probability is 0.9958, so it is very likely that we choose the right colour here. The expected loss of

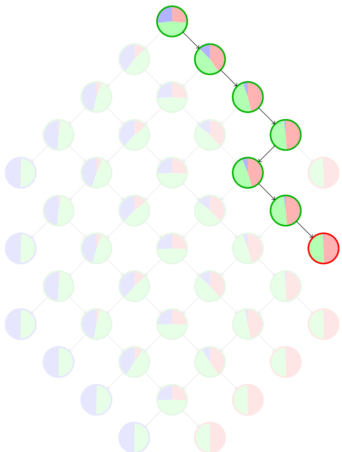

Figure 4.7: An Ideal Observer solution for Trial 2 with $\alpha = 0.01$ and $\gamma = \kappa = 1$. The tree can be interpreted in the same way as Figure 4.2.

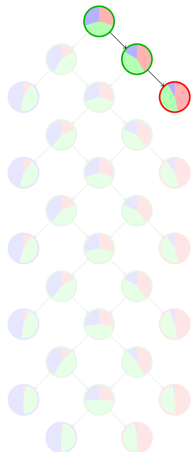

Figure 4.8: An Ideal Observer solution for Trial 2 with $\alpha = 0.05$ and $\gamma = \kappa = 1$. The tree can be interpreted in the same way as Figure 4.2.

choosing red is then 0.0042, whereas the expected loss of opening the next box is 0.0134. We see that if a participant has an $\alpha = 0.05$, then an Ideal Observer would choose after two boxes are opened in Trial 2, as shown in Figure 4.8. Then, the penalty for opening the next box is much higher than in Figure 4.7. Thus, the IO would choose after only two boxes are opened.

As we have presented Ideal Observer solutions in the unlimited case for different values of α, the next step is to show some of the solutions in the limited case.

4.1.3 An Ideal Observer Solution in the Limited Case

Having the expected losses in the limited case as given in (3.16), (3.17) and (3.27), we can visualise them in the same way as Figure 4.2. In the limited trials, we have two parameters, α and β. Thus, we have solutions with both of these parameters varying.

In Figures 4.9 and 4.10, we see two solutions, both with $\alpha = 0.01$. They have different values of β, namely $\beta = 0.6$ and 0.4, respectively. Both trees have the same width, but we see that the tree with the higher β value is shorter. Thus, an Ideal Observer with $\beta = 0.6$ would open fewer boxes than one with $\beta = 0.4$, which is what we would imagine as β is the loss of the test terminating.

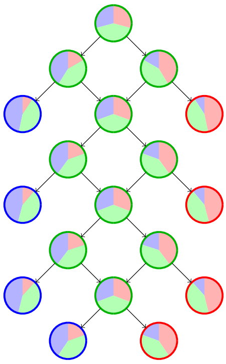

Figure 4.9: A decision tree for an unlimited trial with $\alpha = 0.01$, $\beta = 0.6$ and $\gamma = \kappa = 1$. We can interpret this tree in the same way as Figure 4.2 where the circles around each node shows which decision the Ideal Observer would do.

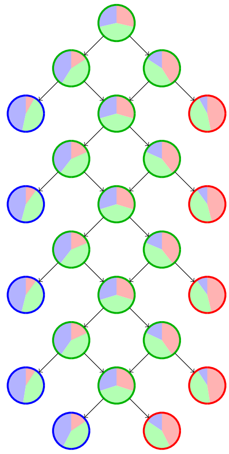

Figure 4.10: A decision tree in the limited case with $\alpha = 0.05$, $\beta = 0.4$ and $\gamma = \kappa = 1$ that can be understood in the same way as Figure 4.2 where the circles around each node shows which decision the Ideal Observer would do.

The trees in the limited case are, in general, slimmer than in the unlimited case. That is because of the penalty we get when the test terminates before we have made a decision. The expected losses for opening another box are bigger than in the unlimited case. Hence, in the limited case, these surpass the expected losses of choosing blue or red as the majority colour earlier than in the unlimited case. Small values of α and β in combination makes the trees wider, as in Figure 4.11.

As in the unlimited case, there are trials where the Ideal Observer solution is to choose the majority colour after one box is opened. This is the case in Figures 4.12 and 4.13. The only thing that has changed from Figure 4.10 to Figure 4.12 and from Figure 4.9 to Figure 4.13 is the value of α, which has increased from 0.01 to 0.05. In the cases with the higher values of α, the expected loss for opening the second box is larger than for choosing the majority colour. That results from these expected losses being dependent on the next expected losses, which again depend on the expected losses for opening another box after that and so on. Additionally, these could potentially be large if the amount of red and blue boxes are close to each other, meaning that we, for example, first open a red box, then a blue, then a red and so forth.

In Figure 4.14, we see the IO solution for Trial 8 where $\alpha = 0.01$ and $\beta = 0.6$. We see that an Ideal Observer would choose the majority colour after seven

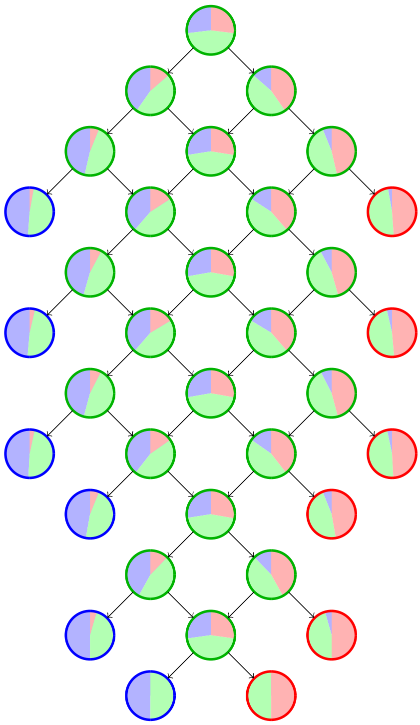

Figure 4.11: A decision tree for an unlimited trial with $\alpha = 0.01$, $\beta = 0.2$ and $\gamma = \kappa = 1$. We can interpret this tree in the same way as Figure 4.2 where the circles around each node shows which decision the Ideal Observer would do.

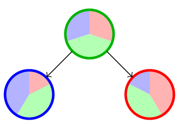

Figure 4.12: A decision tree for a limited trial with $\alpha = 0.05$, $\beta = 0.4$ and $\gamma = \kappa = 1$. It can bee interpreted as the tree in Figure 4.2.

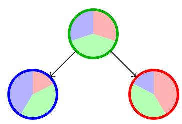

Figure 4.13: A decision tree for a limited trial with $\alpha = 0.05$, $\beta = 0.6$ and $\gamma = \kappa = 1$ that can be interpreted in the same way as the tree in Figure 4.2.

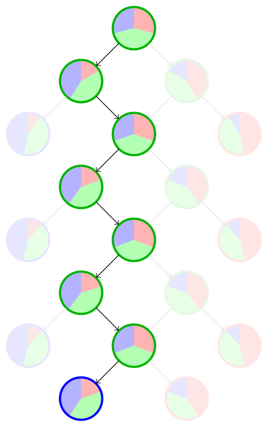

Figure 4.14: An Ideal Observer solution for Trial 8 with $\alpha = 0.01$, $\beta = 0.6$ and $\gamma = \kappa = 1$. The tree can be interpreted in the same way as Figure 4.2

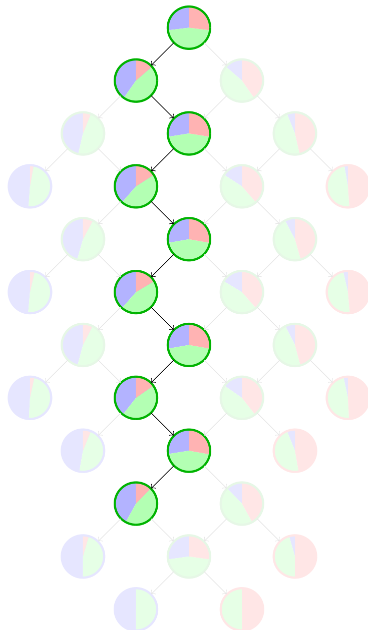

Figure 4.15: An Ideal Observer solution for Trial 8 with $\alpha = 0.01$, $\beta = 0.2$ and $\gamma = \kappa = 1$. The tree can be interpreted in the same way as Figure 4.2

boxes are opened, where four are blue, and three are red. In Figure 4.15, we see another IO solution for Trial 8, where $\alpha = 0.0001$ and $\beta = 0.2$. Here, an Ideal Observer would not choose before the test terminates, and that would be a failed trial. Thus, the Ideal Observer is not perfect, as it is based on expected losses based on the previously opened boxes and not based on what is actually going to happen.

We have now presented some of the Ideal Observer solutions we have found and will continue to present some of the results of the decision model we have defined.

4.1.4 Maximum Likelihood Estimates

As we have presented some Ideal Observer solutions, we will now look at the parameter estimates we have found for each participant. Recall that we find the maximum likelihood estimates (MLEs) by minimising the negative log-likelihood using the L-BFGS-B algorithm as described in Chapter 3.5.

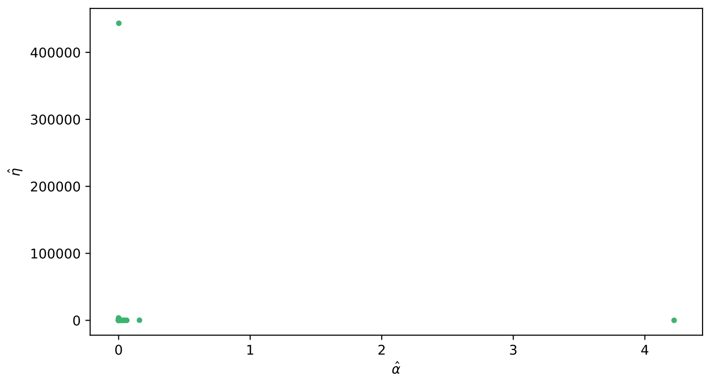

Figure 4.16: The MLEs for all the participants plotted for the unlimited case of the box task with $\gamma = \kappa = 1$. $\hat{\alpha}$ is on the horizontal axis and $\hat{\eta}$ on the vertical.

Unlimited

We start with the unlimited case, where we have the parameters α and η. For each participant, we have found the maximum likelihood estimates of both of these parameters, denoted $\hat{\alpha}$ and $\hat{\eta}$. These are plotted in Figure 4.16. We see one extreme value of each of the estimates, $\hat{\alpha}$ and $\hat{\eta}$. To get a better picture of the values that are not extreme, we zoom in closer to zero for both parameters. This is done in Figure 4.17, and we have zoomed even more in Figure 4.18. Many of the participants have $\hat{\alpha}$ equal to or close to zero, meaning that they have none or little loss of opening boxes. This is not surprising as the task is neither long nor hard to complete, and we would imagine that many of the participants open many boxes.

Recall that high values of η indicate that the participant tends to make the decisions with the least expected losses, and η is infinity when the participant makes the decisions with the least excepted loss each time a box is opened. The log-likelihood function is almost flat for high values of η. That means that the stopping criterion for the optimisation algorithm will be met several places for high values of η. Thus, if the true value of η is infinity, we could find that $\hat{\eta}$ is, for example, 10000 because the stopping criterion is met there. Therefore, we find a threshold for η where we can say that all values above that threshold are so high that they go to infinity. Then, we can say that if a participant has $\hat{\eta}$ above that threshold, she tends to always make the decisions with the least expected loss. This threshold depends on the difference in the two lowest expected losses, but we can, for example, find one for when the difference is 0.01. If we, for example, have a majority of red boxes, then the expected loss of choosing red as the major-

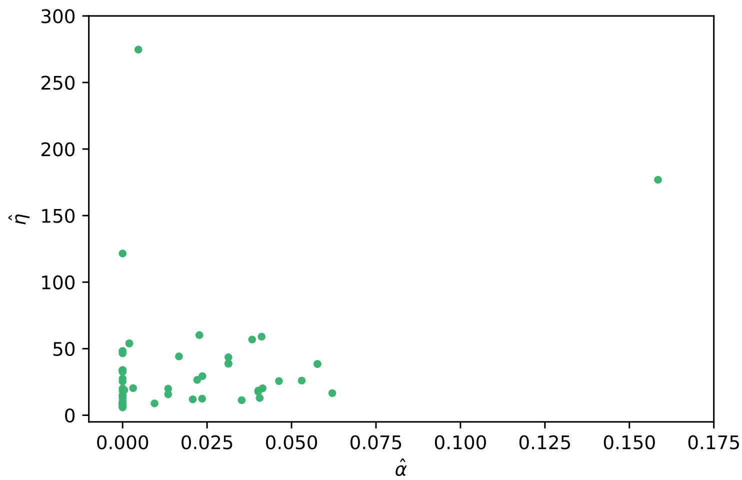

Figure 4.17: The MLEs for all the participants plotted for the unlimited case of the box task with $\gamma = \kappa = 1$. Here we have zoomed in closer to zero on the plot in Figure 4.16.

