Master's thesis

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

Ida Marie S. Falnes

Informational cascades where individuals have uncertain knowledge about others' competence

Master's thesis in Applied Physics and Mathematics
Supervisor: Håkon Tjelmeland
June 2020

NTNU
Norwegian University of
Science and Technology

Ida Marie S. Falnes

Informational cascades where individuals have uncertain knowledge about others' competence

Master's thesis in Applied Physics and Mathematics
Supervisor: Håkon Tjelmeland
June 2020

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

Preface

This thesis constitutes my work for the course 'TMA4900 - Industrial Mathematics, Master's Thesis' at the Department of Mathematical Sciences at NTNU, and completes my Master of Science degree. It is the result of my work in the last semester of the five-year study programme 'Applied Physics and Mathematics', where I have specialized in 'Industrial Mathematics'. The work here is a continuation of the work done in my specialization project for the course TMA4500, written during autumn 2019.

First and foremost, I would like to thank my supervisor Professor Håkon Tjelmeland for introducing me to the interesting topic of observational learning and informational cascades, and for his guidance throughout the semester. The last part of my thesis was written during the special time of the coronavirus outbreak and the strong restrictions that followed, and Håkon has continually been following me up and seamlessly provided feedback and guidance both before and during the lockdown.

I would also like to thank my fellow students for the wonderful time in Trondheim. A special thank you goes to Aleksander Gjersvoll, with whom I have shared home-office during the last months of work on the thesis. I am grateful for your support, patience and good company.

Ida Marie Falnes
June 10, 2020
Trondheim

Abstract

Informational cascades occur when rational individuals consider it optimal to ignore their private knowledge, and rather choose to copy the behaviour of their predecessors when making a decision. The phenomenon is closely related to the process of gaining information through observation of other individuals' actions - what is called observational learning.

In this thesis, we present a model of sequential decision-making with a binary action space. We assume that prior to making his or her decision, each individual observes the decisions of the previous decision-makers. We further assume that each decision-maker has a personal competence related to the decision at hand. All individuals have perfect knowledge about their own competence, but have only uncertain knowledge about other individuals' competences. We define the model in a mathematical fashion, and derive a general expression for the probability of both possible decisions for each individual.

The model is numerically implemented, and sequences of decisions are simulated. Using the Metropolis-Hastings algorithm, we investigate if there is enough information in the observed decisions alone to be able to estimate parameters from the model. Results from the simulation study indicate that there is not enough information in the observed decisions to obtain sufficiently accurate estimates for the model parameters. We suggest to improve the algorithm in order to increase the rate of convergence to the limiting distribution, in addition to allow more information to enter the system.

Sammendrag

Informasjonskaskader er et fenomen knyttet til beslutningstakning. De oppstår når rasjonelle individer betrakter det som optimalt å ignorere sin egen, private kunnskap, og heller velger å kopiere atferden til tidligere beslutningstakere. Fenomenet er tett knyttet til observasjonslæring; å samle informasjon gjennom observasjon av andre individers handlinger.

I denne rapporten presenterer vi en modell for sekvensiell beslutningstakning med et binært handlingsrom. Vi antar at hvert individ observerer beslutningen til tidligere beslutningstakere før han eller hun tar sin egen beslutning. I tillegg antar vi at hver beslutningstaker har en personlig kompetanse knyttet til beslutningen som skal tas. Hvert individ kjenner til sin egen kompetanse, men har kun et usikkert estimat på andres kompetanse. Vi definerer modellen matematisk, og utleder et generelt uttrykk for sannsynligheten for de to mulige beslutningene til hvert individ.

Modellen er numerisk implementert, og kjeder av beslutninger simuleres. Ved å bruke Metropolis-Hastings-algoritmen undersøker vi om det er tilstrekkelig informasjon om systemet i de observerte beslutningene til å estimere modellparametre. Resultatene fra simuleringsstudien tyder på at det ikke er nok informasjon i de observerte beslutningene alene til å oppnå rimelige estimater på modellparametrene. Vi foreslår å tillate mer informasjon i systemet, i tillegg til å forbedre algoritmen for å oppnå raskere konvergens til målfordelingen.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

An informational cascade is a phenomenon related to decision-making and *observational learning* - the process of gaining information through observation of other individuals' decisions. The phenomenon occurs in situations where rational individuals, after observing the decisions of other individuals, consider it most optimal to follow the existing pattern of behaviour and ignore their own, private knowledge. A familiar situation where an information cascade can occur is presented by Banerjee (1992). Imagine two restaurants placed next to each other, restaurant A and restaurant B. Upon arriving, you are not familiar with either restaurant, but have done some research and intend to go to restaurant A. However, you see that there are no customers in restaurant A, while restaurant B has many customers. Now, you might consider it optimal to go to restaurant B because you infer that the other customers have information that is unknown to you. Individuals arriving after you are also likely to go through a similar process of reasoning, and an informational cascade will occur. If all customers have uncertain information about which restaurant is the better, but restaurant B is randomly chosen by the first few customers, there can occur a cascade where all subsequent guests choose restaurant B even though restaurant A might be the objectively better restaurant. In particular, informational cascades may not be favourable for the public, because the lack of diversity in the observed decisions will fail to reflect the private knowledge of each decision-maker, leaving this information unknown for the public. The separate papers of Bikhchandani et al. (1992) and Banerjee (1992) are often mentioned as the first to describe the concepts of informational cascades. The phenomenon has later been a subject of extensive research, and is of interest in fields ranging from psychology and biology (Zentall, 2006) to behavioural economics and network analysis (Rosas et al., 2017).

We will in this thesis consider a model of sequential decision-making with a binary action space. The model is based on the model introduced by Bikhchandani et al. (1998), where decision-makers in sequence choose one out of two possible actions: Either to adopt or reject. Prior to their decision, they will observe a private signal indicating what the correct action is, in addition to the probability that their signal was correct. The lower this probability is, the more uncertain the decision-maker will be on his or her decision. Each of the decision-makers will also observe the actions of the previous decision-makers. This model is referred to as the 'observable-actions' model. Results are compared to the 'observable-signals' model, a benchmark model where all subsequent decision-makers has perfect knowledge both about the private signals and the decisions of the previous decision-makers. Bikhchandani et al. (1998) shows that under the observable-signals paradigm, all individuals will eventually make the same, correct decision. For the 'observable-actions' model, all individuals will conform on the same decision, either correct or wrong, if two individuals in a row chooses the same action. In both models, it is assumed that each decision-maker act rational with regard to the posterior probability. In Chapter 2.1, we look closely into an example of an informational cascade and the underlying rationale of each decision-maker in order to understand why this phenomenon occurs.

In the model to be considered in this thesis, we adapt the idea of using probabilities as a measure of uncertainty in the observed, private signal of each individual. In the model by Bikhchandani et al. (1998), each individual is assigned a probability that their private signal was correct, and all individuals have the same probability for observing the correct signal. Hence, this probability will govern the overall rate of wrong cascades. In our model, we will assume that each decision-maker observes different probabilities. If a decision-maker observes a private signal with a high probability that the signal was correct, this individual can be said to be better informed than an individual with a lower such probability. The concept of an individually assigned probability as an uncertainty in the observed signal can be interpreted as a personal competence, where the individual with the higher competence related to his private signal will have a greater prerequisite for making the correct choice. In a population of decision-makers, it is natural to assume that different individuals will have somewhat different competences.

The aim of this thesis is to continue the work done in Falnes (2019), where two models of sequential decision-making with a binary action space was introduced. It was assumed that each decision-maker would act rationally with regard to the posterior probability. In the two models, all decision-makers have different competences. In the first model, each decision-maker has perfect knowledge about all previous competences and decisions, but not the corresponding, private signals. In the second model, each individual has information about previous decisions, but not the corresponding signals or competences. However, in this model each individual will have his or her own uncertain estimate for the previous individuals' competences, but these estimates are independent of the true competence. The analysis showed that informational cascades occur for both models. In this thesis we will introduce and define a model based on the second model. The aim is to model the situation where each individual's competence estimates are correlated to the true competence. This model is implemented numerically. From this implementation, chains of decisions can be simulated under different choices of parameters. A Bayesian method for parameter estimation is derived. In order to assess this method, we estimate parameters from simulated data.

The report is structured as follows: In Chapter 2, we take a closer look at informational cascades and the mechanisms behind them in order to understand why this phenomenon occurs and to motivate further study. Additionally, some statistical background theory used later is also introduced here. In Chapter 3, the model of study is introduced and mathematically defined. In Chapter 4, we derive a method for Bayesian parameter estimation, and the results of the simulation study are presented and discussed in Chapter 5. In Chapter 6, the thesis is summed up, and suggestions for further work is presented.

Chapter 2

Background

In this chapter we present some relevant background theory. We first take a closer look at informational cascades. Further, we introduce some statistical concepts that we use later. These include Bayesian parameter estimation and the Metropolis-Hastings algorithm. We will also discuss how directed acyclic graphs (DAGs) can be used to represent conditional independence relations.

2.1 Informational cascades

To better understand what an informational cascade is and the mechanisms behind it seen from both the individual and public perspective, we begin by introducing the laboratory experiment conducted by Anderson and Holt (1997) with human test subjects. They constructed a game with two urns filled with balls of two different labels. In urn A, 2/3 of the balls were labeled 'a' and the last 1/3 were labeled 'b', while urn B contained 2/3 balls with the label 'b' and the last 1/3 were labeled 'a'. With an equal (prior) probability, one of these urns is chosen to be the correct urn. The aim of the participants is to correctly identify this urn, and to help them, they will get to observe one ball drawn at random (with replacement) from the correct urn. This is what we define as a *private signal*, and this information is hidden from the other participants. However, the decision of each decision-maker is announced publicly. This means that participants will know the decisions, but not the private signals of all previous participants. Sequentially, each participant receives his or her private signal and then make a guess on the correct urn based on the information available. The experiment showed that very often, individuals will tend to conform on one guess, despite the fact that their private signals suggested that the opposite decision was the correct one. This tendency was particularly prominent when the first few individuals conformed on one decision.

According to the definition of Bikhchandani et al. (1992), the participants in the above experiment are in an informational cascade. If we take a closer look at the rationale behind each decision, we will see that despite the fact that each individual decision-maker acts optimally and rationally with regard to the posterior probability, there is a positive probability that the public conforms on the inferior decision. We will in the following assume the same prior probability for all individuals, and that individuals act rationally with regard to the posterior probability.

For the first participant, the posterior probability that a ball is drawn from urn A, given that it is labeled 'a', is $P(A|a) = 2/3$. The first guess will then reveal the private signal of the first individual, which is now a part of the *public information*. The next participant will observe the first decision, and as a result, he or she can infer the private signal of the first individual. There are now two possible situations.