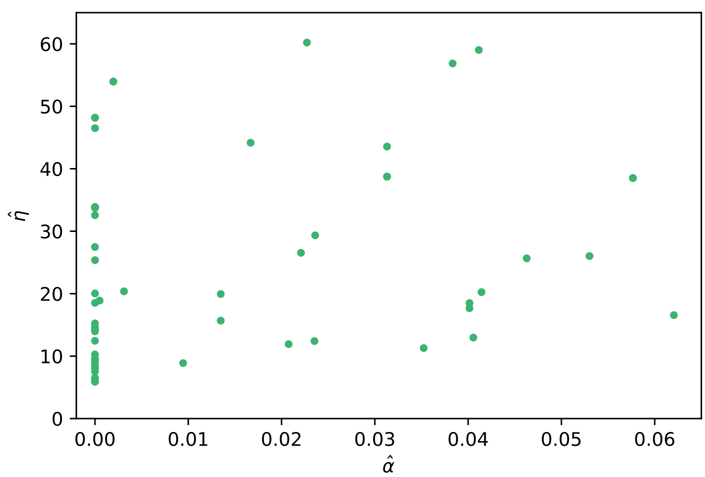

Figure 4.18: The MLEs for all the participants plotted for the unlimited case of the box task with $\gamma = \kappa = 1$. Here we have zoomed in closer to zero even more than in Figure 4.17.

ity colour is quite low, but so could the expected loss of opening the next box be. Consider, a situation where we have $\mathcal{E}_0^i(\varphi) = 0.98$, $\mathcal{E}_1^i(\varphi) = 0.02$ and $\mathcal{E}_2^i(\varphi) = 0.01$. Then, the decision to open the next box has the lowest expected loss, but the decision to choose red as the majority colour is not far away. We then want the threshold value of η to be so high that the probability that the participant chooses to open the next box is close to one. If $\eta = 1000$, the probability that the participant chooses to open the next box given these expected losses, is 0.99995. Thus, we set $\eta = 1000$ as a threshold for when the participant tends to always makes the choices that have the least expected loss.

If we look at the extreme values in Figure 4.16, we find one participant with a very high value of $\hat{\eta}$ and small value of $\hat{\alpha}$. This is individual 58, who has MLEs

$$\begin{aligned} \hat{\alpha} &= 0.0016, \\ \hat{\eta} &= 443422.7. \end{aligned} \tag{4.1}$$

In each of the three unlimited trials, this individual chooses the majority colour exactly when there are six of one of the colours, which is when we can be completely sure what the true majority colour is. That means that individual 58 chooses after seven boxes are opened in Trial 2, and she then chooses red. In Trials 3 and 4, she chooses after 10 and 9 boxes are opened, respectively. Thus, she always chooses the decision with the least expected loss, which are the decisions an Ideal Observer would make, and $\hat{\eta}$ is therefore above the threshold value of 1000. Individual 58 has $\hat{\alpha} = 0.0016$, which is relatively small. That might be because she does not open any more boxes than necessary. She might then have a small loss of opening boxes or some reward of finishing early. This $\hat{\alpha}$ value is so small that it would give an IO solution similar to the one in Figure 4.2, such that one always chooses after six boxes of one of the colours have been opened.

Looking at the other extreme value in Figure 4.16, which is individual number 13, we see that she has a high value of $\hat{\alpha}$ and a small value of $\hat{\eta}$. The values are

$$\begin{aligned} \hat{\alpha} &= 4.2224, \\ \hat{\eta} &= -0.4290. \end{aligned} \tag{4.2}$$

In Trial 2, she chooses after opening two boxes, where both boxes are red. In both Trial 3 and Trial 4, she chooses when three boxes are opened, where two of them are blue, and one is red. She then chooses red as the majority colour despite the fact that choosing blue as the majority colour has a lower expected loss. Thus, she tends to choose decisions with higher expected losses, and therefore has a negative estimate of η. She also chooses quite early; thus, she gets a high estimate of α. However, an Ideal Observer with this high value of α would always choose after one box is opened.

In Figure 4.17 we see another participant that has a high value of $\hat{\alpha}$. This is participant 44 that has MLEs

$$\begin{aligned} \hat{\alpha} &= 0.1585, \\ \hat{\eta} &= 176.9. \end{aligned} \tag{4.3}$$

In all three unlimited trials, she chooses what she thinks is the dominant colour after opening one box. This is why she has a higher value of $\hat{\alpha}$ than most of the other participants. She always chooses the colour of that first box as the majority colour, meaning that if the first box is red, she chooses red as the majority colour. Then the expected loss of choosing red for this participant is 0.2083, whereas the expected loss of opening the next box is 0.3373. Thus, she also chooses the alternative with the least expected loss and gets a pretty high value of $\hat{\eta}$.

If we look at a more typical person, we find, for example, individual number 61. She has

$$\begin{aligned} \hat{\alpha} &= 0.0135, \\ \hat{\eta} &= 19.9432. \end{aligned} \tag{4.4}$$

In the three unlimited trials, she chooses what she thinks is the majority colour after five, six and four boxes are opened. She chooses the majority colour before she can be entirely sure and therefore has $\hat{\alpha}$ larger than zero. However, when she chooses majority colour, she chooses the colour with the least expected loss and thus has a positive $\hat{\eta}$. An IO with the estimates of individual 61 would choose after 3, 10 and 9 boxes were opened. Thus her decisions do not coincide with the IO decisions, but they are not far away.

Having looked at the MLEs in the unlimited case, we now continue with the limited trials.

Limited

In the limited version, we have three parameters, α, β and η, and we have found maximum likelihood estimates for these for all of the 76 participants. The MLEs of α and η for all participants are plotted in Figure 4.20, and we have zoomed in on that plot in Figure 4.21. We see that many of the participants have $\hat{\alpha} = 0$. In fact, all participants except four have $\hat{\alpha} = 0$. We also see that four of the participants have $\hat{\eta}$ higher than the threshold value of 1000, meaning that they make good choices each time a box is opened. What is not so easy to see is that the two participants with the highest values of $\hat{\alpha}$ have negative values of $\hat{\eta}$. This indicates that they make choices with high expected losses and that they have high costs of opening new boxes.

In Figure 4.22 we have plotted the MLEs of α and β for all participants.

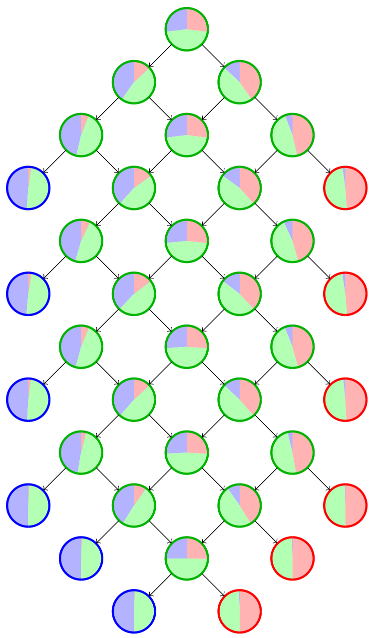

Figure 4.19: An Ideal Observer solution of the box task in the unlimited case for individual number 61 with a uniform prior, such that $\gamma = \kappa = 1$. She has $\hat{\alpha} = 0.0135$. This tree can be interpreted in the same way as the tree in Figure 4.2.

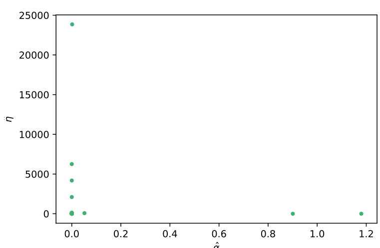

Figure 4.20: Maximum likelihood estimates of α and η for all participants in the limited version of the box task. $\gamma = \kappa = 1$.

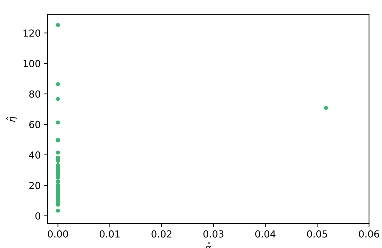

Figure 4.21: MLEs of α and η in the limited case. Zoomed in on the plot in Figure 4.20.

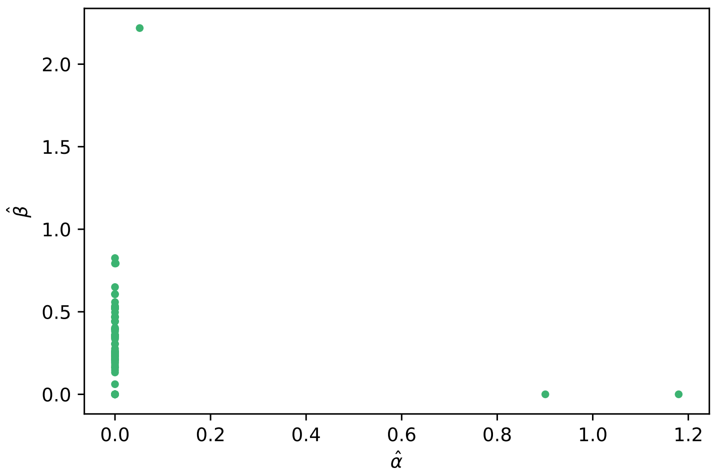

Figure 4.22: Maximum likelihood estimates of α and β for all participants in the limited version of the box task. $\gamma = \kappa = 1$.

Again we see that many participants have $\hat{\alpha} = 0$. This might indicate that α is unnecessary to include in the limited version. Both α and β are values tied to whether we choose the majority colour early or not. Thus, it might be enough only to include β. It is also not obvious how the MLEs give weight to α compared to β.

We have plotted the MLEs of β and η together in Figure 4.23. Again we see that four of the participants have high values of $\hat{\eta}$. Three of these have values of $\hat{\beta}$ close to one, which might indicate that they are afraid of the test terminating and thus choose early, but that they choose the colour that is most likely to be in the majority. Zooming in on the plot as in Figure 4.24, we see that the majority of the participants have values of $\hat{\eta}$ between 10 and 80 and values of $\hat{\beta}$ between zero and 0.7. We also see that some participants have $\hat{\beta}$ equal to zero.

In Figures 4.20 and 4.23 we see a participant with a very high value of $\hat{\eta}$. This is individual number 70, and she has parameter estimates

$$\begin{aligned} \hat{\alpha} &= 0.0015, \\ \hat{\beta} &= 0.7929, \\ \hat{\eta} &= 23851.9. \end{aligned} \tag{4.5}$$

She chooses what she thinks is the majority colour after either two or three boxes are opened in the limited trials. When there are two boxes of the

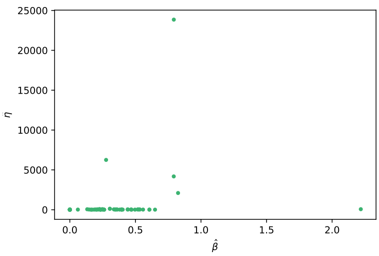

Figure 4.23: Maximum likelihood estimates of β and η for all participants in the limited version of the box task. $\gamma = \kappa = 1$.

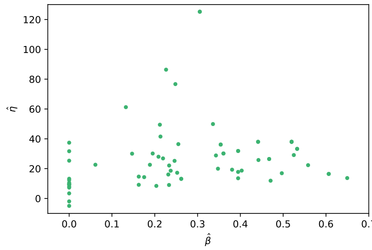

Figure 4.24: MLEs of β and η in the limited case. Zoomed in on the plot in Figure 4.23.

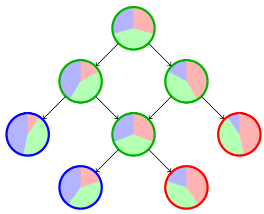

Figure 4.25: An Ideal Observer solution for individual number 70 in the limited version of the box task with $\gamma = \kappa = 1$.

same colour, she chooses that colour as the majority colour. In Figure 4.25 we see the Ideal Observer solution for an individual with the values given in (4.5). We see that the IO would do the exact same thing as individual number 70 has done, that is, choose when we have two boxes of the same colour, and then choose that colour as the dominant colour. In Trial 5, that is after two boxes are opened as the first two boxes are blue. This is visualised in Figure 4.26. In Trial 6, both the IO and participant number 70 chooses after three boxes are opened, as seen in Figure 4.27.

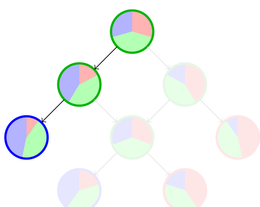

Figure 4.26: An Ideal Observer solution of Trial 5 for individual number 70 where $\gamma = \kappa = 1$.

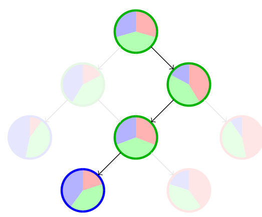

Figure 4.27: An Ideal Observer solution of Trial 6 for individual number 70 where $\gamma = \kappa = 1$.

Figure 4.28: An Ideal Observer solution for individual number 11 in the limited version of the box task with $\gamma = \kappa = 1$.

We have two individuals with high values of $\hat{\alpha}$. These are participants 11 and 13. They both have $\hat{\beta} = 0$ and negative values of $\hat{\eta}$. These two individuals can be interpreted the same way. Therefore, we present only the participant with the highest $\hat{\alpha}$ here, participant number 11. Her parameter estimates are

$$
\begin{aligned}
\hat{\alpha} &= 1.1798, \\
\hat{\beta} &= 0.0, \\
\hat{\eta} &= -1.9538.
\end{aligned} \tag{4.6}
$$

She chooses after two or three boxes are opened, and she tends to choose the colour in the minority, not majority, out of the opened boxes. That is the reason for $\hat{\eta}$ being negative. $\hat{\beta}$ is a measure of the loss one gets when the test terminates. If the test terminates, this counts as a failed trial. It also counts as a failed trial if the participant chooses the wrong colour as the majority colour. This individual does not seem to care if she chooses the wrong colour as the majority colour, and we might believe that, in the same way, she does not care whether the test terminates or not. That might be the reason that $\hat{\beta}$ is zero. However, as she chooses after two and three boxes are opened, this indicates some loss of opening boxes; thus, we get the high value of $\hat{\alpha}$. Earlier, we discussed that α might be unnecessary in the limited trials as so many participants have $\hat{\alpha} = 0$. However, for individual number 11 and the other participant with the high value of $\hat{\alpha}$, this parameter might be needed to express some penalty of opening boxes as we have that $\hat{\beta} = 0$. In Figure 4.28 we see that an Ideal Observer with the same estimates as individual number 11 would choose the majority colour before any boxes are opened. The expected loss of choosing the majority colour then is 0.5 as it is the probability that the opposite colour is in the majority. The expected loss of opening the first box is much higher due to the high value of $\hat{\alpha}$.

In Figures 4.22 and 4.23 we see participant number 75 that has a high value of $\hat{\beta}$. Her parameter estimates are

$$
\begin{aligned}
\hat{\alpha} &= 0.0517, \\
\hat{\beta} &= 2.219, \\
\hat{\eta} &= 70.87.
\end{aligned} \tag{4.7}
$$

This participant chooses after one box is opened in all six limited trials. This is the reason for the high value of $\hat{\beta}$. At the same time, individual 75 always chooses the colour of that first box as the majority colour, which is the colour with the least expected loss. The expected loss of opening the next box might be lower than the expected loss of choosing that as the majority colour. However, these expected losses are often close to each other, whereas the expected loss of choosing the other colour as the majority is often further away. Thus, she tends to choose decisions with low expected losses, but they might not be the decisions with the lowest expected loss. The value of $\hat{\eta}$ is, therefore, relatively high.

If we look at a more typical person, we find, for example, individual number 40. She has parameter estimates

$$\begin{aligned} \hat{\alpha} &= 0.0, \\ \hat{\beta} &= 0.4414, \\ \hat{\eta} &= 38.00. \end{aligned} \tag{4.8}$$

She opens between 2 and 5 boxes in all the trials except Trial 8. In all of these trials, she chooses the most probable colour as the colour that she believes is the majority. Thus, $\hat{\eta}$ is reasonably high. In Trial 8, she opens nine boxes and tries to open the tenth when the test terminates. In that trial, there are never two boxes of the same colour opened after each other, and the probability that one of the colours is the majority colour is quite low.

As the MLEs are presented for the different parameters, we continue with presenting their respective confidence intervals.

4.1.5 Confidence Intervals

We have presented an Ideal Observer solution of the box task and the maximum likelihood estimates for different participants. Now, we will discuss some of the confidence intervals we have found. Recall that these are found by finding the MLEs of 1000 bootstrap samples for each participant and then finding the 5-th and 95-th percentiles, as discussed in Chapter 3.6.

Unlimited

We start with the confidence intervals in the unlimited case. In Figure 4.29 we see the confidence intervals for α. The whole interval for individual number 13 is not included as it is very long. Recall that the MLE of α for person 13 is very large, and we also have a very large value of the upper limit of the CI for α, $\hat{\alpha}_{1000}^{(95)}$. We see that many of the CIs include zero, meaning that many participants might not have that small loss of opening the next box.

In Figure 4.30 we see all the CIs in the unlimited case for η. We see that

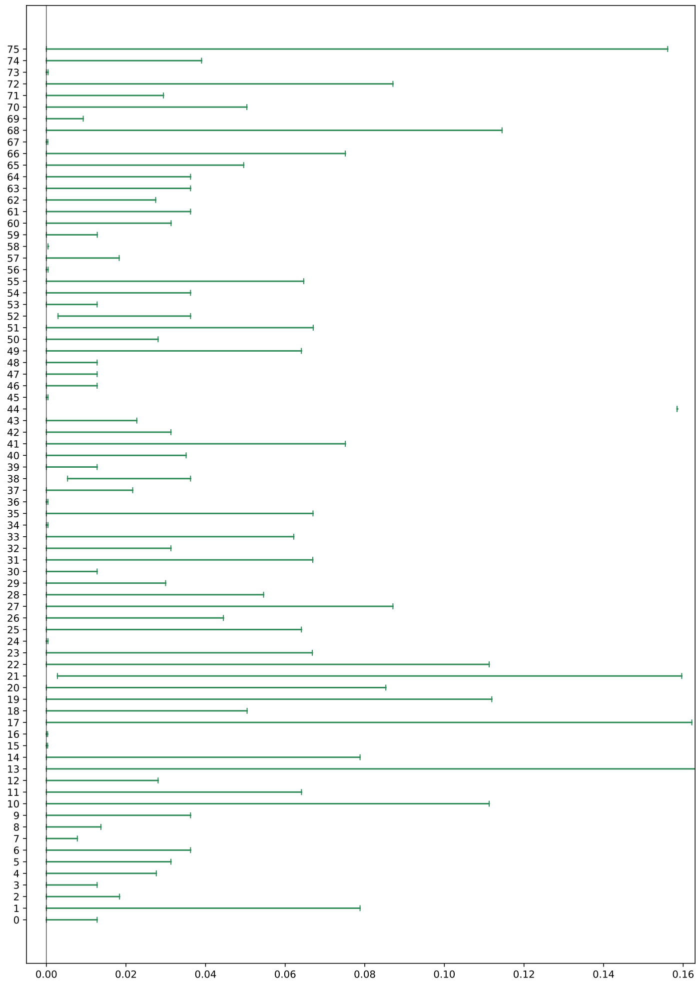

Figure 4.29: Confidence intervals for all the participants for α in the unlimited version of the box task. We see that individual number 13 has an interval outside the range of this plot. $\gamma = \kappa = 1$.

some of the participants have the whole CI above 1000, the threshold value we defined for $\hat{\eta}$. We can for these participants conclude that they always make decisions with small expected losses. We also see that many of the CIs include the threshold value. In Figure 4.31 we have zoomed in on the CIs and added a line at zero. Only one CI include zero, the CI for individual number 13. Recall that she has a negative value of $\hat{\eta}$ and a very high $\hat{\alpha}$, as seen in (4.2).

In Chapter 4.1.4 we discussed individual 61 that is a more typical person. Her MLEs are given in (4.4). We have 1000 bootstrap samples and thus 1000 values of both $\hat{\alpha}$ and $\hat{\eta}$ for this participant. These are plotted in Figure 4.32. There, we have also plotted the confidence interval as black lines. We see that most of the bootstrap samples are accumulated in the bottom left corner. Thus, we zoom in there in Figure 4.33, where we also have plotted the CI. We have that

$$\left[\hat{\alpha}_{1000}^{*(5)}, \hat{\alpha}_{1000}^{*(95)}\right] = [0, 0.0362],$$
$$\left[\hat{\eta}_{1000}^{*(5)}, \hat{\eta}_{1000}^{*(95)}\right] = [10.0449, 110.3661].$$

We see that the confidence interval for α includes zero. Recall that having $\alpha = 0$ means that one does not have any loss of opening boxes. Additionally, we see that the whole interval for η is above zero, which means that individual number 61 makes decisions with low expected losses. However, they might not always be the decisions with the lowest expected loss.

In Chapter 4.1.4 we also discussed individual number 13 that has a high value of $\hat{\alpha}$, as seen in (4.2), compared to the other participants. The 1000 bootstrap samples for participant number 13 are plotted in Figure 4.34, and zoomed on the $\hat{\eta}$ axis in Figure 4.35. The CIs are

$$\left[\hat{\alpha}_{1000}^{*(5)}, \hat{\alpha}_{1000}^{*(95)}\right] = [0, 1246.3510],$$
$$\left[\hat{\eta}_{1000}^{*(5)}, \hat{\eta}_{1000}^{*(95)}\right] = [-9.9272, 18.3986].$$

We see that many of the values of $\hat{\alpha}$ are very high, much higher than the MLE; thus, $\hat{\alpha}_{1000}^{*(95)}$ is very high. That means that the estimate for this individual is uncertain. Thus, the model might not be a good fit for this participant and the choices she makes. That might be because the expected loss is based on the next decisions that will be made. We have decided that these decisions are the decisions that an Ideal Observer would make, the decisions with the least expected losses. However, this participant tends to make decisions with high expected losses, meaning that the decisions we include in the model are not so relevant for this individual. Thus, we might conclude that this model is not a great fit for participants that make decisions with high expected losses.

Figure 4.30: Confidence intervals for each participant for η in the unlimited version of the box task. $\gamma = \kappa = 1$.

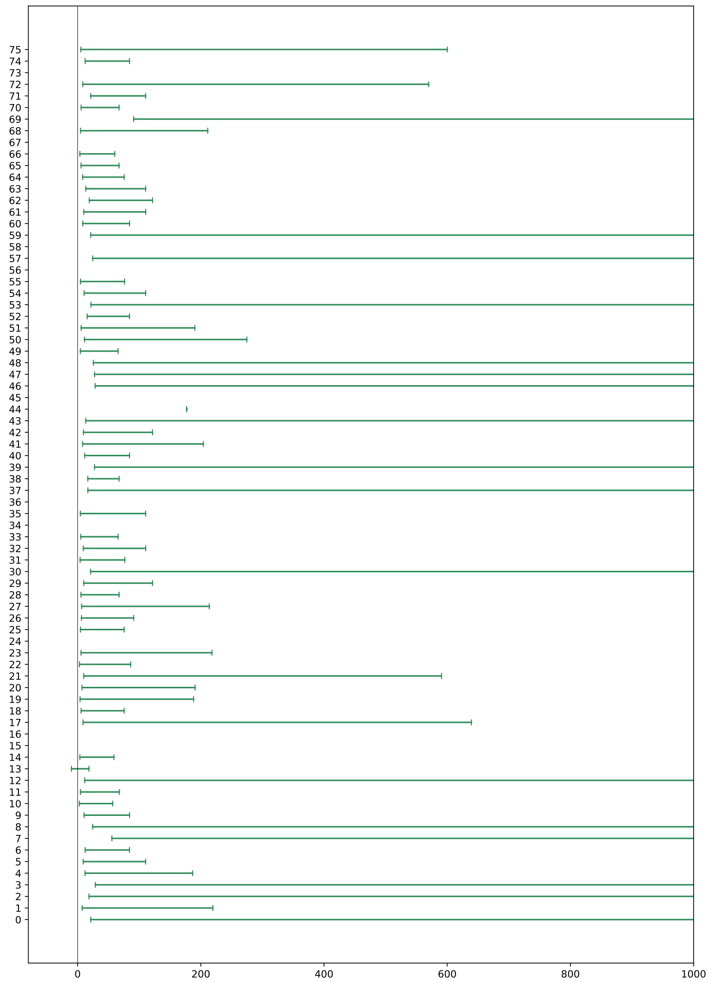

Figure 4.31: Confidence intervals of η for each participant in the unlimited trials of the box task with $\gamma = \kappa = 1$, zoomed.

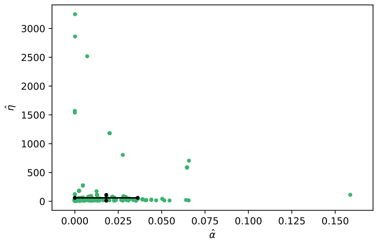

Figure 4.32: All of the MLEs of the 1000 bootstrap samples plotted for individual number 61 in the unlimited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

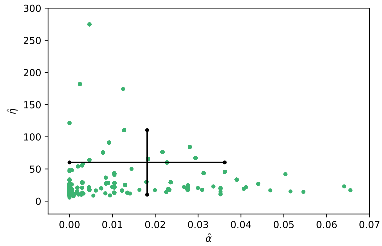

Figure 4.33: Here we have zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.32. This is for individual number 61 in the unlimited case with $\gamma = \kappa = 1$.

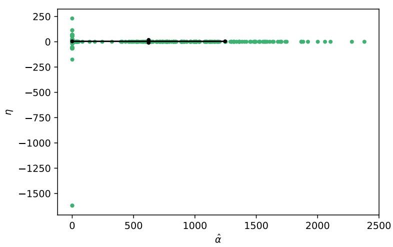

Figure 4.34: All of the MLEs of the 1000 bootstrap samples plotted for individual number 13 in the unlimited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

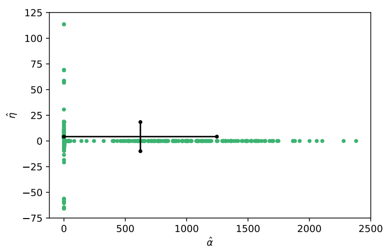

Figure 4.35: Zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.32. This is for individual number 13 in the unlimited case with $\gamma = \kappa = 1$.

Another participant with extreme values we looked at in Chapter 4.1.4 is individual number 58. The MLEs are given in (4.1). Her value of $\hat{\eta}$ is above the threshold value, which makes the probabilities in (3.76) be approximate either zero or one. Thus, when we draw decisions based on those probabilities, we will always end up with the same decisions, and those are the decisions that the participant have made. All the simulated decisions are identical to the decisions the participant has made, and the MLEs will be identical. Thus, the length of the two CIs will be zero, and the values will be equal to the MLEs. Hence, the CIs are

$$\left[\hat{\alpha}^{*(5)}_{1000}, \hat{\alpha}^{*(95)}_{1000}\right] = [0.0016, 0.0016],$$
$$\left[\hat{\eta}^{*(5)}_{1000}, \hat{\eta}^{*(95)}_{1000}\right] = [443422.7, 443422.7].$$

Thus, parametric bootstrapping might not be a satisfactory method to find confidence intervals for the participants that tend to make the decisions with the least expected loss.