1. If the first guess was 'A' and individual 2 observed 'a' as his or her private signal, he or she knows that there has been two private signals indicating that 'A' is the correct urn. The similar holds if the first guess was 'B', and the private signal of individual 2 was a ball labeled 'b'.

2. If the first guess was 'A' ('B'), but individual 2 observed 'b' ('a'), he or she now sees urn A and B as equally likely events. In his or her eyes, there has now been a total of one 'a'-signal and one 'b'-signal. Consequently, his or her final decision will be random. This can be confirmed with a simple calculation using Bayes' rule, noting that each private signal is drawn independently from the correct urn:

$$P(A|a,b) = \frac{P(a,b|A)P(A)}{P(a,b)} = \frac{0.5 \cdot 2/3 \cdot 1/3}{0.5 \cdot 1/3 \cdot 2/3 + 0.5 \cdot 1/3 \cdot 2/3} = \frac{1}{2}.$$

The third participant will face one of the following situations.

1. If the two previous decisions were 'A' ('B'), the third individual will infer that there has been two 'a'-signals ('b'-signals) in total, or that there has been one of each, but the second person chose 'A' ('B') at random. Individual 3 and all of the succeeding participants are probable to infer that there has been a majority of one of the signals, and sees it rational to make the same decision as their predecessors, irrespective of their own, private signal.

2. If there has been one decision of each, the third participant knows that there has been one 'a'-signal and one 'b'-signal in total, and decides based on his or her private information. The next participant will face the same situation as participant 2, as the game effectively restarts when there has been one of each decision.

From the above example we see that as long as each decision differs from the previous, the public is able to draw information about private signals based on the corresponding decisions of each participant. As soon as a few individuals begin to favour one decision, new, similar decisions will not add any new information about decision-makers' private knowledge. As each new decision is uninformative for the succeeding decision-makers, each new individual will simply do the exact same reasoning as the individual before him. Effectively, succeeding individuals will infer that there is a majority of private signals favouring one of the labels, based only on the two first decisions in the cascade. This causes the counter-intuitive situation where each single decision is a result of logical processing of the available information, but the public as a whole may conform on the less optimal choice. According to Bikhchandani et al. (1992), informational cascades are most likely to occur in situations where individuals' knowledge alone is not sufficient to make an optimal choice, and where other individuals' decisions can be observed. The more uncertain an individual is, the more rational it is to see other individuals' decisions as more informative than his or her own information. In recent years, research on informational cascades concern, among other things, online shopping habits. Duan et al. (2009) argue that informational cascades are particularly prominent on the Internet because of the large number of products and the information overload. This makes it difficult for individuals to acquire the knowledge to make the optimal choice. At the same time, other individuals' choices are easily available as for example best-seller lists and other ranking systems, making it both rational and efficient to follow the choices of others.

2.2 Bayesian inference and MCMC

Statistical inference consists of methods to draw generalizations about populations (Walpole et al., 2012), and includes among other things estimation of unknown quantities. In this thesis, we are interested in estimation of parameters that are present in statistical models. Statistical models are typically defined through assumptions concerning relationships between random variables or observed data, and often consist of collections of probability distributions. The properties of such distributions are governed by its parameters. Bayesian statistical modelling is based on Bayes' theorem. From the Bayesian point of view, the parameters of interest are considered to be stochastic variables. This as opposed to the frequentist, or classical perspective, where parameters are treated as fixed constants. The main objective of Bayesian parameter estimation is to analyse the posterior distribution of the parameters using prior knowledge about the parameters in combination with the observed data. As an example, we denote θ as the vector of the parameters of interest. The knowledge or prior belief of the parameters, before any data is observed, is summarized in what is called the *prior distribution* $p(\theta)$. Further, we let z denote the observed data, and define the *likelihood* $p(z|\theta)$. With Bayes' theorem, we define the posterior distribution by

$$p(\theta|z) = \frac{p(z|\theta)p(\theta)}{p(z)} \propto p(z|\theta)p(\theta),$$

where we write the last transition as $p(z)$ does not depend on θ. The posterior distribution can be considered as an adjustment to our prior knowledge of the parameters after data is observed.

2.2.1 The Metropolis-Hastings Algorithm

This subsection is meant to serve as a reminder of the Metropolis-Hastings algorithm. For a more thorough introduction to the topic, the reader is referred to Gamerman and Lopes (2006). The Metropolis-Hastings (M-H) algorithm is a Markov chain Monte Carlo

(MCMC) technique first described by Metropolis et al. (1953) and later generalized by Hastings (1970). In general, MCMC methods are a collection of algorithms with the objective of sampling from probability distributions that often are high-dimensional or in other manners complex and hence difficult to sample from using direct methods. As the name suggests, MCMC-methods are based on the theory of Markov chains. Informally, the idea is to construct a Markov chain that has the desired distribution as its limiting distribution. As the number of iterations increases, the drawn states of the Markov chain become increasingly closer to the stationary distribution and can be considered approximate draws from the limiting distribution.

In this thesis, we will use the M-H algorithm to estimate the joint posterior distribution $p(\theta|z)$ where z is the observed data and $\theta = (\theta_1, \dots, \theta_l)$ a collection of parameters. The distribution we want to sample from is often referred to as the *target distribution*. We will use *single-site* updates. This means that we update one and one element of the parameter vector θ. We let $\theta^{(t)}$ denote the sample at iteration t in the algorithm. A proposal distribution $q(\theta^*|\theta^{(t)})$ has to be defined. The proposal distribution has to be chosen such that the constructed Markov chain is aperiodic and irreducible. These are sufficient conditions for convergence to a unique limiting distribution (Roberts and Smith, 1994). However, the convergence properties of the chain will be highly dependent on the choice of proposal distribution (Givens and Hoeting, 2013).

Samples from the target distribution - in our case $p(\theta|z)$ - is obtained by first defining the proposal distribution $q(\theta^*|\theta^{(t)})$. An initial value $\theta^{(0)}$ is set. This value will need to fulfill the condition $p(\theta^{(0)}|z) > 0$, but can otherwise be chosen arbitrarily. A proposal θ^* from $q(\theta^*|\theta^{(t)})$ is drawn and accepted with the M-H acceptance probability defined by

$$a = \min\left(1, \frac{p(\theta^*|z)}{p(\theta^{(t)}|z)} \cdot \frac{q(\theta^*|\theta^{(t)})}{q(\theta^{(t)}|\theta^*)}\right). \tag{2.1}$$

If the proposal is accepted, we set $\theta^{(t+1)} = \theta^*$. Otherwise, $\theta^{(t+1)} = \theta^{(t)}$ and a new value is proposed.

Since the drawn states from the first iterations typically depend on the initial value, they are not considered draws from the distribution of interest. Because of this, they should not be included when doing inference on the generated samples from an MCMC algorithm. The period characterized by the drawn states before the chain has reached its equilibrium distribution is often called the *burn-in period* and consists of a given number of iterations m. One of the main difficulties when using MCMC methods is to decide the number of iterations m, and hence verify whether or not the constructed Markov chain has converged sufficiently close to the limiting distribution. Theoretically, we need an infinite number of iterations to obtain samples from the target distribution. There exist many methods to assess convergence based on both visual inspections and statistical properties of the sampled distribution. In this report we consider it sufficient to utilize some visual inspections. One of these includes to run several chains from different initial values $\theta^{(0)}$ and investigate if they have the same behaviour after m iterations. If the chain is independent of the starting value, it is an indication that it has converged.

Graphically, this inspection can be performed by assessing a *trace plot*. This is a plot with the number of iterations t on the x-axis, and the corresponding state at time t, $\theta^{(t)}$, on the y-axis. Inspection of trace plots can give us an indication of which value for m we should choose to ensure that the samples we use for inference are representative for the limiting distribution. We can also use the trace plots to get a sense of the *mixing properties* of the chain. When evaluating the performance of an implemented MCMC-method, we are interested in how fast the chain converges and how well the target distribution is explored. This is related to the dependence between two drawn states and the number of iterations apart these states need to be before they are considered independent. If the drawn states seem to move rapidly around an equilibrium, we say that the mixing is good. On the other hand, if few values are accepted and the chain stays in the same state for many iterations in a row, the mixing is poor.

2.3 Conditional independence and DAGs

Two random variables, X and Y, are said to be conditionally independent given a third random variable Z, if and only if

$$f_{X,Y|Z}(x, y|z) = f_{X|Z}(x|z) \cdot f_{Y|Z}(y|z), \tag{2.2}$$

meaning they are independent in their conditional probability distribution given Z (see for example Dawid (1979)). A shorter notation is $X \perp\!\!\!\perp Y|Z$. As an intuitive explanation of the above expression, we can say that Y offers no additional knowledge about X when Z is known.

A directed acyclic graph (DAG) is a useful way to represent and visualise conditional independence relations among random variables. The following presentation about DAGs and conditional independence is inspired by Højsgaard et al. (2012, Ch. 1). We define a graph as the pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a set of vertices or nodes and $\mathcal{E}$ is a set of edges. In a DAG, the edges are *directed*, and the graph is *acyclic*, see Figure 2.1. A node a is the parent node of node b if there is a directed edge $a \to b$, and we denote the parental set of node b as $\text{pa}(b)$. Similarly, a node c is an ancestor of node b if it exists a directed path $c \mapsto b$, and the set of all ancestors of node b is denoted $\text{an}(b)$. We say the DAG $\mathcal{G}$ with vertices $(X_v)_{v\in\mathcal{V}}$ represent the probability distribution for $\mathcal{V}$ if

$$f(\mathbf{x}_{\mathcal{V}}) = \prod_{v\in\mathcal{V}} f(x_v|x_{\text{pa}(v)}), \tag{2.3}$$

where $f(x)$ is the probability function. Using the above in an example, we can write joint probability of the system in Figure 2.1 as factors of conditional probabilities

$$f(x_1, x_2, x_3, x_4, x_5, x_6) = f(x_1)f(x_2|x_1)f(x_3|x_1)f(x_4|x_2, x_3)f(x_5|x_2)f(x_6|x_3).$$

Let $\hat{X}_v$ denote all random variables except the variables represented by the descendants and parents of node v. From the definition of conditional probability given in (2.2) and from (2.3), it follows that $X_v \perp\!\!\!\perp \hat{X}_v|X_{\text{pa}(v)}$. We use the graph depicted in Figure 2.1 to

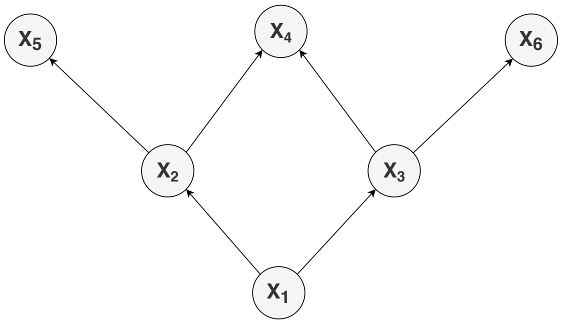

Figure 2.1: An example of a DAG. All edges are *directed*, illustrated by a one-way arrow. The graph is *acyclic*, meaning there are no directed cycles in the graph. Node 1 has no parent nodes, but is the parent node of nodes 2 and 3. The set $\{1, 2, 3\}$ constitute the ancestral set an(4) of node 4.

illustrate this property, by showing that the random variables X_1 and X_4 are independent given X_2 and X_3. The system can be factorized to

$$f(x_1, x_4, x_5, x_6|x_2, x_3) \propto f(x_1)f(x_2|x_1)f(x_3|x_1)f(x_4|x_2, x_3)f(x_5|x_2)f(x_6|x_3).$$

Integrating out irrelevant variables, we obtain

$$f(x_1, x_4|x_2, x_3) \propto f(x_1)f(x_2|x_1)f(x_3|x_1)f(x_4|x_2, x_3),$$

and we observe that x_1 and x_4 are independent in their conditional densities, and thus conditionally independent given x_2 and x_3 according to the definition.

Chapter 3

A model with uncertainty in others' competence

The aim of this chapter is to present the model of study. We will introduce the necessary notation and describe the model in a mathematical fashion. Based on the model description, we will further derive a general expression for the probability of each possible action of individual i. This expression is implemented numerically and will be used in the simulation studies presented in Chapter 5.

3.1 Definition and notation

As explained in the introduction, we wish to model the situation where individuals sequentially make decisions with a binary action space $\{0, 1\}$, and where one of the decisions is defined as the correct one. The correct decision, or the true value, is denoted $X \in \{0, 1\}$. This value will be the same, but is unknown, for all individuals $i = 1, \ldots, n$ in the sequence of decision-makers. Each individual in this sequence will first get to observe a private signal which we denote $Y_i \in \{0, 1\}$. Additionally, each individual is given the probability that their observed signal y_i is correct (equal to x). We denote this probability as p_i, and define it as

$$P(Y_i = x | X = x) = p_i.$$

This probability can be interpreted as a measure on how competent each individual is. Each individual i will only get to observe his or her own competence p_i. However, individual i will have his or her own estimates of the previous individuals' competences. Further, individual i will regard his or her competence estimates as the true competence of the prior individuals. We let j denote the index of a previous individual, such that $j = 1, \ldots, i - 1$. Individual i's estimate for p_j, the unknown competence of individual j, is denoted $\tilde{p}_j^i$. Finally, the decision of individual i is stochastic, and denoted $Z_i \in \{0, 1\}$. It is a guess based on his or her knowledge about x, which is summarized in the posterior

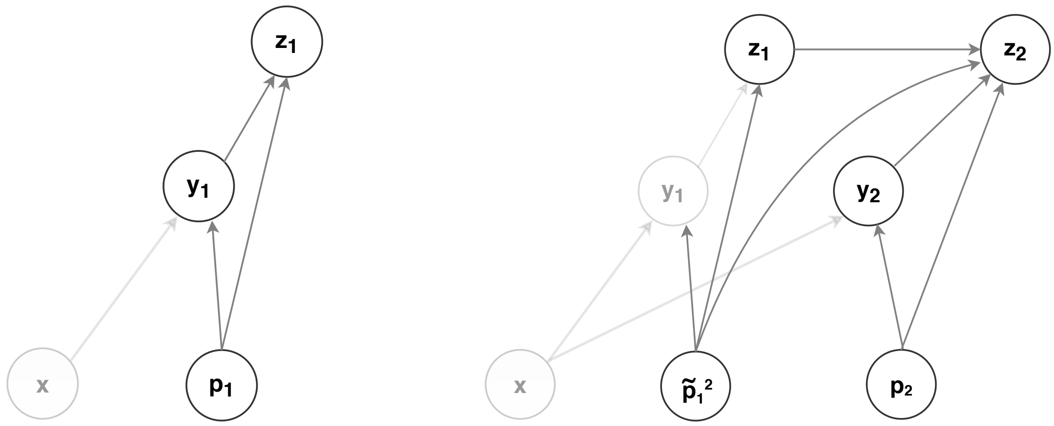

(a) Individual 1. (b) Individual 2.

Figure 3.1: Illustration of dependencies between all parameters for the first two decisions z_1 (a) and z_2 (b). Here, y_1 and y_2 denotes the private signals of individuals 1 and 2. The true value is denoted x, and the probabilities p_1 and p_2 denotes the true competences of each individual. The shaded nodes illustrates unknown variables. In addition to x, which is unknown for all individuals, individual 2 will not observe the private signal, y_1, of individual 1. Additionally, individual 2 will not observe the true competence of individual 1, but regards his or her estimate $\tilde{p}_1^2$ as the true competence.

probability

$$P(X = x|Y_i = y_i, Z_{i-1} = z_{i-1}, \ldots, Z_1 = z_1, p_i, \tilde{p}_{i-1}^i, \ldots, \tilde{p}_1^i). \tag{3.1}$$

Figure 3.1 visualises the model with graphs seen from the perspective of the first two individuals. In (a), we see that the decision of individual 1, z_1, only depends on his or her own competence p_1 and private signal y_1. The second individual will observe the first decision and take it into account when the next decision in the chain is made. However, individual 1's private signal y_1 is not available, and he or she only has an uncertain estimate $\tilde{p}_1^2$ of the competence of the first decision-maker. Since individual number 2 regards his competence estimate as the true competence of individual 1, p_1 is not included in the model graph for the second decision.

As noted in Chapter 2.3, model graphs are useful for visualising conditional independence relations among a set of random variables. In our model, we assume that each private signal y_i is conditionally independent of each other given the corresponding competence p_i or corresponding competence estimate $\tilde{p}_i$ and the true value x. This means that

$$f(y_1, \ldots, y_n|x, \tilde{p}_1, \ldots, \tilde{p}_{n-1}^n, p_n) = f(y_1|x, \tilde{p}_1^n) \ldots f(y_{n-1}|x, \tilde{p}_{n-1}^n) f(y_n|x, p_n), \tag{3.2}$$

a property we state for later reference. As seen in (3.1), each decision z_i will depend on all of the available information individual i has. The model is defined in a sequential fashion. This means that the decision at index i can only depend on previous decisions z_j for $j < i$. We assume that each decision z_i depends on the corresponding private signal of individual i, but not other individuals' private signals. Similarly, z_i depends on the competence and competence estimates of individual i. As a result, the decision z_i will be conditionally independent of the true value x given y_i, p_i and $\tilde{p}_1^i, \dots, \tilde{p}_{i-1}^i$. In terms of equations, we use the model graphs to see that the joint conditional distribution of the full system of previous decisions from individual n's point of view, given all private signals and competence estimates are given by

$$\begin{aligned}&f(z_1, \dots, z_{n-1}, x, p_n | y_1, \dots, y_n, \tilde{p}_1^n, \dots, \tilde{p}_{n-1}^n) \\ &\propto f(x) f(z_1, \dots, z_{n-1} | y_1, \dots, y_{n-1}, \tilde{p}_1^n, \dots, \tilde{p}_{n-1}^n) f(p_n),\end{aligned} \tag{3.3}$$

where we have used that the decisions up to z_{n-1} must be independent of individual n's private signal y_n. Integrating over p_n on both sides of the above expression, we get

$$\begin{aligned}&f(z_1, \dots, z_{n-1}, x | y_1, \dots, y_n, \tilde{p}_1^n, \dots, \tilde{p}_{n-1}^n) \\ &\propto f(x) f(z_1, \dots, z_{n-1} | y_1, \dots, y_{n-1}, \tilde{p}_1^n, \dots, \tilde{p}_{n-1}^n),\end{aligned} \tag{3.4}$$

and we observe that x and $z_1, \dots, z_{n-1}$ are independent in their conditional densities. Finally, we note that the i'th decision is independent of future decisions, competences and competence estimates. Using this, we can repeatedly apply the definition on conditional probability (sometimes known as the chain rule of probability) to rewrite the joint conditional distribution over all decisions $z_1, \dots, z_{n-1}$ above. We obtain

$$\begin{aligned}&f(z_1, \dots, z_{n-1} | \tilde{p}_1^n, \dots, \tilde{p}_{n-1}^n, y_1, \dots, y_{n-1}) \\ &= f(z_1 | y_1, \tilde{p}_1^n) \prod_{i=2}^{n-1} f(z_i | z_1, \dots, z_{i-1}, \tilde{p}_1^n, \dots, \tilde{p}_i^n, y_i).\end{aligned} \tag{3.5}$$

3.2 Competences and competence estimates

Above, the probability p_i is introduced as a measure of the competence of individual i. We elaborate on this description by noting that the closer p_i is to 1, the more probable it is that individual i's observed signal is correct. Hence, if p_i is large, individual i will have a greater prerequisite to make the correct choice compared to individuals with lower competences. Since we assume that individuals come from the same population, we let all competences p_i for $i = 1, \dots, n$ be independently and identically distributed. In our model, it is natural to assume that each competence $p_i \in (0.5, 1)$. If individual i observes a signal y_i with probability $p_i < 0.5$, he or she will draw the conclusion that the observed signal was most probably wrong, and that the opposite of the signal, $1 - y_i$, is correct with probability $1 - p_i$. Because of this symmetry, we will limit the competences to the interval $(0.5, 1)$. In our model, this is done by assuming that the competences come from a transformed beta-distribution with parameters α and β. A standard beta distributed variable p have bounds $(0, 1)$, and in general we can transform this to having bounds (a, b) by the relationship

$$p_i = (b - a)p + a. \tag{3.6}$$

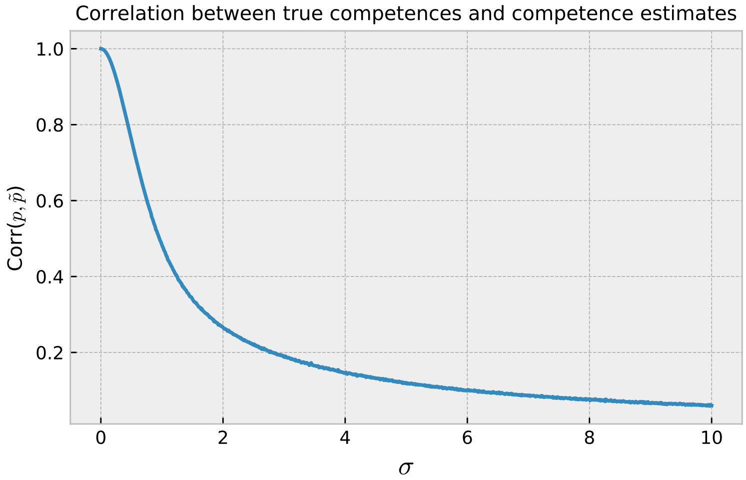

Figure 3.2: Simulated competences p and competence estimates $\tilde{p}$ for different values of σ. The original competences are beta distributed with parameters $\alpha = 5$ and $\beta = 10$.