Having shown some of the confidence intervals for the unlimited trials, we continue with the limited trials.

Limited

We will now have a look at the confidence intervals in the limited case. In Figure 4.36 we see the CIs for α for all of the participants. Individual number 11 has a much larger upper limit of the CI than the range of the axis here; thus, the CI continues outside the plot. We also see that participant 13 has a higher interval than the other participants. These two are also the individuals with the high values of $\hat{\alpha}$ in Figure 4.20. We also see that all participants except these two have upper limits smaller than 0.2, meaning that most participants have none or a only a small loss of opening the next box.

In Figure 4.37, we have plotted all the confidence intervals for β in the limited case. Again, we see that individual 11 has a large interval that spans beyond the axis. Many of the intervals include zero. Having $\hat{\beta} = 0$ means that one does not feel any loss if the test terminates, although this is defined as a failed trial. We also see that most participants have upper limits below one, meaning that failed trials are not disastrous for most participants.

We have also plotted the confidence interval for η for each participant. This is shown in Figure 4.38. The CI for participant 70 is not included here as it is much higher than the axis range. There are only two CIs that include zero, the ones for participants 11 and 13. The rest of the intervals are above zero. That indicates that most participants make choices with little

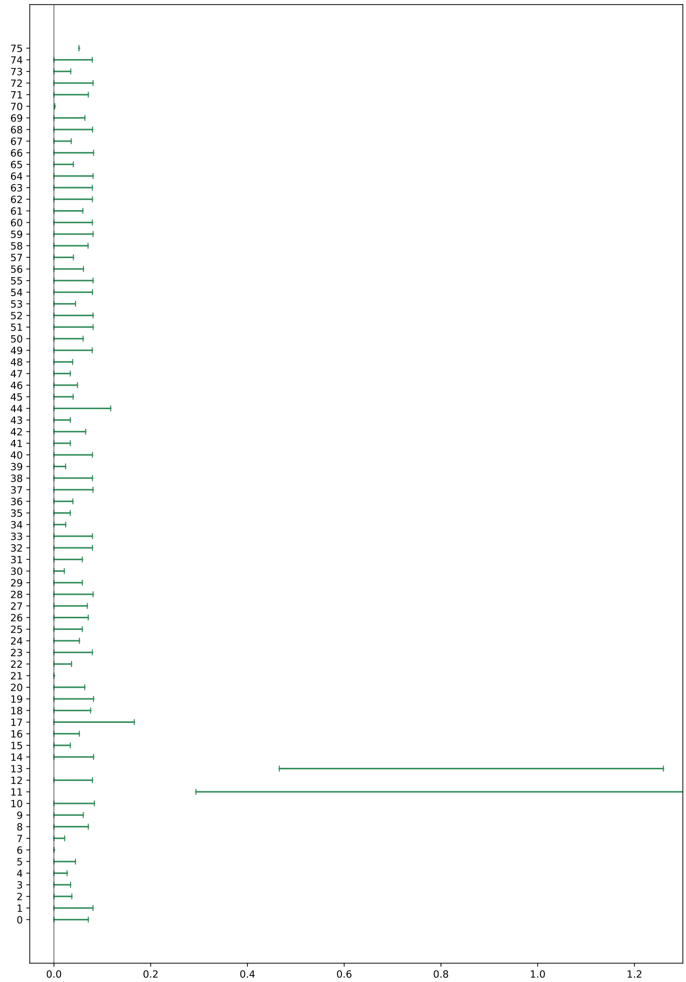

Figure 4.36: Confidence intervals for α in the limited case of the box task with $\gamma = \kappa = 1$ for each participant. Individual number 11 has a much higher value of the upper limit than the range of the axis here.

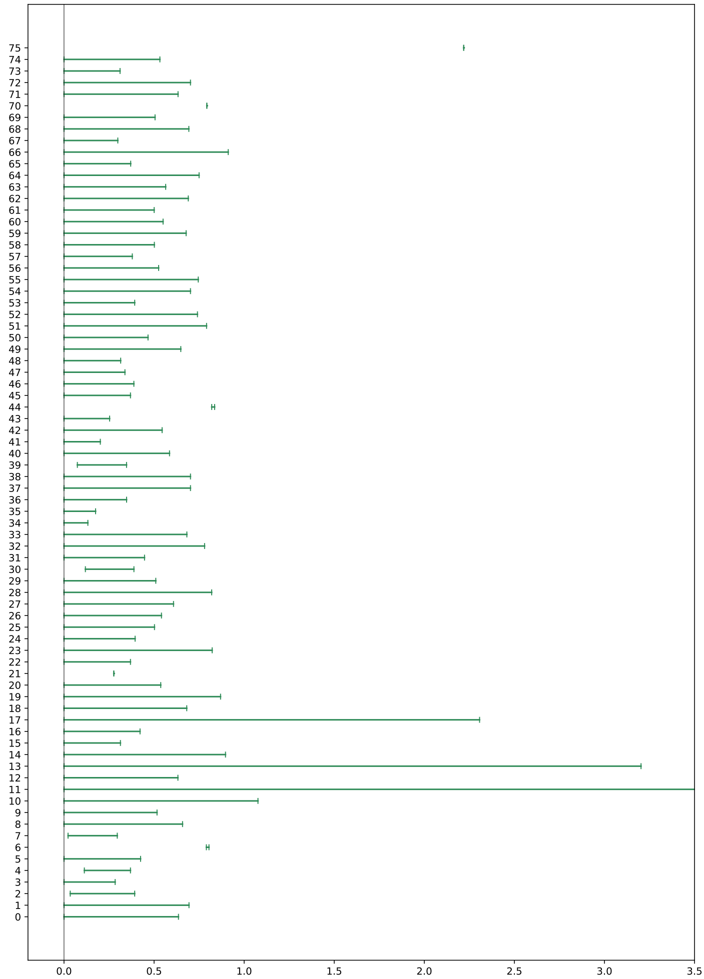

Figure 4.37: Confidence intervals for β in the limited case of the box task with $\gamma = \kappa = 1$ for each participant. As in Figure 4.36 participant 11 has a long interval.

expected loss, whereas individuals 11 and 13 might tend to make decisions with higher expected losses. Many participants have long intervals, which might mean that there is uncertainty tied to their estimates.

We have three participants where the confidence intervals for all three parameters have length zero. These are individuals 21, 70 and 75. As for individual 75 in the unlimited case, they have very high values of $\hat{\eta}$, which makes the probabilities of each of the decisions as seen in (3.76) be approximately zero or approximately one. Thus, all the simulated decisions are the same as the decisions they have made. Therefore, all of the 1000 MLEs are the same, and finding the percentiles gives the same value as the MLEs. It should be noted that individual 75 has $\hat{\eta} = 70.87$ as seen in (4.7), which is far below the threshold value of 1000 that we defined for η. She has a pretty high estimate of β, which makes the expected loss of opening the next box relatively high. Therefore, the expected losses differ considerably, and it takes a much lower value of $\hat{\eta}$ for the probabilities in (3.76) to be approximately zero or one. Thus, $\hat{\eta} = 70.87$ is a high value for individual number 75.

In Figures 4.36 and 4.37 we see that individual 11 has long confidence intervals for α and β. In Figure 4.39 we have plotted $\hat{\alpha}$ and $\hat{\beta}$ of the 1000 bootstrap samples for individual 11 together with the confidence intervals for α and β. We have zoomed in on that plot in Figure 4.40. Additionally, we have plotted the MLEs and the confidence intervals of α and η in Figure 4.41 and zoomed in on that plot in Figure 4.42. The confidence intervals are

$$\begin{aligned}
\left[\hat{\alpha}_{1000}^{*(5)}, \hat{\alpha}_{1000}^{*(95)}\right] &= [0.2935, 649.1], \\
\left[\hat{\beta}_{1000}^{*(5)}, \hat{\beta}_{1000}^{*(95)}\right] &= [0, 16.74], \\
\left[\hat{\eta}_{1000}^{*(5)}, \hat{\eta}_{1000}^{*(95)}\right] &= [-10.45, -0.0013].
\end{aligned}$$

We see also here that there are many high values of $\hat{\alpha}$, which again give long intervals that indicate an uncertain estimate of α. Furthermore, there are some extreme values of $\hat{\beta}$. The CI has an upper limit at 16.74, which is relatively high compared to the other participants. Thus, there is some uncertainty tied to the estimate of both α and β for individual 11, indicating that this model is not a good fit. As discussed for individual 13 in the unlimited version, this might be because we assume Ideal Observer decisions when we find the expected losses of opening the next box. These decisions are not that relevant for participant 11 in the limited case as she tends to make decisions with high expected losses. She chooses the colour in the minority multiple times, meaning that she chooses the decisions with the highest expected loss.

Figure 4.38: Confidence intervals for η in the limited case of the box task with $\gamma = \kappa = 1$ for each participant. Participant 70 has a too high CI to be included here.

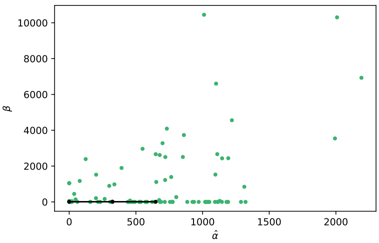

Figure 4.39: All of the MLEs for α and β of the 1000 bootstrap samples plotted for individual number 11 in the limited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

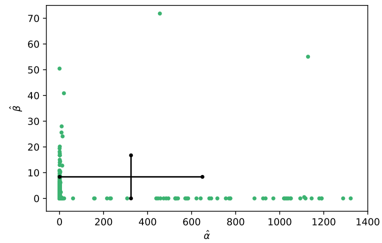

Figure 4.40: Here we have zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.39. This is for individual number 11 in the limited case with $\gamma = \kappa = 1$.

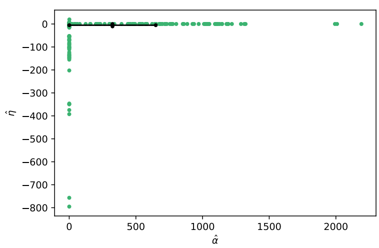

Figure 4.41: All of the MLEs for α and η of the 1000 bootstrap samples plotted for individual number 11 in the limited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

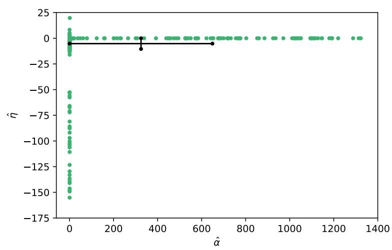

Figure 4.42: Here we have zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.41. This is for individual number 11 in the limited case with $\gamma = \kappa = 1$.

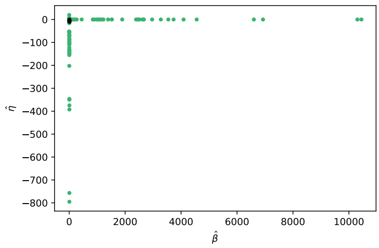

Figure 4.43: All of the MLEs for β and η of the 1000 bootstrap samples plotted for individual number 11 in the limited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

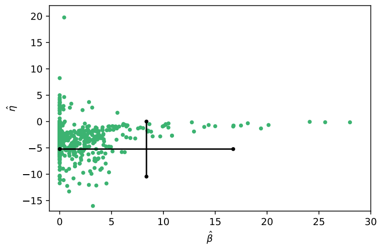

Figure 4.44: Here we have zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.43. This is for individual number 11 in the limited case with $\gamma = \kappa = 1$.

In Figure 4.43 we see $\hat{\beta}$ and $\hat{\eta}$ of participant number 11 for the 1000 bootstrap samples plotted, and zoomed in Figure 4.44. We see that the whole interval for η is below zero. This is a strong indicator that individual 11 makes many choices with high expected losses.

In Figure 4.45, we have plotted the MLEs of α and β for the 1000 bootstrap samples of individual 40 that is a more typical person. We see that $\hat{\alpha}$ and $\hat{\beta}$ tend to not be zero at the same time, indicating that she gets some loss of opening boxes, but it differs whether the MLE is putting weight on α or β. In Figure 4.46 we have plotted all $\hat{\alpha}$ and $\hat{\eta}$ together, and zoomed in on that plot in Figure 4.47. Here we see that there are many high values of $\hat{\eta}$, and that the confidence interval is long. The values of $\hat{\beta}$ and $\hat{\eta}$ are plotted in Figure 4.48 and we have zoomed in Figure 4.49. The confidence intervals are

$$
\begin{aligned}
\left[\hat{\alpha}^{*(5)}_{1000}, \hat{\alpha}^{*(95)}_{1000}\right] &= [0, 0.0796], \\
\left[\hat{\beta}^{*(5)}_{1000}, \hat{\beta}^{*(95)}_{1000}\right] &= [0, 0.5857], \\
\left[\hat{\eta}^{*(5)}_{1000}, \hat{\eta}^{*(95)}_{1000}\right] &= [21.29, 2716.5].
\end{aligned}
$$

We have now looked at the results when we have a uniform prior, meaning that $\gamma = \kappa = 1$. Next, we look at how the results are influenced when we use a different prior. Thus, we look at the sensitivity of to the hyperparameters γ and κ.