The probability density function of a nonstandard beta-distribution with bounds (a, b) is found by normalising the standard beta density $f((X-a)/(b-a))/(b-a)$ and is given by

$$f(p) = \frac{(p-a)^{\alpha-1}(b-p)^{\beta-1}}{B(\alpha,\beta)(b-a)^{\alpha+\beta-1}},$$

where we use $(a, b) = (0.5, 1)$.

According to the model description, decision-makers will only know the exact value of their own private competence, but have uncertain estimates of the competences of previous decision-makers. These estimates should be correlated to the original competences in order to model how well individuals in the system know each other or their ability to judge other individuals' competences. This is modelled by first taking the log-odds (also known as logit) transform to create a map of the probability values from $(0, 1)$ to $(-\infty, \infty)$. By adding noise from a known distribution to this transformation, we can control the correlation between the probability estimates and the original probabilities. We obtain

$$t_i = \text{logit}\left((p_i - 0.5)/0.5\right) + \varepsilon, \tag{3.7}$$

where $\text{logit}(p_i) = \log(p_i/(1-p_i))$ and where we assume that $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. We then use the inverse logit function to re-transform probabilities back to $(0.5, 1)$. As a result, individual n's estimates for the previous individuals' competences are given by

$$\tilde{p}_i^n = \left(\frac{e^{t_i}}{1+e^{t_i}}\right) \cdot (1-0.5) + 0.5,$$

where we have used (3.6). Figure 3.2 illustrates how the correlation between the original competences p and the estimated competences decreases as σ increases. Figure 3.3 shows

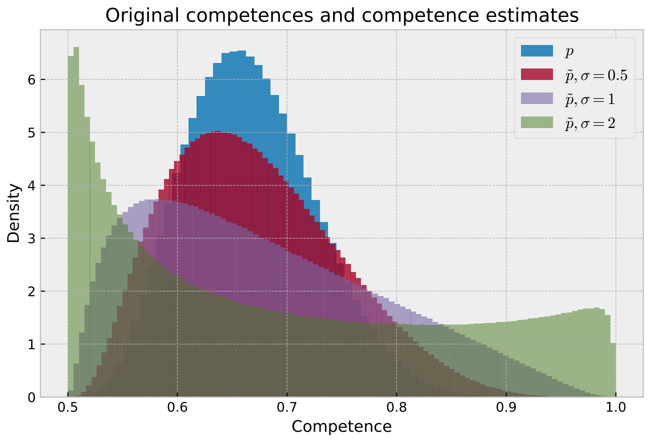

Figure 3.3: Simulated competences p and competence estimates $\tilde{p}$ for different values of σ. The original competences are beta distributed with parameters $\alpha = 5$ and $\beta = 10$.

the distribution of the original competences and the competence estimates for some values of σ.

For later reference, we note that given p_i, t_i given by (3.7) follows a normal distribution with $\mathrm{E}(t_i) = \mu = \mathrm{logit}((p_i - 0.5)/0.5)$ and $\mathrm{Var}(t_i) = \sigma^2$. In general, a random variable U whose logit transformation follows a normal distribution with mean μ and standard deviation σ, follows the logitnormal distribution. This distribution has probability density function (pdf) given by

$$f(u) = \frac{1}{\sqrt{2\pi}\sigma} \cdot \frac{1}{u(1-u)} \exp\left\{ -\frac{1}{2} \left(\frac{\mathrm{logit}(u) - \mu}{\sigma} \right)^2 \right\}, \tag{3.8}$$

see for example Frederic and Lad (2008). Since $\mathrm{logit}((\tilde{p}_i^n - 0.5)/0.5) = t_i$, the distribution of the competence estimates shifted to $(0, 1)$ given the true competences is logitnormal with location parameter $\mu = \mathrm{logit}((p_i - 0.5)/0.5)$ and scale parameter σ.

It is not obvious how different choices of the parameter σ will affect the resulting chains of decisions. When σ is chosen rather low, the competence estimates will be close to the true competences in the population. Hence, the decisions will depend on the true distribution of p. The situation where $\sigma \to 0$ resembles Model 1 in Falnes (2019), where all individuals knows the exact value of each others competence. If σ is very large, a given individual i will get very high competence-estimates for some of the previous decision-makers (close to 1), and very low for others (close to 0.5). This will happen at random and independently of the actual competences of the previous individuals.

3.3 Sub-optimal decisions

Based on the model definition, the *optimal* decision of individual n is deterministic, and given by

$$z_{n,\text{optimal}} = \underset{x}{\operatorname{argmax}}\, P(X = x | Y_n = y_n, Z_{n-1} = z_{n-1}, \ldots, Z_1 = z_1, p_n, \tilde{p}^n_{n-1}, \ldots, \tilde{p}^n_1). \tag{3.9}$$

In order to simplify notation, we use boldface letters to denote the vector of previous decisions $\mathbf{z_{n-1}} = (z_1, \ldots, z_{n-1})$ and the vector of individual n's competence estimates $\tilde{\mathbf{p}}^{\mathbf{n}} = (\tilde{p}^n_1, \ldots, \tilde{p}^n_{n-1})$.

We assume that decision-makers have a positive probability of making the least optimal decision. This will make our model more realistic. Decision-makers - human beings, for instance - are usually not able to perform exact calculations of complex quantities like (3.1) and deterministically make the optimal decision. There is naturally an uncertainty in most decision-making processes, and we can model this by adding noise to our model. In order to do this, we introduce the softmax function, which we use to weight the two posterior probabilities $P(X = 0 | y_n, p_n, \tilde{\mathbf{p}}^{\mathbf{n}}, \mathbf{z_{n-1}}))$ and $P(X = 1 | y_n, p_n, \tilde{\mathbf{p}}^{\mathbf{n}}, \mathbf{z_{n-1}}))$. The decision of individual n is then random, and expressed in terms of the probability

$$\begin{aligned} &P(Z_n = z_n | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}) \\ &= \frac{\exp\left(P(X = z_n | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}})/\tau\right)}{\exp\left(P(X = z_n | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}})/\tau\right) + \exp\left(P(X = 1 - z_n | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}})/\tau\right)}. \end{aligned} \tag{3.10}$$

The parameter τ is a control parameter used to model the degree of randomness in the final decisions. This is illustrated in Figure 3.4 where the probability for the n'th decision in a simulated decision-chain is plotted against τ. The figure shows that as $\tau \to 0$, the probability of choosing the optimal decision, given by (3.9), goes to 1. Since $P(Z_n = z_n | \cdot) + P(Z_n = 1 - z_n | \cdot) = 1$, we have that as $\tau \to \infty$, the probability of each decision approaches 0.5, and the final decision is random.

3.4 Derivation of the n'th decision

We are interested in simulating chains of decisions. In order to do that, we need to be able to calculate the expression given in (3.10). Hence, we take a closer look at each of the probabilities

$$P(X = 0 | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}) \quad \text{and} \quad P(X = 1 | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}),$$

where $P(X = 0 | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}) + P(X = 1 | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}) = 1$. From these probabilities, we will derive expressions that only include the known quantities presented in sections 3.1-3.3. We begin by using the definition on conditional probability. We can write

$$P(X = x | y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}) = \frac{P(x, y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}})}{P(y_n, p_n, \mathbf{z_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}})}. \tag{3.11}$$

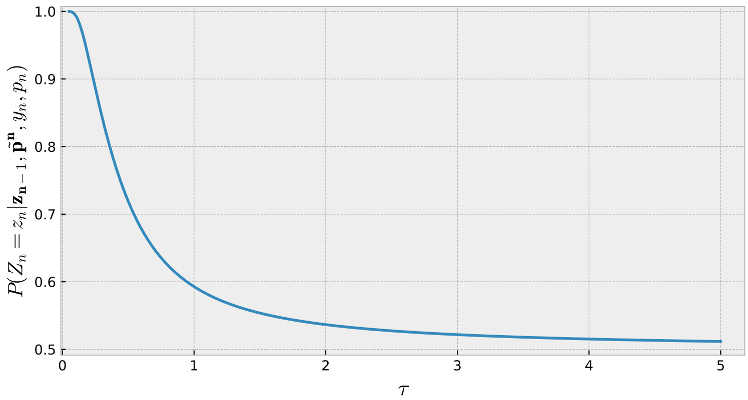

Figure 3.4: Illustration of how the parameter τ controls the weights of the probability of a given decision, $P(Z_n = z_n|y_n, p_n, \mathbf{z_{n-1}}, \mathbf{\tilde{p}^n})$. When $\tau \to 0$, individuals follows the optimal behaviour given by (3.9), and as $\tau \to \infty$, the probability of choosing each of the two possible actions approaches 1/2.

Continuing, we use the definition on joint probability in terms of conditional distributions to rewrite the fraction in (3.11). We look at the expressions in the numerator and denominator separately, and obtain

$$P(X = x|y_n, p_n, \mathbf{z_{n-1}}, \mathbf{\tilde{p}^n}) = \frac{P(\mathbf{z_{n-1}}, y_n|x, p_n, \mathbf{\tilde{p}^n})P(x, p_n, \mathbf{\tilde{p}^n})}{P(\mathbf{z_{n-1}}, y_n|p_n, \mathbf{\tilde{p}^n})P(p_n, \mathbf{\tilde{p}^n})}.$$

From the dependency graphs in Figure 3.1, we can observe that x is independent of p_n and $\mathbf{\tilde{p}^n}$. Hence, we get $P(x, p_n, \mathbf{\tilde{p}^n}) = P(x)P(p_n, \mathbf{\tilde{p}^n})$ in the numerator, and cancel the common factors in the numerator and the denominator. This yields

$$P(X = x|y_n, p_n, \mathbf{z_{n-1}}, \mathbf{\tilde{p}^n}) = \frac{P(\mathbf{z_{n-1}}, y_n|x, p_n, \mathbf{\tilde{p}^n})P(x)}{P(\mathbf{z_{n-1}}, y_n|p_n, \mathbf{\tilde{p}^n})} := \frac{c}{d}. \tag{3.12}$$

We will now focus on the expression in the numerator of (3.12). As noted in Section 3.1, the model is defined sequentially. As a result, we must have that the private signal of individual n, y_n, are conditionally independent of the previous decisions $\mathbf{z_{n-1}}$ given p_n and x. This follows trivially from the model definition. Hence we have that

$$P(\mathbf{z_{n-1}}, y_n|x, p_n, \mathbf{\tilde{p}^n})P(x) = P(y_n|x, p_n)P(\mathbf{z_{n-1}}|x, \mathbf{\tilde{p}^n})P(x).$$

Let $y_1, \ldots, y_{n-1}$ denote the previous individuals' private signals. These are unknown to individual n. To account for this, we need to use the law of total probability. In this way, we introduce sums over each previous y_i to obtain an expression we can compute. This gives us an expression for the numerator in 3.12

$$c = P(y_n|x, p_n)P(x) \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}=0}^{1} P(\mathbf{z_{n-1}}|y_1, \ldots, y_{n-1}, \mathbf{\tilde{p}^n})P(y_1, \ldots, y_{n-1}|x, \mathbf{\tilde{p}^n}),$$

where we have used that $\mathbf{z_{n-1}}$ is conditionally independent of x given $\mathbf{y_{n-1}} = (y_1, \ldots, y_{n-1})$ and $\tilde{\mathbf{p}}^{\mathbf{n}}$, as discussed in Section 3.1 and shown in equations (3.3) and (3.4).

Now shifting focus to the denominator of (3.12). We are in a similar situation as previously, as the previous private signals $y_1, \ldots, y_{n-1}$ are unknown. Unlike before, the variable x is not given. Hence, we use the law of total probability over both the previous private signals and x. Doing this, we obtain

$$d = \sum_{x'=0}^{1} \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}=0}^{1} P(\mathbf{z_{n-1}}, y_n | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}, x', y_1, \ldots, y_{n-1}) P(x', y_1, \ldots, y_{n-1} | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}).$$

As before, we use that the current private signal y_n is independent of the previous decisions $\mathbf{z_{n-1}}$, and get

$$\sum_{x'=0}^{1} \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}=0}^{1} P(\mathbf{z_{n-1}}, y_n | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}, x', y_1, \ldots, y_{n-1}) P(x', y_1, \ldots, y_{n-1} | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}) =$$
$$\sum_{x'=0}^{1} P(y_n | x', p_n) \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}=0}^{1} P(\mathbf{z_{n-1}} | \tilde{\mathbf{p}}^{\mathbf{n}}, x', y_1, \ldots, y_{n-1}) P(x', y_1, \ldots, y_{n-1} | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}).$$

We again use that $\mathbf{z_{n-1}}$ is conditionally independent of x' given $\mathbf{y_{n-1}}$ and $\tilde{\mathbf{p}}^{\mathbf{n}}$. We can then write

$$d = \sum_{x'=0}^{1} P(y_n | x', p_n) \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}=0}^{1} P(\mathbf{z_{n-1}} | \tilde{\mathbf{p}}^{\mathbf{n}}, y_1, \ldots, y_{n-1}) P(x', y_1, \ldots, y_{n-1} | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}).$$

Focusing on the last factor $P(x', \mathbf{y_{n-1}} | p_n, \tilde{\mathbf{p}}^{\mathbf{n}})$, we use the definition on conditional probability and obtain

$$P(x', \mathbf{y_{n-1}} | p_n, \tilde{\mathbf{p}}^{\mathbf{n}}) = \frac{P(x', \mathbf{y_{n-1}}, p_n, \tilde{\mathbf{p}}^{\mathbf{n}})}{P(p_n, \tilde{\mathbf{p}}^{\mathbf{n}})}.$$

We then use the definition on joint probability in terms of conditional probabilities in the numerator to obtain

$$\frac{P(x', \mathbf{y_{n-1}}, p_n, \tilde{\mathbf{p}}^{\mathbf{n}})}{P(p_n, \tilde{\mathbf{p}}^{\mathbf{n}})} = \frac{P(\mathbf{y_{n-1}} | x', p_n, \tilde{\mathbf{p}}^{\mathbf{n}}) P(x', p_n, \tilde{\mathbf{p}}^{\mathbf{n}})}{P(p_n, \tilde{\mathbf{p}}^{\mathbf{n}})} = P(\mathbf{y_{n-1}} | x', p_n, \tilde{\mathbf{p}}^{\mathbf{n}}) P(x'),$$

since x' and $p_n, \tilde{\mathbf{p}}^{\mathbf{n}}$ are independent. Further, the private signals of individuals prior to individual n, $\mathbf{y_{n-1}}$, must be independent of individual n's competence. Using this, we can write $P(\mathbf{y_{n-1}} | x', p_n, \tilde{\mathbf{p}}^{\mathbf{n}}) = P(\mathbf{y_{n-1}} | x', \tilde{\mathbf{p}}^{\mathbf{n}})$.

As a result of all of the above, we can write the fraction in (3.11) as

$$P(X = x | y_n, z_1, \ldots, z_{n-1}, p_n, \tilde{p}_1^n, \ldots, \tilde{p}^n) = \frac{P(\mathbf{z_{n-1}}, y_n | x, p_n, \tilde{\mathbf{p}}^{\mathbf{n}}) P(x)}{P(\mathbf{z_{n-1}}, y_n | p_n, \tilde{\mathbf{p}}^{\mathbf{n}})}$$
$$= \frac{P(y_n | x, p_n) P(x) \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}}^{1} P(\mathbf{z_{n-1}} | y_1, \ldots, y_{n-1}, \tilde{\mathbf{p}}^{\mathbf{n}}) P(y_1, \ldots, y_{n-1} | x, \tilde{\mathbf{p}}^{\mathbf{n}})}{\sum_{x'=0}^{1} P(y_n | x', p_n) P(x') \sum_{y_1=0}^{1} \cdots \sum_{y_{n-1}=0}^{1} P(\mathbf{z_{n-1}} | \tilde{\mathbf{p}}^{\mathbf{n}}, y_1, \ldots, y_{n-1}) P(y_1, \ldots, y_{n-1} | x', \tilde{\mathbf{p}}^{\mathbf{n}})}. \tag{3.13}$$

The above expression can be simplified by rewriting the joint probabilities of the decisions $\mathbf{z_{n-1}}$ and private signals $\mathbf{y_{n-1}}$. Beginning with the decisions, we note by the model definition that each decision z_i depends only on the previous decisions $z_1, \ldots, z_{i-1}$, and not future decisions as shown in (3.5). Using this, we can write the joint, conditional probability

$$\begin{aligned} P(\mathbf{z_{n-1}}|\mathbf{y_{n-1}}, \tilde{\mathbf{p}}^{\mathbf{n}}) &= P(z_1|y_1, \tilde{p}_1^n)P(z_2|z_1, y_2, \tilde{p}_1^n, \tilde{p}_2^n) \ldots P(z_{n-1}|\mathbf{z_{n-2}}, \tilde{\mathbf{p}}^{\mathbf{n}}, y_{n-1}) \\ &= \prod_{i=1}^{n-1} P(z_i|\mathbf{z_{i-1}}, \tilde{\mathbf{p}}_{\mathbf{i}}^{\mathbf{n}}, y_i). \end{aligned}$$

Given the true value x and the current competence estimate $\tilde{p}_i^n$, the private signals are conditionally independent of each other as showed in (3.2). Hence, we can write

$$P(\mathbf{y_{n-1}}|\tilde{\mathbf{p}}^{\mathbf{n}}, x) = \prod_{i=1}^{n-1} P(y_i|\tilde{p}_i^n, x).$$

We use the above, and finally write (3.13) as

$$\begin{aligned} &P(X = x|y_n, z_1, \ldots, z_{n-1}, p_n, \tilde{p}_1^n, \ldots, \tilde{p}^n) \\ &= \frac{P(y_n|x, p_n)P(x)\prod_{i=1}^{n-1}\sum_{y_i=0}^{1} P(z_i|\mathbf{z_{i-1}}, \tilde{\mathbf{p}}_{\mathbf{i}}^{\mathbf{n}}, y_i)P(y_i|x, \tilde{p}_i^n)}{\sum_{x'=0}^{1} P(y_n|x', p_n)P(x')\prod_{i=1}^{n-1}\sum_{y_i=0}^{1} P(z_i|\mathbf{z_{i-1}}, \tilde{\mathbf{p}}_{\mathbf{i}}^{\mathbf{n}}, y_i)P(y_i|\tilde{p}_i^n, x')}. \end{aligned}$$

In order to calculate the above expression, one need the recursion given by the expression inside the product sum in both the numerator and the denominator. This is found by calculating the probability for each of the previous decisions given by the softmax function in (3.10).

3.5 Simulating chains of decisions

The model described and derived in the previous subsections is implemented in the programming language Python. Random sampling from known probability distributions is done using routines from the SciPy Statistics library[1]. In the implementation of the model, we have assumed that the prior distribution of the unknown, true value X is $f(x) = 0.5$. This means that prior to observing the private signal and other individuals' decisions, each decision-maker sees each value of X as equally likely.

Figure 3.5 shows examples of simulated chains of decisions for different choices of the parameters α and β, which governs the distribution of the true competence p_i for each individual. In the figure, values for τ and σ are held constant. We use $\tau = 0.05$, which is rather low and corresponds to a high probability of making the optimal choice. This mimics the situation where individuals make mostly rational decisions. We have

[1]SciPy v1.3.3 Reference Guide `https://docs.scipy.org/doc/scipy/reference/index.html`, accessed 14.05.2020

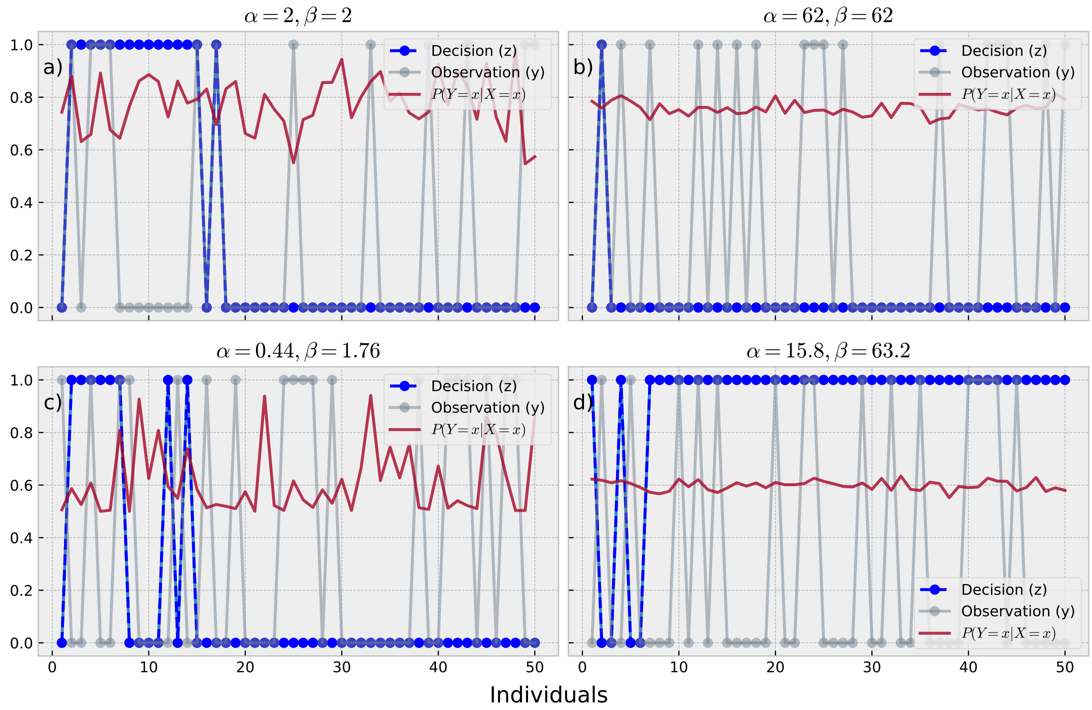

Figure 3.5: Simulated chains of decisions of $n = 50$ individuals for four different choices of the parameter-pairs α and β. The true value is $x = 0$ in all simulations. In plots a) and b) in the upper row, the mean value of the simulated competences is $E(p_i) = 0.75$. For the lower row, the mean value of the simulated competences is $E(p_i) = 0.6$. The standard deviation in plot a) and c) is $SD(p_i) = 0.1$, while for b) and d) $SD(p_i) = 0.02$.