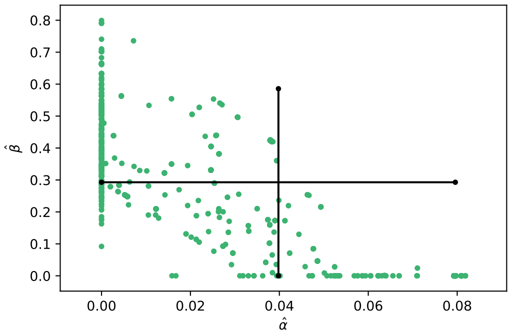

Figure 4.45: All of the MLEs for α and β of the 1000 bootstrap samples plotted for individual number 40 in the limited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

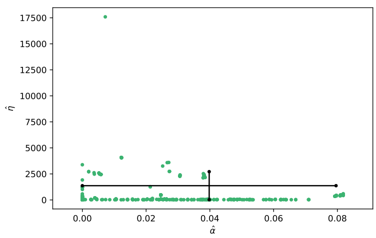

Figure 4.46: All of the MLEs for α and η of the 1000 bootstrap samples plotted for individual number 40 in the limited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

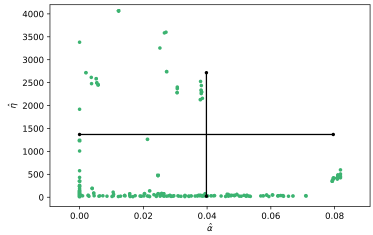

Figure 4.47: Here we have zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.46. This is for individual number 11 in the limited case with $\gamma = \kappa = 1$.

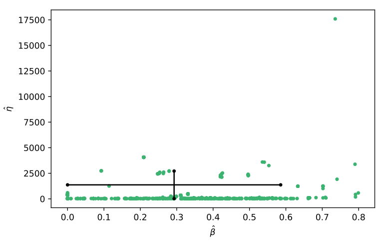

Figure 4.48: All of the MLEs for β and η of the 1000 bootstrap samples plotted for individual number 40 in the limited case with $\gamma = \kappa = 1$. The confidence intervals for the two parameters are also included.

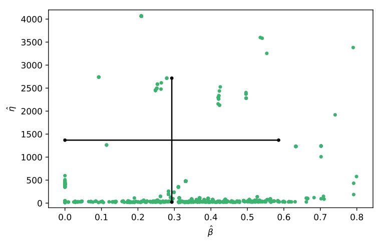

Figure 4.49: Here we have zoomed in on the bootstrap samples and confidence intervals that are plotted in Figure 4.48. This is for individual number 11 in the limited case with $\gamma = \kappa = 1$.

4.2 Sensitivity to Hyperparameters

The results we have discussed so far have all been for a uniform prior of Θ. Recall that Θ is the success probability in the Bernoulli distribution we have for the X_i's, that are the colours of the boxes as described in (2.13), where we later conditioned on there not being six blue and six red boxes. We have a beta prior for Θ with hyperparameters γ and κ, seen in (2.9). As we can see in Figure 3.9, having $\gamma = \kappa = 1$ is the same as having a uniform prior for Θ. We will now have a look at how the results are affected when we use a different prior.

As discussed in Chapter 3.1, the participants are told that one of the colours always is in the majority. Therefore, one might think that the probability that there are more of one of the colours is higher than the probability that there are six of each of the colours. Thus, we want a prior more like the purple or orange lines in Figure 3.9. We choose the purple line where $\gamma = \kappa = 0.5$.

Having a uniform prior for Θ, we have found MLEs for all the parameters and confidence intervals tied to each of them for each of the participants. We do the same with a prior having $\gamma = \kappa = 0.5$, and compare the results. With these values of the hyperparameters, the model put more weight on the colours of the boxes that are opened. For example, suppose many of the opened boxes are blue. In that case, the probability that blue is the majority colour is higher in the model with the new prior, resulting in a lower expected loss of choosing blue as the majority colour.

4.2.1 Unlimited

When we look at the sensitivity to hyperparameters, we start with the results for the unlimited version.

In Figure 4.50 we have plotted all the MLEs for both priors. The green dots are the MLEs in the model with the uniform prior for Θ, and the orange points are the MLEs in the model with the new prior. We see that the participants that had extreme values with the uniform prior still have extreme values with the new prior, but the values differ considerably. The highest value of $\hat{\eta}$ is remarkably lower, but it is still above that threshold value we defined for η. That might be because of the flat log-likelihood function, making the high estimates of η somewhat uncertain, as discussed in Chapter 4.1.4. Thus, the large change in $\hat{\eta}$ might not be because of the new prior but because of the numerical optimisation. If we zoom in on the plot in Figure 4.50, as we have done in Figure 4.51, we see that the participants with high, but not extreme, values of $\hat{\eta}$ or $\hat{\alpha}$ have shifted their

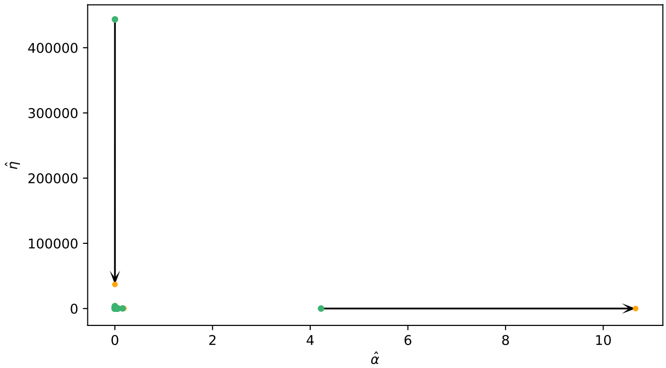

Figure 4.50: Maximum likelihood estimates in the unlimited version of the box task. The green dots represent MLEs where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange dots are the MLEs where $\gamma = \kappa = 0.5$.

MLEs substantially. However, zooming even more, as in Figure 4.52, we see that the MLEs of the participants that do not have extreme or high values are shifted a little, but not considerably, with some exceptions. All participants who did not have $\hat{\alpha} = 0$ before have shifted the values of $\hat{\alpha}$ closer to zero with the new prior. Most of the participants with $\hat{\alpha} = 0$ in the first model have the same in the alternative model. We also see that most of these participants have the same or about the same value of $\hat{\eta}$, perhaps slightly shifted upwards, with the new prior. For the other participants, there are, with some exceptions, no considerable differences in the values of $\hat{\eta}$. Some participants get higher values of $\hat{\eta}$, and others get lower values without any apparent pattern.

In Figure 4.53 we see the confidence intervals for α. The green lines are the CIs we found earlier with the uniform prior, and the orange lines represent the CIs when we have the beta prior for Θ with $\gamma = \kappa = 0.5$. As for the MLEs, we see that the confidence intervals are close to each other. Many of the upper limits are slightly shifted to the left, with some exceptions. Again, we see that the interval of individual 13 is long, such that the upper limit is outside of this plot. These upper limits for the old and new prior are quite similar.

Figure 4.54 displays the confidence interval for η for each participant. We also have here that the orange lines are the CIs in the case with the new prior and the green lines are the CIs we found in the case with the uniform prior. In Figure 4.55 we have zoomed in closer to zero to get a better view of the participants with lower values of $\hat{\eta}$. These confidence intervals

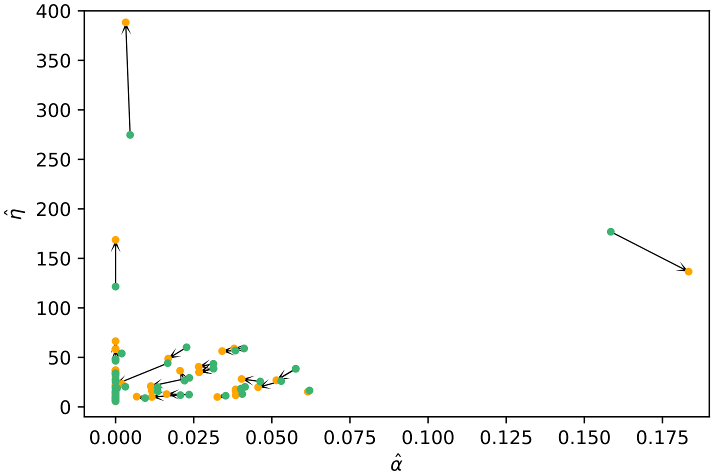

Figure 4.51: MLEs using two different priors in the unlimited case of the box task. This is the same plot as in Figure 4.50 zoomed in.

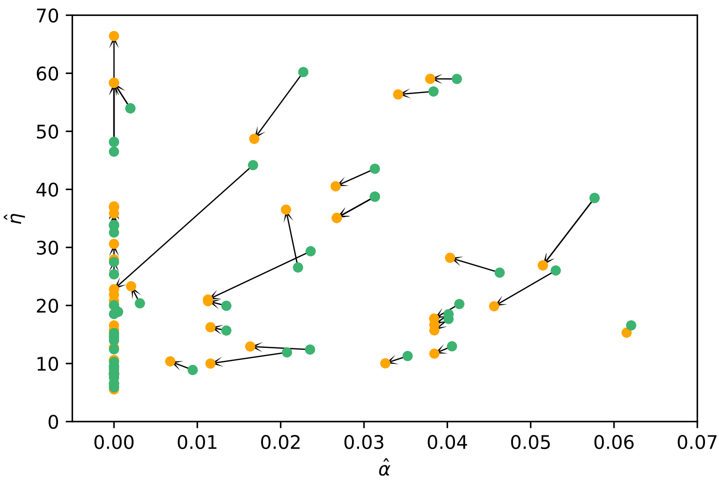

Figure 4.52: The plot of MLEs in Figure 4.51 zoomed.

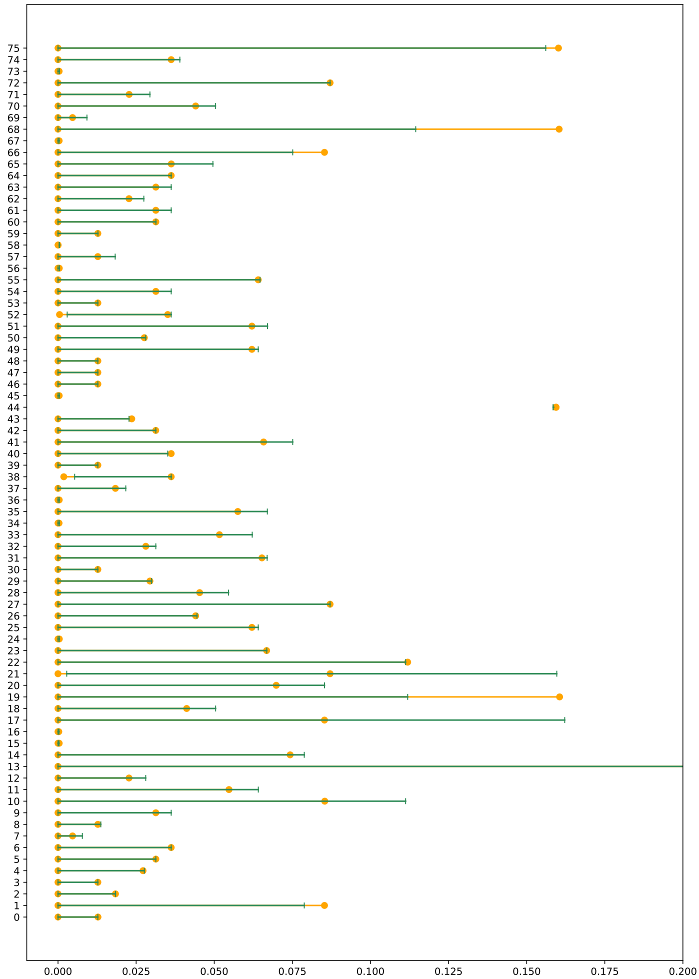

Figure 4.53: Confidence intervals for α for all of the 76 participants in the unlimited version of the box task. The green intervals represent the situation where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange lines are the intervals for a prior where $\gamma = \kappa = 0.5$.

differ more than the intervals for α. Some are shorter than before, and others are longer, without any clear pattern of which intervals that is. Some intervals are slightly shifted to the right, but those are mostly the intervals that are longer than the original intervals. The upper limits are mainly higher than the original, and the new lower limits are often close to the old ones. However, these higher upper limits might be because of the flat log-likelihood function as discussed in Chapter 4.1.4.

Thus, we see that the unlimited version is not very sensitive when changing the hyperparameters from 1 to 0.5. The MLEs of η do not seem to change much, whereas the MLEs of α tend to be a bit smaller. There are no substantial changes in the confidence intervals.

4.2.2 Limited

We continue looking at the estimates and confidence intervals using a new prior for Θ, this time in the limited case of the box task.

In the previous results, where we have a uniform prior, we have that many values of $\hat{\alpha}$ are zero in the limited version. In Figure 4.56 we have plotted both the new and the old estimates of α and η. The green dots are the estimates for the model with the uniform prior, and the orange dots are the new estimates with the beta prior having hyperparameters $\gamma = \kappa = 0.5$. We see that all participants that have $\hat{\alpha} = 0$ with the uniform prior, still have that with the non-uniform prior. The two participants with high values of $\hat{\alpha}$, individuals 11 and 13, get even higher values of $\hat{\alpha}$. Their values of $\hat{\eta}$ are slightly higher with the new prior. There are three participants that had values of $\hat{\alpha}$ slightly higher than zero that with the new prior have $\hat{\alpha} = 0$. Most participants get a lower value of $\hat{\eta}$ with the non-uniform prior. The majority of the values are only shifted a little downwards, except some that get a huge difference in the values. These are participants that have huge values of $\hat{\eta}$ from before with the uniform prior. We discussed earlier that the log-likelihood function is flat for high values of η and that the stopping criterion in the optimisation can be met several places when η is high. That might be the reason for the large change in the value of $\hat{\eta}$. A few participants get higher values of $\hat{\eta}$, but if the values are higher it is not by a lot.