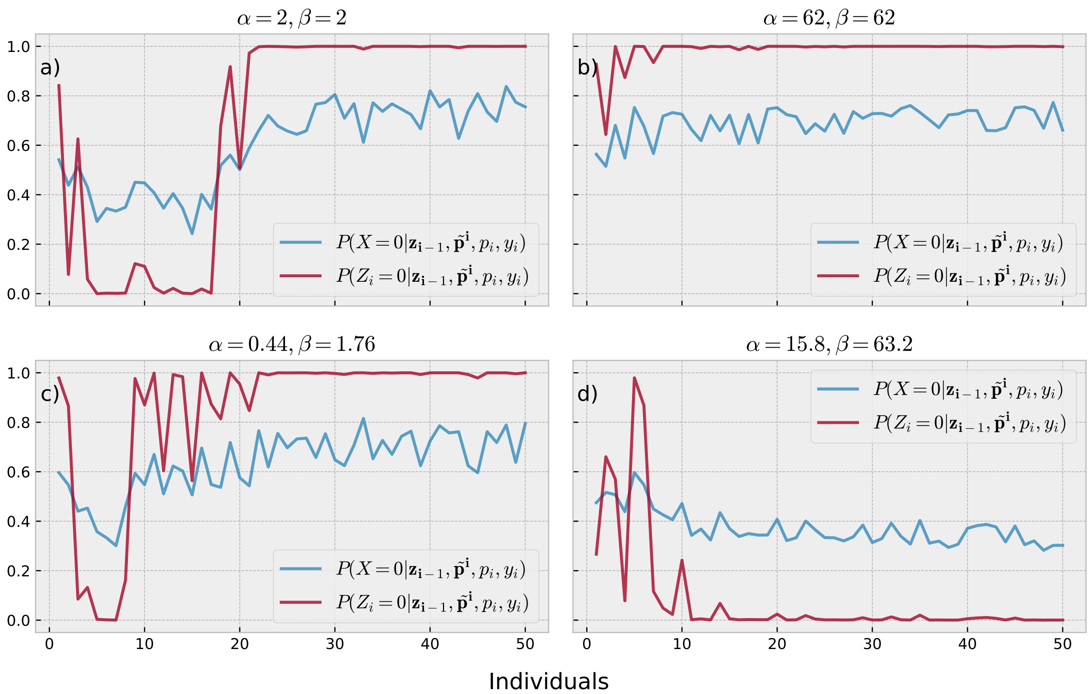

Figure 3.6: Probabilities for $X = 0$ for each individual before and after applying the softmax function. Corresponds to the simulation in Figure 3.5.

let $\sigma = 1$, which makes the correlation between the competence estimates and the true competences $Corr(p, \tilde{p}) = 0.47$. Competence estimates are thus somewhat close to the true competences. The true value $X = 0$ in all situations. The plots show that for these specific examples, cascades happen in all cases. In general, the typical behaviour of a chain seems to be that the first few decisions will to various degrees vary between 0 and 1. This is as expected, as decision-makers have observed few previous decisions, and will naturally emphasize their own, private knowledge. As more and more individuals make their decisions, decision-chains will tend to stabilize on one decision for the rest of the decision-makers.

In Figure 3.5 *d*), we can observe that decision-makers eventually settle on the wrong decision. This is what we define as a *wrong cascade*. In *a*), *b*) and *c*), individuals conform on the correct decision. We note that one should be careful to draw conclusions from single samples. However, there seems to be a tendency that the larger the variation in the simulated competences (red curve), the more varying will the simulated decisions be (blue dots). This is a result confirmed by Falnes (2019) for two similar models, and may be connected to the fact that some individuals will receive particularly strong private signals, and choose to emphasize their private knowledge to a larger extent than others with weaker such signals. It is also reasonable to assume that the higher the mean value of the true competences, the more probable it is that the resulting cascade is correct. A larger fraction of the simulated signals y will be correct, and individuals are also more confident on their private knowledge. With σ chosen such that the correlations between competence estimates and the true competences are rather high, individuals are also likely to think that others' competences are high, and view previous decisions as informative. Figure 3.6 shows the probabilities for choosing $X = 0$ for the simulation in Figure 3.5. The figure visualizes the effect the softmax-weighting has on the final probabilities for choosing a specific value. We remind that $P(X = 0|\mathbf{z_{i-1}}, \mathbf{\tilde{p}^i}, p_i, y_i) + P(X = 1|\mathbf{z_{i-1}}, \mathbf{\tilde{p}^i}, p_i, y_i) = 1$ and similarly $P(Z_i = 0|\mathbf{z_{i-1}}, \mathbf{\tilde{p}^i}, p_i, y_i) + P(Z_i = 1|\mathbf{z_{i-1}}, \mathbf{\tilde{p}^i}, p_i, y_i) = 1$. As the number of individuals grow, the probability of choosing a specific value slowly increases towards 1 and is weighted to be very close to 1 by the softmax-function for this value of τ.

For the parameters $\alpha = 62$ and $\beta = 62$ depicted in plot *b*) of Figure 3.5, we illustrate in Figure 3.7 the effect different values of τ have on the chain of final decisions. In all simulations in the figure, the same seed has been used. This means that the illustrated competences, competence estimates and private signals are the same in all situations. As seen in Figure 3.4, small changes to τ give rather large changes to the degree of randomness in each decision. As we will discuss later, this plot can also give us an idea of how to choose reasonable parameter values. Plot *a*) illustrate a value of τ close to 0, which means that most individuals will make the optimal choice. As mentioned previously, the mean value of the competences and competence estimates are rather high, making individuals consider the observed decisions of previous individuals as informative. Only a few individuals in the beginning of the chain act opposite of the predominant behaviour, before enough decisions is observed and individuals conform on the correct value. The simulation depicted in *b*) illustrates the same situation, but with a higher value for τ. Decisions seem more random, and the chain uses a longer time to converge.

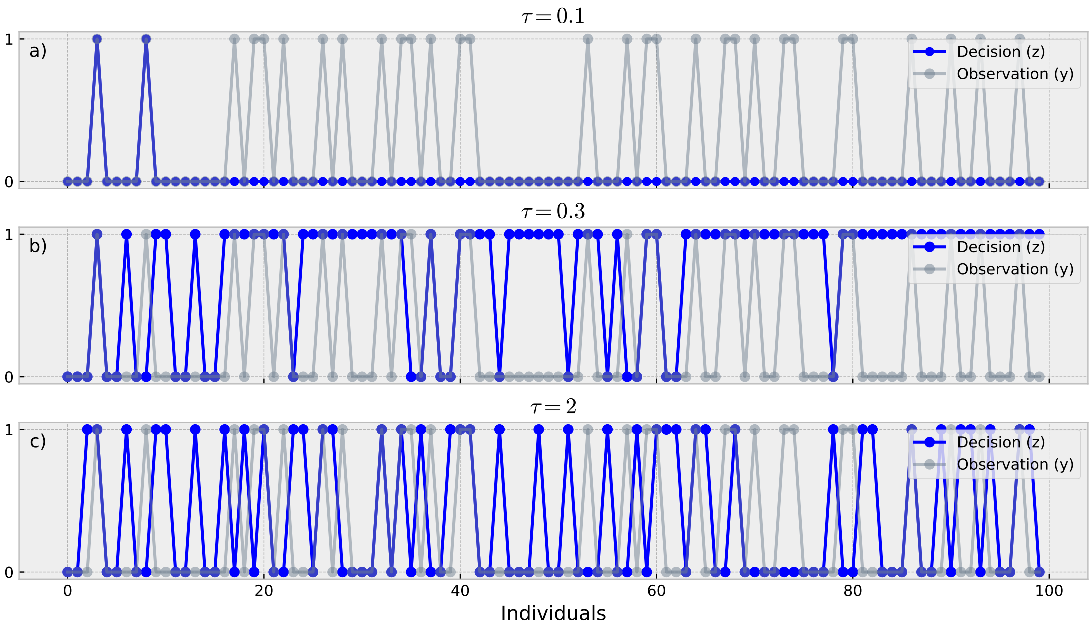

Figure 3.7: Simulated chains of decisions of $n = 100$ individuals for three different choices of the parameter τ. The true value is $x = 0$ in all simulations. In all three plots, the same seed has been used. The other parameters are fixed, and $\alpha = 62$, $\beta = 62$ and $\sigma = 1$.

Since individuals have high estimates of other individuals competences, many individuals in a row will make the same, wrong decision after observing previous individuals. Even when the chain seems to have converged, there will still be single individuals that act differently than the rest. The last plot illustrate a high value of τ. The probabilities of each final decision is weighted to be close to 0.5, which means that most decisions are arbitrary, and the individuals do not seem to conform on one decision.

Chapter 4

Parameter estimation

We have now defined a model for sequential decision making, implemented it numerically and are able to simulate chains of decisions. The parameter choices of α and β govern the overall competence in the population, while σ governs the accuracy of the knowledge individuals have about the competence of the other decision-makers. Finally, the parameter τ reflects the degree of randomness in the individuals' decisions. As seen in the previous chapter, different choices of parameters will result in decision chains with different characteristics. The objective of this chapter is to derive a method for parameter estimation. One way to do this is to use maximum likelihood estimation by for example deriving the likelihood of the system and use numerical optimization with respect to the different parameters of interest. However, we will see that the large number of unknown variables in the system makes this approach computationally infeasible, and we will in stead address the Bayesian approach. In particular, we will make use of the M-H algorithm and give as input simulated decision chains as the observed data. The goal is to investigate whether or not there is enough information in the decision-chains to be able to simulate from the posterior distribution of the parameters.

4.1 The posterior distribution of the system

Based on the model definition, we can formulate the posterior distribution of the system. First, we let θ denote the parameters of interest, $\theta = (\tau, \sigma, \alpha, \beta)$. These are the hyperparameters of the system. The only observed variables in the model are the decisions of individuals $1, \ldots, n$, denoted $\mathbf{z} = (z_1, \ldots, z_n)$. The posterior distribution is given by

$$p(\theta|\mathbf{z}) \propto f(\theta)(\mathbf{z}|\theta). \tag{4.1}$$

In order to evaluate the likelihood $f(\mathbf{z}|\theta)$, we need to account for the unobserved variables in our system. These are the personal competences of each individual, the private signals and the competence estimates of the previous decision-makers, in addition to the true value x. We denote the personal competences of all individuals $\mathbf{p} = (p_1, \ldots, p_n)$ and the private signals $\mathbf{y} = (y_1, \ldots, y_n)$. The competence estimates are denoted $\tilde{\mathbf{p}} =$

$(\tilde{p}_1^2, \dots, \tilde{p}_1^n, \dots, \tilde{p}_{n-1}^n)$. summing up, the collection of unknown variables and parameters in our system is the hyperparameters in θ, in addition to $x, \mathbf{p}, \mathbf{y}, \tilde{\mathbf{p}}$. The analytical expression for the likelihood is given by

$$f(\mathbf{z}|\theta) = \int_{\mathbf{p}} \int_{\tilde{\mathbf{p}}} \prod_{i=1}^{n} \sum_{y_i=0}^{1} f(\mathbf{z}, \mathbf{p}, \tilde{\mathbf{p}}, y_1, \dots, y_n|\theta) d\mathbf{p} d\tilde{\mathbf{p}}.$$

The complex nature of the above expression makes it infeasible to calculate analytically as it consists of integrals of high dimension. In particular, we have n integrals over each competence p_i, and $1/2 \cdot n \cdot (n-1)$ integrals over the competence estimates $\tilde{p}_j^i$, in addition to the n sums over each private signal y_i. This is why we instead turn to the MCMC approach and approximate the full posterior of the variables up to a proportionality constant. We make use of the fact that the posterior in (4.1) is proportional to the joint posterior of the system given the observed data,

$$f(\theta|\mathbf{z}) \propto f(\theta, x, \mathbf{p}, \tilde{\mathbf{p}}, \mathbf{y}|\mathbf{z}),$$

and use the more convenient, latter distribution as our target distribution in the M-H algorithm. By the model formulation, we have that

$$f(\theta, x, \mathbf{p}, \tilde{\mathbf{p}}, \mathbf{y}|\mathbf{z}) \propto f(\theta) f(x) f(\mathbf{p}|\theta) f(\mathbf{y}|p, \theta) f(\tilde{\mathbf{p}}|\mathbf{p}, \theta) f(\mathbf{z}|\mathbf{y}, \mathbf{p}, \theta)$$

where $f(\theta)$ denotes the prior distribution of the parameters of interest, and the other distributions are defined previously. The marginal posterior for a given parameter is found by integrating out the other variables. We use component-wise updates, and write out the posterior distribution in terms of univariate distributions. We get

$$\begin{aligned} &f(\theta, x, \mathbf{p}, \tilde{\mathbf{p}}, \mathbf{y}|\mathbf{z}) \propto \\ &f(\tau) f(\sigma) f(\alpha) f(\beta) f(x) \prod_{i=1}^{n} f(p_i|\alpha, \beta) \prod_{j=1}^{i-1} f(\tilde{p}_j^i|p_i, \sigma) f(y_i|x, p_i, \tilde{p}_j^i) f(z_i|\mathbf{z_{i-1}}, y_i, p_i, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau), \end{aligned} \tag{4.2}$$

where $\mathbf{z_{i-1}} = (z_1, \dots, z_{i-1})$ and $\tilde{\mathbf{p}}^{\mathbf{i}}$ denotes individual i's estimates of the previous individual's competences. We now derive a method to approximate this distribution by random sampling with the M-H algorithm.

4.2 Prior distributions

As mentioned in Chapter 2.1, the choice of prior distributions should reflect the knowledge we have about the parameters of interest prior to observing any data. Ideally, information like this can for example come from domain experts or others that have specific knowledge about the situation under study. Sometimes it may be difficult to obtain such knowledge and we need to use a vague prior distribution for the parameters. One approach is to choose priors so the resulting posterior becomes a known distribution it is trivial to sample from, and one can use the Gibbs sampling procedure, see for example

Gamerman and Lopes (2006, Ch. 5). This is not possible in our case, and there is no obvious choice of prior. Our aim is to investigate whether or not we are able to sample from the posterior distribution of the parameters given the observed decisions. Hence, we consider it convenient to use the flexible gamma distribution for our purpose. The hyperparameters can be adjusted in various manners to resemble different cases of prior knowledge, and thus give us an idea of how well the method works on real data.

In our model, we assume that the parameters in θ are independent, and write

$$f(\theta) = f(\tau)f(\sigma)f(\alpha)f(\beta).$$

Each of the parameters are defined to be positive, which is ensured in the prior distribution when we assume gamma priors. We state that the prior distributions are given by

$$\begin{aligned}
f(\tau) &\propto \tau^{g_\tau - 1} e^{-\tau/h_\tau}, \\
f(\sigma) &\propto \sigma^{g_\sigma - 1} e^{-\sigma/h_\sigma}, \\
f(\alpha) &\propto \alpha^{g_\alpha - 1} e^{-\alpha/h_\alpha}, \\
f(\beta) &\propto \beta^{g_\beta - 1} e^{-\beta/h_\beta},
\end{aligned}$$

where g_{θ_i} and h_{θ_i} are fixed parameters of the gamma distribution, and θ_i denotes element i in θ. We omit the normalizing constants as we are interested in the M-H acceptance probability stated in (2.1), which is a ratio where all constants cancel.

4.3 Proposal distributions

The M-H algorithm requires a proposal distribution $q(\theta^*|\theta^{(t)})$ with the objective of proposing a new state θ^* for the Markov chain, given the current state $\theta^{(t)}$. We will use a single-site M-H algorithm, which means that we will iteratively propose a new state for one and one component in the system given by (4.2). Hence, we will define a univariate proposal distribution for each component. In general, the rate of convergence and the mixing properties of the chain will depend on the choice of proposal distributions. We do not want the proposals to be too far from, nor too close to the current state. A small step size will lead to small changes and a high acceptance rate, but highly correlated samples and a slow exploration of the target, and the chain will use many iterations to converge. Similarly, too large steps will lead to a low acceptance rate.

We will first consider the parameters of main interest, namely the hyperparameters in $\theta = (\tau, \sigma, \alpha, \beta)$. These parameters are all defined to be positive, and each proposal distribution has to ensure that the proposals at all times stay within this domain. In order to model this, we will let v denote a gamma distributed variable with expected value close to 1 and a small variance. By letting the proposed value for parameter θ_i be given by $\theta_i^* = \theta_i^{(t)} \cdot v$, we ensure that all proposals are positive and that the proposed states are sufficiently close to the current state. We here explicitly state the proposal distribution for the parameter τ, but note that the proposal distribution of all parameters in θ will have the same form. If $v \sim \text{Gamma}(a, b)$, then it can be shown that the proposal distribution for $\tau^* = \tau^{(t)} \cdot v$ where the current state $\tau^{(t)}$ is given, is

$$q(\tau^*|\tau^{(t)}) \propto \frac{(\tau^*)^{a-1}}{\tau^{(t)^a}} \cdot e^{-\tau^*/(b \cdot \tau^{(t)})}. \tag{4.3}$$

Hence, $\tau^*|\tau^{(t)} \sim \text{Gamma}(a, b \cdot \tau^{(t)})$.

Now turning to the updates of the other variables in (4.2). The true value x and the private signals $\mathbf{y}$ are discrete and binary, so a good proposal for these variables is to simply propose the opposite of the current state. Hence,

$$q(x^*|x^{(t)}) = 1 \cdot \text{I}(x^* = 1 - x) \quad \text{and} \quad q(y_i^*|y_i^{(t)}) = 1 \cdot \text{I}(y_i^* = 1 - y_i), \quad i = 1, \ldots, n,$$

where $\text{I}(\cdot)$ denotes the indicator function.

When it comes to updates on the competences and the competence estimates, we need to ensure that the proposed values $\mathbf{p}^* \in (0.5, 1)$ and $\tilde{\mathbf{p}}^* \in (0.5, 1)$. For this reason, we will use the same trick as for modelling the generation of the competence estimates in Section 3.2. We let

$$\zeta_i^{(t)} = \text{logit}((p_i^{(t)} - 0.5)/0.5) + e_i,$$

where $e_i \sim \mathcal{N}(0, \nu_p^2)$, meaning e_i follows a zero-mean normal distribution with standard deviation ν_p. Transforming the above quantity back to the interval $(0.5, 1)$, we get the proposals

$$p_i^* = \left(\frac{e^{\zeta_i^{(t)}}}{1 + e^{\zeta_i^{(t)}}} \right) \cdot 0.5 + 0.5.$$

As stated in Section 3.2, $(p_i^* - 0.5)/0.5|p_i^{(t)}$ follows the logitnormal distribution with density given by (3.8) with location parameter $\mu = \text{logit}((p_i^{(t)} - 0.5)/0.5)$ and scale parameter ν_p. We use the same proposal distribution for each of the components of the competence estimates $\tilde{\mathbf{p}}$, where we denote the scale parameter of the proposal distribution $\nu_{\tilde{p}}$.

4.4 The M-H acceptance probability

With the prior and proposal densities defined, we are now ready to define the Metropolis-Hastings acceptance ratio given in (2.1). As previously mentioned, the full posterior distribution is our target distribution. We define the ratio from the M-H acceptance probability given in (2.1) as

$$r(\phi^*|\phi^{(t)}) := \frac{f(\phi^*|\mathbf{z})}{f(\phi^{(t)}|\mathbf{z})} \cdot \frac{q(\phi^{(t)}|\phi^*)}{q(\phi^*|\phi^{(t)})},$$

where $\phi = (\tau, \sigma, \alpha, \beta, \mathbf{p}, \tilde{\mathbf{p}}, \mathbf{y}, x)$. Since we will only update one parameter at the time from ϕ, the factors not being updated will cancel out in the above fraction. In order to make computations as fast and stable as possible, we simplify the expression for the acceptance probability as much as possible for each of the parameters. For each of the hyperparameters in $\theta = (\tau, \sigma, \alpha, \beta)$, cancelling out constant factors gives