We have plotted the MLEs for α and β together in Figure 4.57. We see also here participants 11 and 13 with the high values of $\hat{\alpha}$. Their values of $\hat{\beta}$ were zero with the uniform prior and stay zero with the new prior. We also see one participant that increases the value of $\hat{\beta}$ considerably. That is individual number 75. Recall that this participant has confidence intervals of length zero because of the high value of $\hat{\eta}$. Her MLEs for the situation where we have a uniform prior are given in (4.7). The log-likelihood function

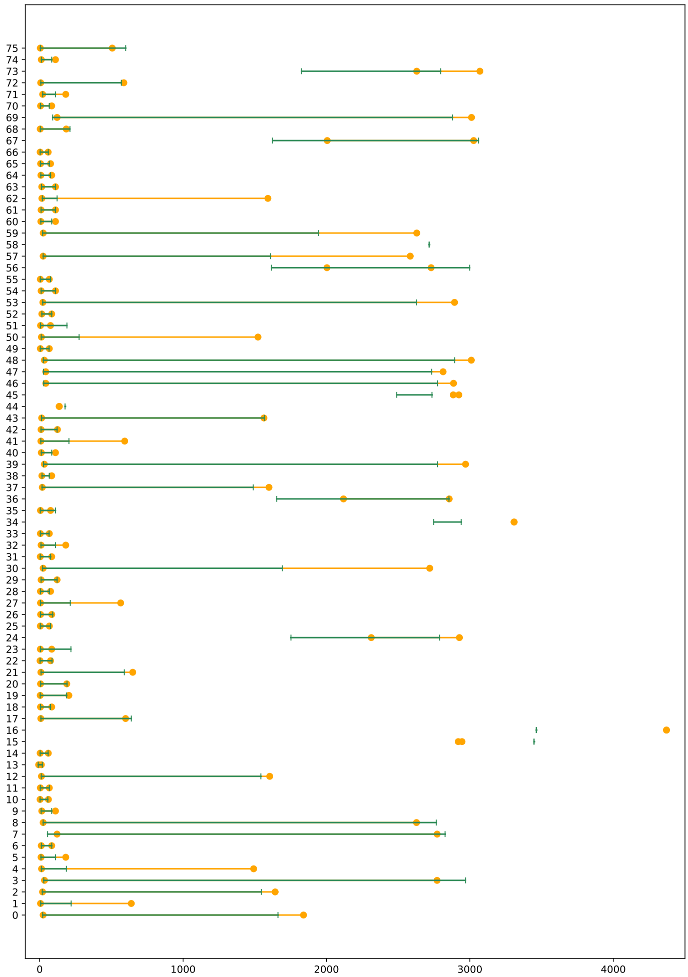

Figure 4.54: Confidence intervals for η for all of the 76 participants in the unlimited version of the box task. The green intervals represent the situation where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange lines are the intervals for a prior where $\gamma = \kappa = 0.5$.

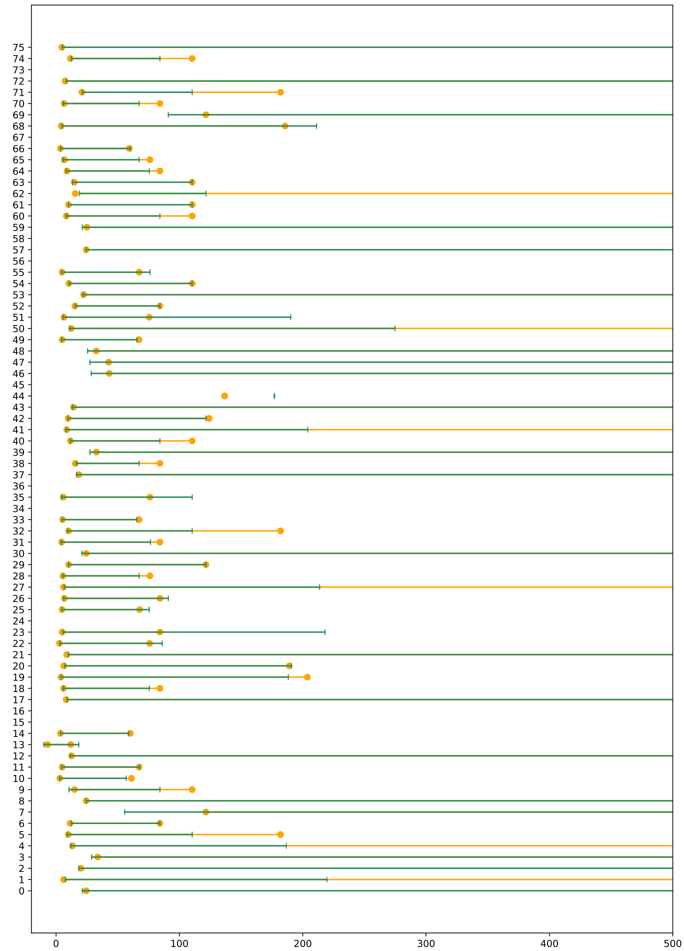

Figure 4.55: The plot in Figure 4.54 zoomed.

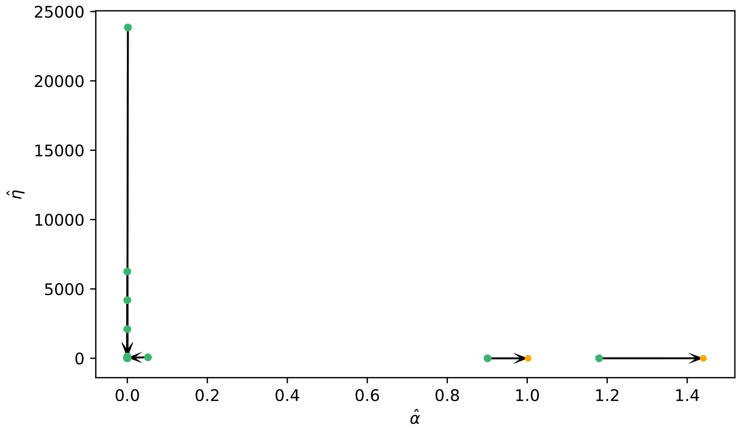

Figure 4.56: Maximum likelihood estimates of α and η in the limited version of the box task. The green dots represent MLEs where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange dots are the MLEs where $\gamma = \kappa = 0.5$.

is also quite flat for big values of α or β, meaning that the same argument for high differences in $\hat{\eta}$ holds for $\hat{\beta}$ as well.

The plot for $\hat{\beta}$ and $\hat{\eta}$ together is shown in Figure 4.58. Here, we again see participant 75 that gets a much higher $\hat{\beta}$. She gets a smaller value of $\hat{\eta}$, but again, that might be because of the flat log-likelihood function.

If we zoom in on the plot in Figure 4.58, as in Figure 4.59, we see that most participants, with some exceptions, get smaller values of both $\hat{\beta}$ and $\hat{\eta}$. All the participants that had $\hat{\beta} = 0$ with the uniform prior still have that with the non-uniform prior. All the other participants, except participant 75, get slightly smaller values of $\hat{\beta}$, resulting in more participants getting $\hat{\beta} = 0$ with the new prior.

In general, we get smaller values of all of the estimates of all three parameters, with some exceptions, with the non-uniform prior for Θ.

In Figure 4.60, we have plotted all the confidence intervals for α in both the case where we have a uniform prior and the case where we have a non-uniform prior. As before, the green lines are the intervals when we have the uniform prior, and the orange lines when the prior is non-uniform. We see that the upper limit of the CI for participant 11 is higher than the upper limit here. A plot including the whole interval for this person can be found in Appendix B. Here we have also included a plot that zooms in on the

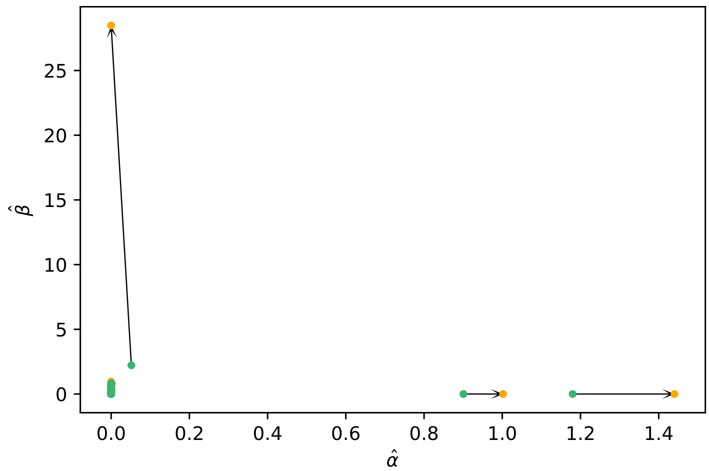

Figure 4.57: Maximum likelihood estimates of α and β in the limited version of the box task. The green dots represent MLEs where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange dots are the MLEs where $\gamma = \kappa = 0.5$.

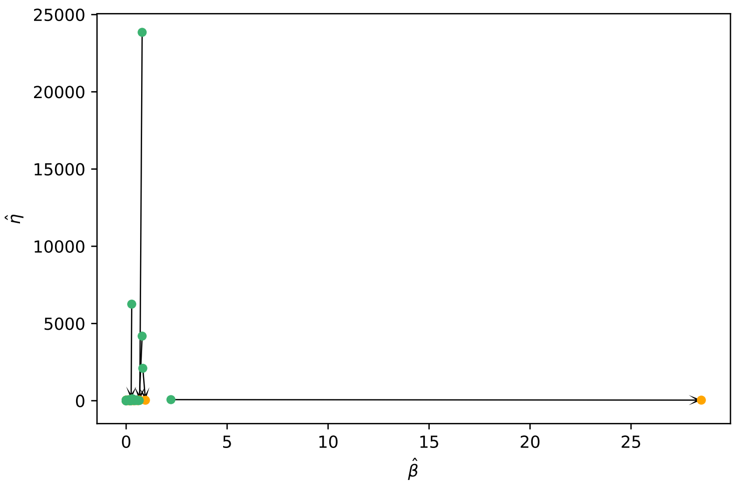

Figure 4.58: Maximum likelihood estimates of β and η in the limited version of the box task. The green dots represent MLEs where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange dots are the MLEs where $\gamma = \kappa = 0.5$.

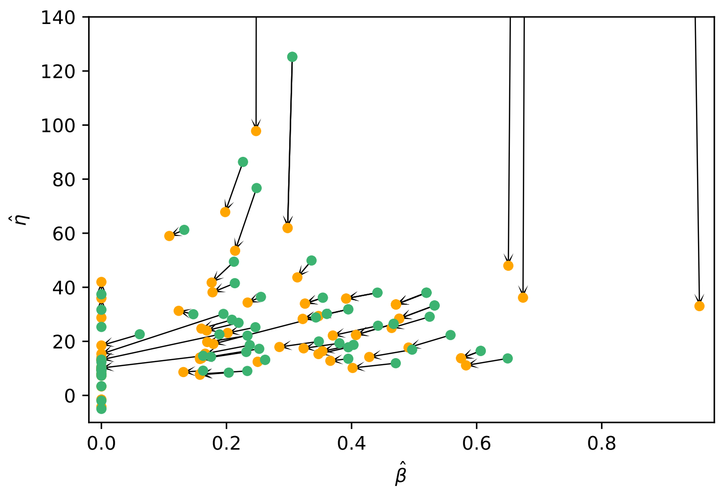

Figure 4.59: The plot in Figure 4.58, zoomed.

x-axis in Figure 4.60. We see that many of the intervals are pretty similar for both priors. However, we see that participants 6, 21 and 70, who had intervals of length zero, actually get an interval with the new prior. That is because their new estimates of η are much lower; thus, the probabilities in (3.76) are not zero or one anymore. In contrast, we have participant 75 that still has length zero of the confidence interval.

We have plotted the confidence intervals for β for all participants with the two different priors in Figure 4.61. As for α, we see that participants 21 and 70 no longer have intervals of length zero, but that individual 75 still has. We also get that participants 6 and 44 get much longer confidence intervals for β with the new prior. The rest of the intervals are quite similar with the two different priors.

The confidence intervals for η are plotted in Figure 4.62. Again, we see that individual 21 and 70 get new intervals with length larger than zero and that individual 75 does not. As for the confidence interval for β for individual 44, the interval of η is much longer with the new prior. Along with this, individual 6 gets non-overlapping confidence intervals for η. The rest of the intervals are close for the two different priors. Some intervals are now shorter, and some are longer.

As we have seen, the parameters estimates in the limited version are more affected by the change of prior than in the unlimited version. In the limited,

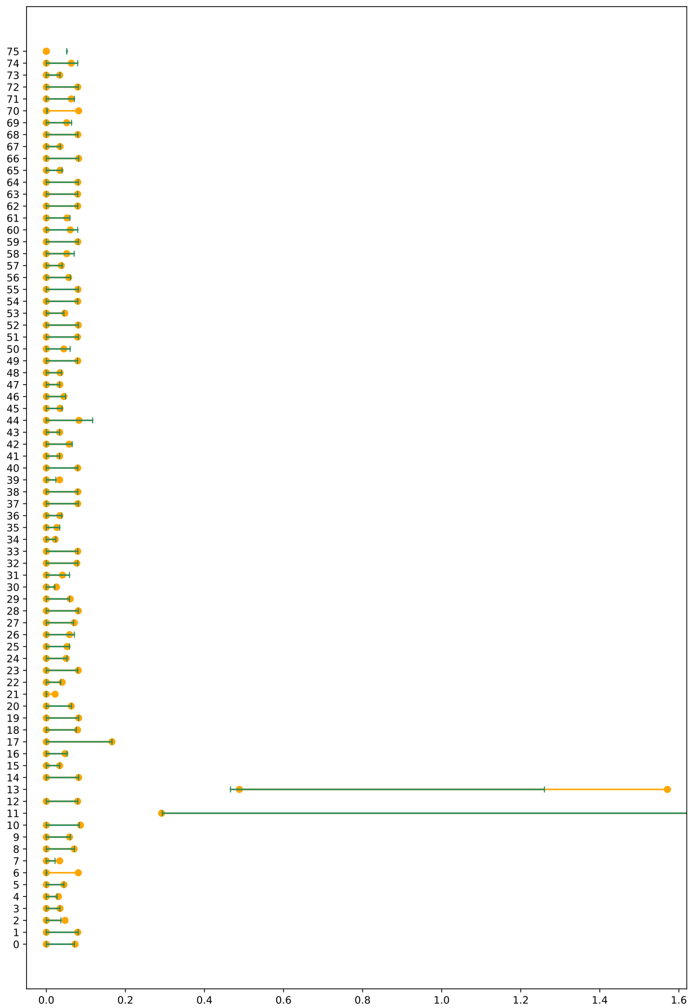

Figure 4.60: Confidence intervals for α for all of the 76 participants in the limited version of the box task. The green intervals represent the situation where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange lines are the intervals for a prior where $\gamma = \kappa = 0.5$.

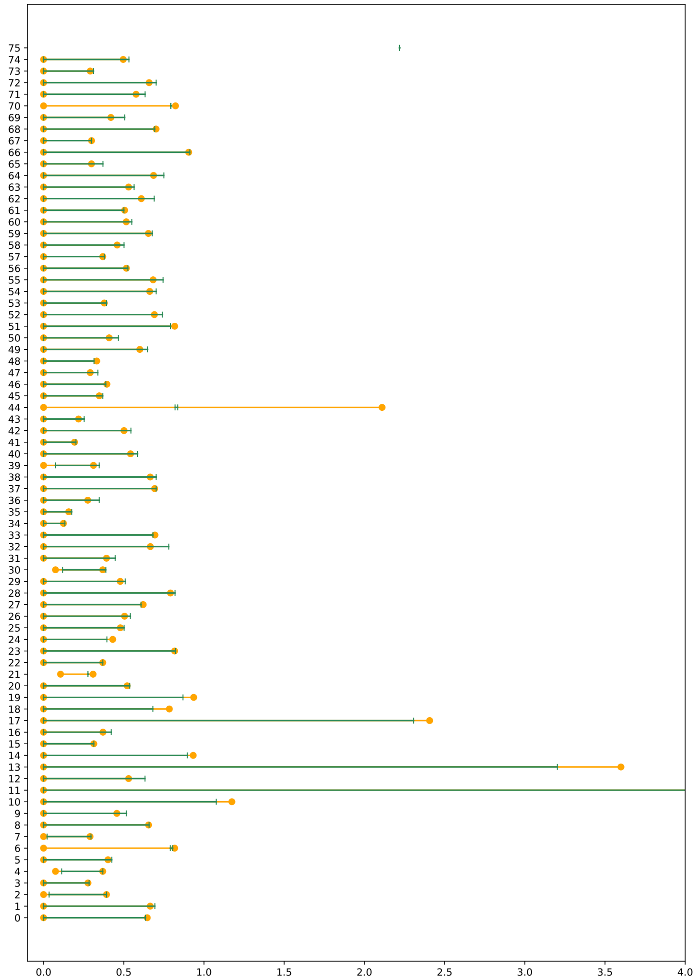

Figure 4.61: Confidence intervals for β for all of the 76 participants in the limited version of the box task. The green intervals represent the situation where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange lines are the intervals for a prior where $\gamma = \kappa = 0.5$.

Figure 4.62: Confidence intervals for η for all of the 76 participants in the limited version of the box task. The green intervals represent the situation where we have a uniform prior for Θ, that is, $\gamma = \kappa = 1$. The orange lines are the intervals for a prior where $\gamma = \kappa = 0.5$. The old interval for person 70 is a scalar that has value 23851.9, thus it is outside the rage of this plot.

all three parameters seem to get slightly smaller values, whereas only $\hat{\alpha}$ and not $\hat{\eta}$ in the unlimited seem to be smaller. However, for both cases, there seem not to be any large changes in the confidence intervals.

Chapter 5

Closing Remarks

In this report, we model decisions in the box task using a softmax model with parameters α, β and η. The parameter α is a small penalty or loss we get each time we open a box, and β is the loss we get when the test terminates before choosing what the majority colour is. The last parameter, η, says something about how far away the choices we make are from the decision with the least expected loss. 76 participants have done several box task trials. We find maximum likelihood estimates for each participant and confidence intervals tied to each parameter using parametric bootstrapping. MLEs are plotted for all participants, in addition to bootstrap samples and CIs for some participants. We find that the model is a good fit for participants who make good choices but worse for those who make decisions with high expected losses. Moreover, parametric bootstrapping makes the confidence intervals for participants that make good choices be zero. Thus, that is not an ideal way of finding CIs for these participants. We also discuss the sensitivity to the hyperparameters in the prior distribution for Θ. The estimates of α seem to be slightly smaller in the unlimited case when we change the hyperparameters from 1 to 0.5. The values of $\hat{\eta}$ changes slightly, but with no pattern of who gets higher and who gets lower values. In the limited version, all three parameters seem to get smaller values of the estimates. However, the confidence intervals in both cases seem to get minor changes, with no clear pattern of the changes for any of the parameters.

In addition to this, we develop an Ideal Observer (IO) solution of the box task. This is done by finding the expected losses for the three choices we have each time a box is opened and then always making the decision with the least expected loss. These expected losses are the expected loss of choosing that blue is the majority colour, choosing that red is the majority colour and the expected loss of opening the next box. These IO solutions are visualised as decision trees that depend on parameters α and β.

In Chapter 1 we discussed that the box task is an alternative to the beads task to assess a 'jumping to conclusions' (JTC) bias. The parameters α and β say something about how reluctant one is to open boxes. As mentioned, it is hard to differentiate between the two parameters and the loss one gets. However, as β is the loss the participant gets when the test terminates in the limited trial, we might believe that α is the parameter that best describes the tendency to jump to conclusions. That small loss one gets each time a box is opened could be tied to the tendency to JTC. We see, for example, that participant 44 opens one box in all the unlimited trials. She gets a high value of $\hat{\alpha}$ as seen in (4.3). Additionally, participant 75 is one of the few to get $\hat{\alpha} \neq 0$ in the limited version, and she opens one box before deciding in all the limited trials. Thus, they both open only one box before choosing what they believe is the majority colour, and they both have high values of $\hat{\alpha}$. However, participant 75 also have a high value of $\hat{\beta}$, which might also indicate a JTC bias.

An adjustment we could do with the model is to remove α in the limited version. As we saw in Chapter 4.1.4, for the majority of the individuals, the $\hat{\alpha}$ is zero in the limited case. However, we saw that including α might be helpful in the situations where the participants make decisions with high expected losses.

We have assumed that participants have one value of both α and η in the unlimited case and then another in the limited case. A second possibility is that each participant has one single value of α and one single value of η that applies to both versions of the box task. We could then combine the two likelihood functions and minimise this one to find these values once instead of twice for each participant.

The parameter α is assumed to be constant and does not depend on the number of boxes that are opened, i. It could be a possibility that α varies with i. The same holds for β. This might be more realistic as participants can get more reluctant to open boxes when more boxes are opened. Moreover, in this thesis, it is assumed that T, the number of opened boxes when the test terminates, is uniformly distributed. We could also make a model where T varies with i.

In Chapter 4.1.5 we discussed that some participants get confidence intervals with length zero. That means that parametric bootstrapping is not adequate in these situations. Thus, it might be useful to find confidence intervals another way, for example, using nonparametric bootstrapping.

In this report, we discussed the sensitivity to hyperparameters using a beta prior for Θ with $\gamma = \kappa = 0.5$. We could try with other hyperparameters and see if the results will be even more different. For example, $\gamma = \kappa = 0.1$ might be an interesting case as this prior will be even more concentrated to the ends of the parameter space for Θ.

Bibliography

American Psychiatric Association (2013). *Diagnostic and statistical manual of mental disorders*. 5th ed. American Psyciatric Publishing.

Balzan, R. P., Ephramus, 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.

Bishop, C. (2013). *Pattern Recognition and Machine Learning: All "just the Facts 101" Material*. Information science and statistics. Springer (India) Private Limited.

Casella, G. and Berger, R. L. (2002). *Statistical Inference*. 2nd ed. Brooks/cole, Cengage Learning.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2009). *Introduction to Algorithms*. 3rd ed. MIT Press.

Dudley, R., Taylor, P., Wickham, S., and Hutton, P. (2016). "Psychosis, delusions and the "jumping to conclusions" reasoning bias: A systematic review and meta-analysis." *Schizophrenia bulletin*, 42, 652–665.

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman & Hall.

Garety, P., Freeman, D., Jolley, S., Ross, K., Waller, H., and Dunn, G. (2011). "Jumping to conclusions: the psychology of delusional reasoning." *Advances in psychiatric treatment*, 17, 332–339.

Givens, G. and Hoeting, J. (2012). *Computational Statistics*. Wiley Series in Computational Statistics. Wiley.

He, X., Samuelson, F., Gallas, B., Sahiner, B., and Myers, K. (2013). "The equivalence of a human observer and an ideal observer in binary diagnostic tasks." In *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 8673, 86730E. Society of Photo-optical Instrumentation Engineers.

Huq, S. F., Garety, P. A., and Hemsley, D. R. (1988). "Probabilistic judgements in deluded and non-deluded subjects." *The Quarterly journal of experimental psychology*, 40, 801–812.

Liese, F. and Miescke, K. (2008). *Statistical Decision Theory: Estimation, Testing, and Selection*. Springer Series in Statistics. Springer.

Moritz, S., Balzan, R. P., Kulagin, S. C., Göritz, A. S., Gaweda, L., and Andreou, C. (2017). "A new paradigm to measure probabilistic reasoning and a possible answer to the question why psychosis-prone individuals jump to conclusions." *Journal of Abnormal Psychology*, 126, 406–415.

Moritz, S. and Woodward, T. S. (2005). "Jumping to conclusions in delusional and non-delusional schizophrenic patients." *The British Journal of Clinical Psychology*, 44, 193–207.

Schay, G. (2016). *Introduction to Probability With Statistical Applications*. Birkhäuser.

Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. (1997). "Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization." *ACM Trans. Math. Softw.*, 23, 550–560.

Appendix A

The Trials and Draws to Decision in the Box Task

As we have data from 76 participants that have done the box task, we are presenting the order of the boxes in the nine trials they have done. We also include histograms of how many boxes the participants opened before they either chose the majority colour or the test terminated. This is called draws to decision.

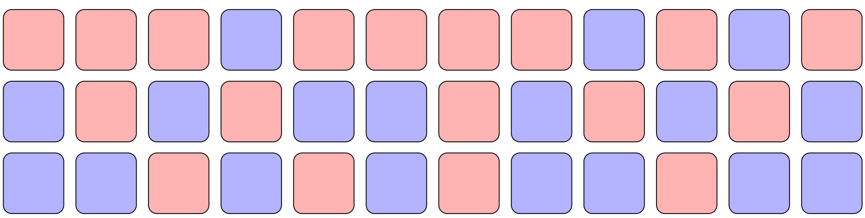

Figure A.1: The order of the boxes in the three unlimited trials. That is, Trial 2, 3 and 4.

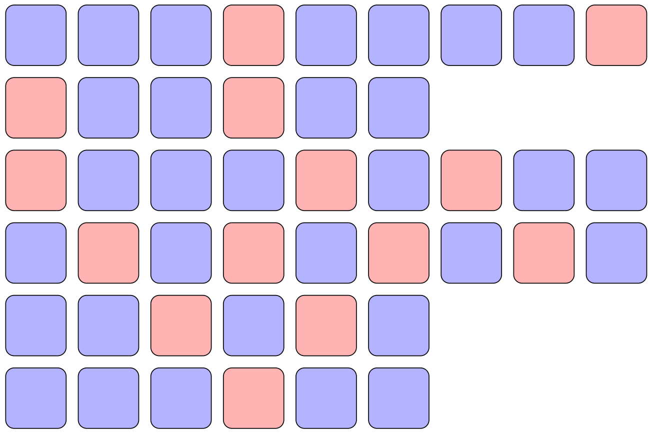

Figure A.2: The order of the boxes in the six limited trials. That is, Trials 5, 6, 7, 8 , 9 and 10. We see that Trials 5, 7 and 8 terminate after nine boxes are opened, whereas the other trials terminate after six boxes are opened.

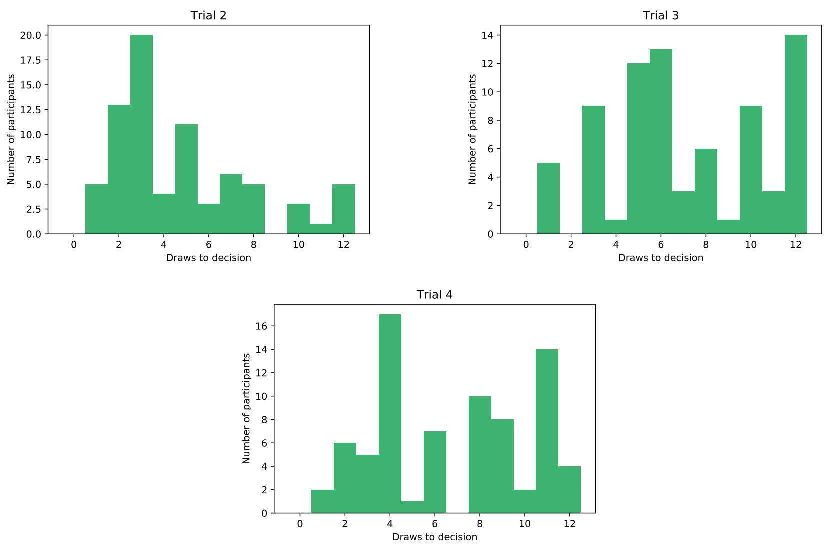

Figure A.3: Histogram of the draws to decisions for all participants in the three unlimited trials.

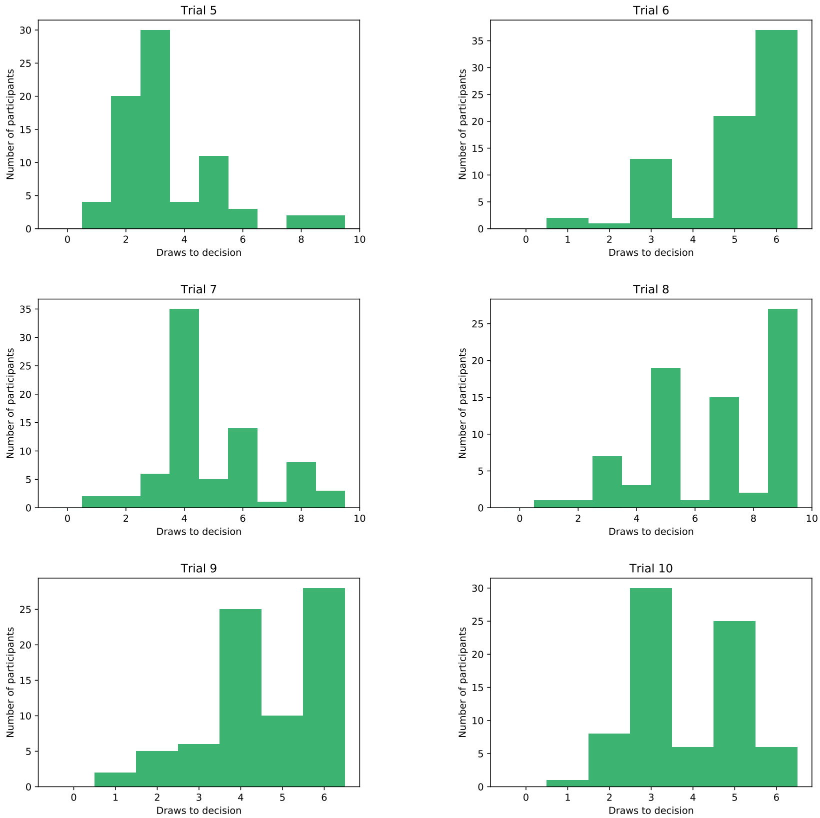

Figure A.4: The draws to decisions for all participants in the six limited trials. That is, how many boxes they open before they choose what they think is the majority colour, or before the test terminates.

Appendix B

Confidence Intervals

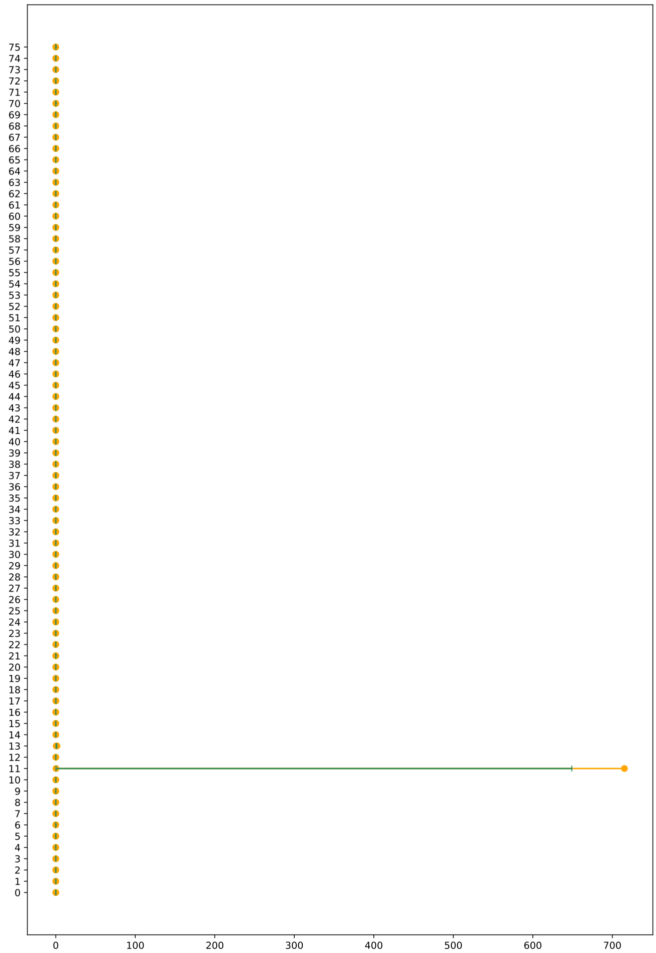

Figure B.1: Confidence intervals of α in the limited case for all participants with two different priors for Θ. The green lines represent CIs in the situation where we use a uniform prior, that is $\gamma = \kappa = 1$ and the orange lines for when we have $\gamma = \kappa = 0.5$.

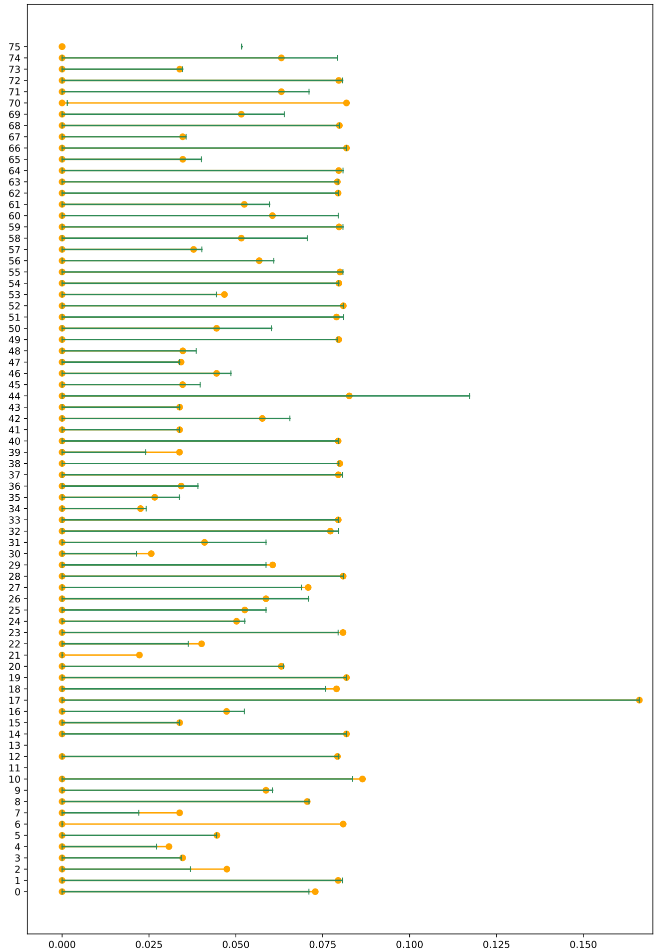

Figure B.2: Confidence intervals of α in the limited case for all participants with two different priors for Θ. Here we have zoomed in more on the plot in Figure 4.60

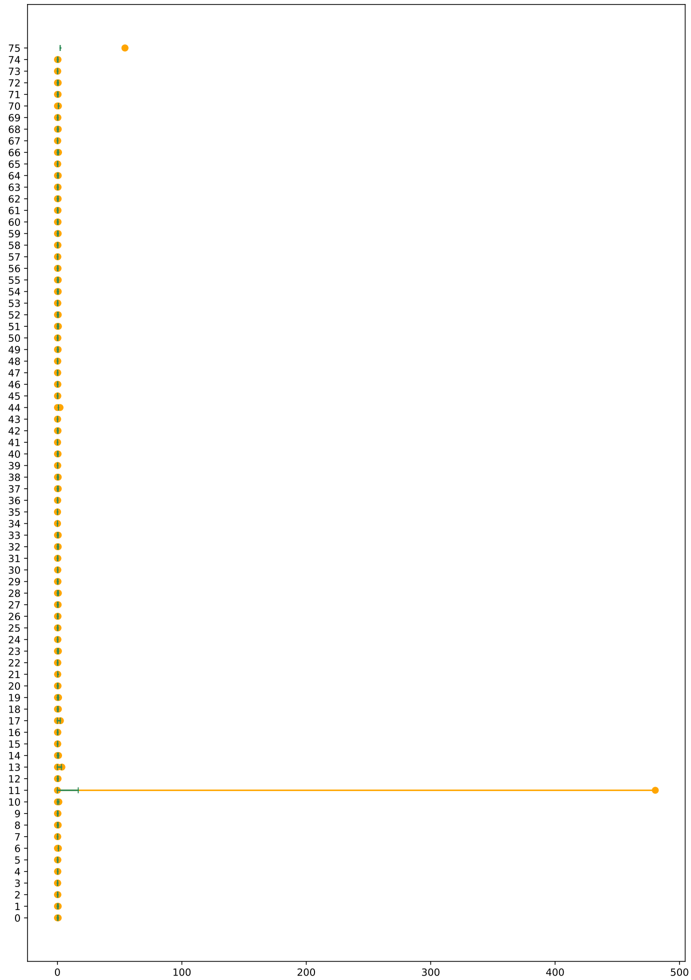

Figure B.3: Confidence intervals of β in the limited case for all participants with two different priors for Θ. The green lines represent CIs in the situation where we use a uniform prior, that is $\gamma = \kappa = 1$ and the orange lines for when we have $\gamma = \kappa = 0.5$.

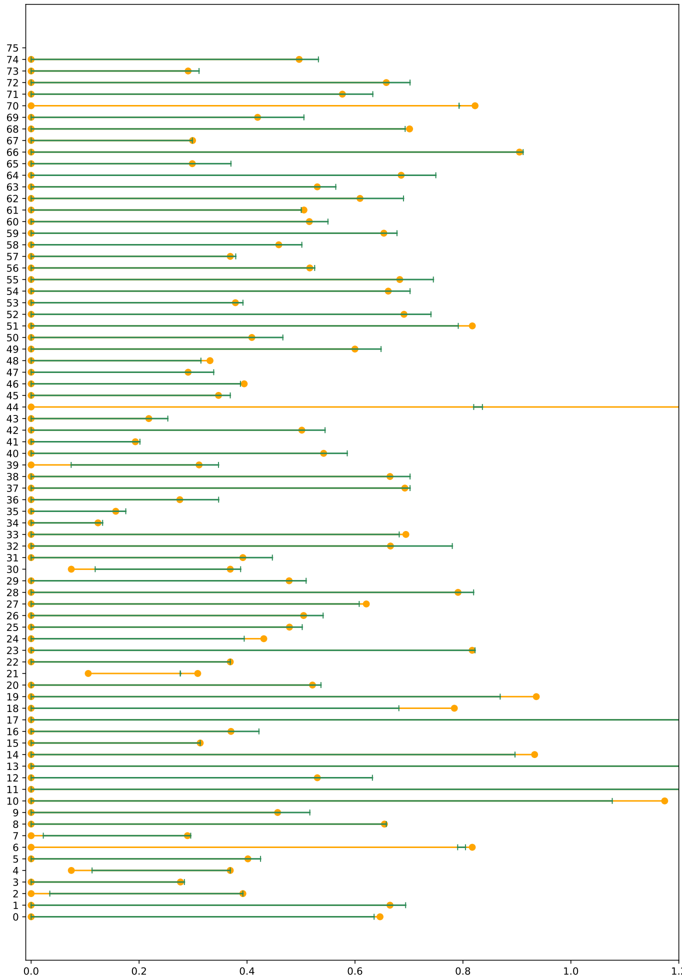

Figure B.4: Confidence intervals of β in the limited case for all participants with two different priors for Θ. Here we have zoomed in on the plot in Figure 4.61

Johanne Skogvang

Modelling Decisions in the Box Task

NTNU
Kunnskap for en bedre verden