$$
\begin{aligned}
r_\tau(\tau^*|\tau^{(t)}) &= \frac{f(\tau^*)}{f(\tau^{(t)})} \cdot \frac{\prod_{i=1}^n f(z_i|\mathbf{z_i}, y_i, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau^*)}{\prod_{i=1}^n f(z_i|\mathbf{z_i}, y_i, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau^{(t)})} \cdot \frac{q(\tau^{(t)}|\tau^*)}{q(\tau^*|\tau^{(t)})}, \\
r_\sigma(\sigma^*|\sigma^{(t)}) &= \frac{f(\sigma^*)}{f(\sigma^{(t)})} \cdot \frac{\prod_{i=1}^n \prod_{j=1}^{i-1} f(\tilde{p}_j^i|p_i, \sigma^*)}{\prod_{i=1}^n \prod_{j=1}^{i-1} f(\tilde{p}_j^i|p_i, \sigma^{(t)})} \cdot \frac{q(\sigma^{(t)}|\sigma^*)}{q(\sigma^*|\sigma^{(t)})}, \\
r_\alpha(\alpha^*|\alpha^{(t)}) &= \frac{f(\alpha^*)}{f(\alpha^{(t)})} \cdot \frac{\prod_{i=1}^n f(p_i|\alpha^*, \beta)}{\prod_{i=1}^n f(p_i|\alpha^{(t)}, \beta)} \cdot \frac{q(\alpha^{(t)}|\alpha^*)}{q(\alpha^*|\alpha^{(t)})}, \\
r_\beta(\beta^*|\beta^{(t)}) &= \frac{f(\beta^*)}{f(\beta^{(t)})} \cdot \frac{\prod_{i=1}^n f(p_i|\alpha, \beta^*)}{\prod_{i=1}^n f(p_i|\alpha, \beta^{(t)})} \cdot \frac{q(\beta^{(t)}|\beta^*)}{q(\beta^*|\beta^{(t)})},
\end{aligned}
$$

where the proposal distributions $q(\theta_i^*|\theta_i^{(t)})$ are given by (4.3). We will also update the other parameters component wise. From (4.2), we can see that for each of the personal competences p_i, the fraction in the acceptance probability will be

$$
r_{p_i}(p_i^*|p^{(t)}) = \frac{f(p_i^*|\alpha, \beta) f(z_i|\mathbf{z_{i-1}}, y_i, p_i^*, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau) \prod_{j=1}^{i-1} f(\tilde{p}_j^i|p_i^*, \sigma)}{f(p^{(t)}|\alpha, \beta) f(z_i|\mathbf{z_{i-1}}, y_i, p_i^{(t)}, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau) \prod_{j=1}^{i-1} f(\tilde{p}_j^i|p_i^{(t)}, \sigma)} \cdot \frac{q(p_i^{(t)}|p_i^*)}{q(p_i^*|p_i^{(t)})},
$$

where j denotes the index of the previous decision-makers. Similarly, the acceptance ratio for the competence estimates of the i'th individual are given by

$$
r_{\tilde{p}_j^i} = \frac{f(\tilde{p}_j^{i*}|p_i, \sigma) f(y_i|x, \tilde{p}_j^{i^*}) f(z_i|\mathbf{z_{i-1}}, p_i, y_i, \tilde{p}_j^{i^*}, \tilde{p}_{-j}^i)}{f(\tilde{p}_j^{i(t)}|p_i, \sigma) f(y_i|x, \tilde{p}_j^{i(t)}) f(z_i|\mathbf{z_{i-1}}, p_i, y_i, \tilde{p}_j^{i(t)}, \tilde{p}_{-j}^i)},
$$

where $\tilde{p}_{-j}^i$ denotes the competence estimates not being updated, i. e. ($\tilde{p}_{-j}^i = \tilde{p}_1^i, \ldots, \tilde{p}_{j-1}^i, \tilde{p}_{j+1}^i, \ldots, \tilde{p}_{i-1}^i$). The acceptance ratios for the private signals y_i is given by

$$
r_{y_i}(y_i^*|y_i) = \frac{f(z_i|\mathbf{z_{i-1}}, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau, y_i^*) f(y_i^*|x, p_i) \prod_{k=i+1}^n f(y_i^*|x, \tilde{p}_i^k)}{f(z_i|\mathbf{z_{i-1}}, \tilde{\mathbf{p}}^{\mathbf{i}}, \tau, y_i^{(t)}) f(y_i^{(t)}|x, p_i) \prod_{k=i+1}^n f(y_i^{(t)}|x, \tilde{p}_i^k)} \cdot \frac{q(y_i^{(t)}|y_i^*)}{q(y_i^*|y_i^{(t)})}.
$$

Finally, the acceptance ratio for x is given by

$$
r_x(x^*|x) = \frac{f(x^*) \prod_{i=1}^n \prod_{j=1}^{i-1} f(y_i|x^*, p_i) f(y_i|x^*, \tilde{p}_j^i)}{f(x^{(t)}) \prod_{i=1}^n \prod_{j=1}^{i-1} f(y_i|f(x^{(t)}, p_i) f(y_i|x^{(t)}, \tilde{p}_j^i)} \cdot \frac{q(x^{(t)}|x^*)}{q(x^*|x^{(t)})}.
$$

All the necessary framework for the M-H algorithm is now derived, and the method is ready to be implemented numerically.

4.5 Notes on the implementation

The derived method for parameter estimation is implemented in the programming language Python. The advantage of using Python is that it is easy to read and implement.

However, Python is an interpreted language, and will have a significant amount of overhead, compared to a compiled language, e.g. C or Fortran. The objective of this thesis is of an exploratory fashion; we are interested in some key properties of the method and our system. This includes assessing whether or not it is sensible to later make us of the derived method for inference on data from for example a proper experiment or other forms of 'real-world data'. For our use, it is reasonable to use Python, but implementing the code in a lower-level, compiled language would probably improve the run-time significantly.

As previously mentioned, we use component-wise updates. This means that we for each iteration in the M-H algorithm update one and one parameter in θ. Hence, in one iteration we visit α, β, σ and τ once. It is reasonable to assume that the more examples of decision-chains simulated with the same parameters, the more information is present in the system when estimating parameters. For this reason, we assume that it is advantageous to use multiple decision-chains when sampling the posterior of a set of parameters. This means that one iteration corresponds to updating the hyperparameters by using *one* of the input decision-chains. In the next iteration, we use the second input decision-chain, and so on. The exact number of chains one should use is one of the properties we will study in the next chapter.

In addition to the hyperparameters, states from the posterior distribution of the unknown variables $\mathbf{y}$, $\mathbf{p}$, $\tilde{\mathbf{p}}$ and x will need to be updated. These variables are specific to each chain of decisions, and we consider them less 'interesting' than the model parameters in θ. Additionally, updating $\mathbf{y}$, $\mathbf{p}$ and $\tilde{\mathbf{p}}$ involves evaluating the likelihood $f(z_i|\mathbf{z}_{i-1}, \tilde{\mathbf{p}}^{\mathbf{i}}, y_i, \tau)$ which is computationally expensive. This is why we choose to update these variables less frequently. Every 4'th time we have been through all decision-chains, we also update the parameters corresponding to each decision-chain. In order to make computations as stable as possible, calculations of the M-H acceptance ratios are done on log-scale. Many of the factors in the probability expressions will often produce very small numbers, and by doing calculations on log-scale we might avoid excessive rounding errors.

Chapter 5

Numerical experiments

The parameter estimation method derived in the previous chapter is implemented. In this chapter, we present a simulation study and discuss the results. In particular, we simulate chains of decisions, and use these decisions as input when we estimate parameters. The aim is to investigate some key properties about the derived method - is there enough information in the system based on the observed decisions alone? How many chains of decisions do we need as input in order to have stable computations and convergence, and to what degree are we able estimate the correct parameters? We compare the performance of the method with decision-chains of different characteristics, and test the sensitivity for different prior distributions. This analysis will be a useful tool if one in the future wish to test the method on real data.

5.1 Generating decisions: Two cases

Throughout the simulation studies conducted in this chapter, we will consider two different cases of simulated decisions. Specifically, we introduce two different sets of the hyperparameters $\tau, \sigma, \alpha, \beta$, and from each set simulate decision-chains. These parameters should be chosen such that the resulting decisions resemble realistic situations. Moreover, each of the cases should to some extent exhibit different characteristics. The purpose of this is to analyse the difference in the performance of the method for two distinct situations. It is natural to assume that the information in the decisions lies in the variability in the decisions before the cascades occur, as cascades have proved to occur relatively fast in most of the simulations. Therefore, we would like one of the cases to resemble a situation where many individuals are uncertain and decisions are varying. This case will be compared to a case where cascades occur during the first few individuals, and there presumably is less information in the simulated decision-chains.

The chosen parameter-values for the two different cases are presented in Table 5.1. For each of these two cases, we simulate 10 chains of decisions, each of length $n = 50$ individuals. Figure 5.1 shows the resulting simulated decisions from each case. In all simulations, we have let the true value $X = 0$. In the first case, parameters α and β are chosen such that the overall competence in the population is rather low (close to

	τ	σ	α	β
Case 1	0.05	1	2	50
Case 2	0.1	2	15	20

Table 5.1: Parameter choices for each of the two cases used in the simulation study.

0.5) for all individuals. Further, we use $\tau = 0.05$, which means that there is a very high chance to make the optimal decision. In other words, most individuals will act rationally. Additionally, we let $\sigma = 1$ for the first case. This corresponds to a correlation of 0.47, between the true competences and the competence estimates. Hence, competence estimates are moderately close to the true competences. The simulated chains from Case 1 is depicted in the upper plot in Figure 5.1. About half of the simulated decision-chains ends up as wrong cascades. This is probably due to the low competences in the population.

For Case 2, α and β are chosen such that competences are higher, with a mean value of 0.71. The variance of the simulated competences is also higher compared to Case 1. This means that to a greater extent than in Case 1, some individuals will have a higher competence, while others will have a lower competence than the mean value. The parameter $\tau = 0.1$, which is higher than in Case 1, but reasonably low, meaning most individuals make the optimal decision. Nevertheless, each individual will have a larger probability of not making the optimal decision in Case 2 compared to Case 1. Also, $\sigma = 2$, which makes the correlation between competences and competence estimates about 0.26. In total and compared to Case 1, individuals in Case 2 will have an overall higher competence, but are more prone to making random decisions and has less information about other individuals' true competence. However, from the simulated decisions seen in the lower plot in Figure 5.1, decisions from Case 2 are to a greater degree correct, but less variable than in Case 1. It is therefore reasonable to believe that there will be more 'difficult' to estimate parameters form the situation resembled by Case 2.

5.2 Experimental setup

In this section we present a plan for the numerical experiments we conduct in this report. We begin by assessing the convergence properties of the method, before we move on to study the resulting posterior distributions for each of the cases.

For all experiments, we have chosen a set of parameters for the prior distributions and a set of tuning parameters for the proposal distributions. These values are summarized in Table 5.2. In the choice of tuning parameters for the prior distributions of α and β, we have assumed it is reasonable to consider it known that α is rather low (somewhere between 0 and 20) and β is higher than α (between 20 and 100). With regard to the decision to be taken, this corresponds to the situation where the overall competence in the population is low. Based on values for the 5 and 95 percentiles for the prior distribution of α and β, we will expect the mean value of the competences to lie in the interval $(0.51, 0.7)$ with expected value $E[p_i] = 0.56$. For Case 2, we assume it is reasonable to assume that

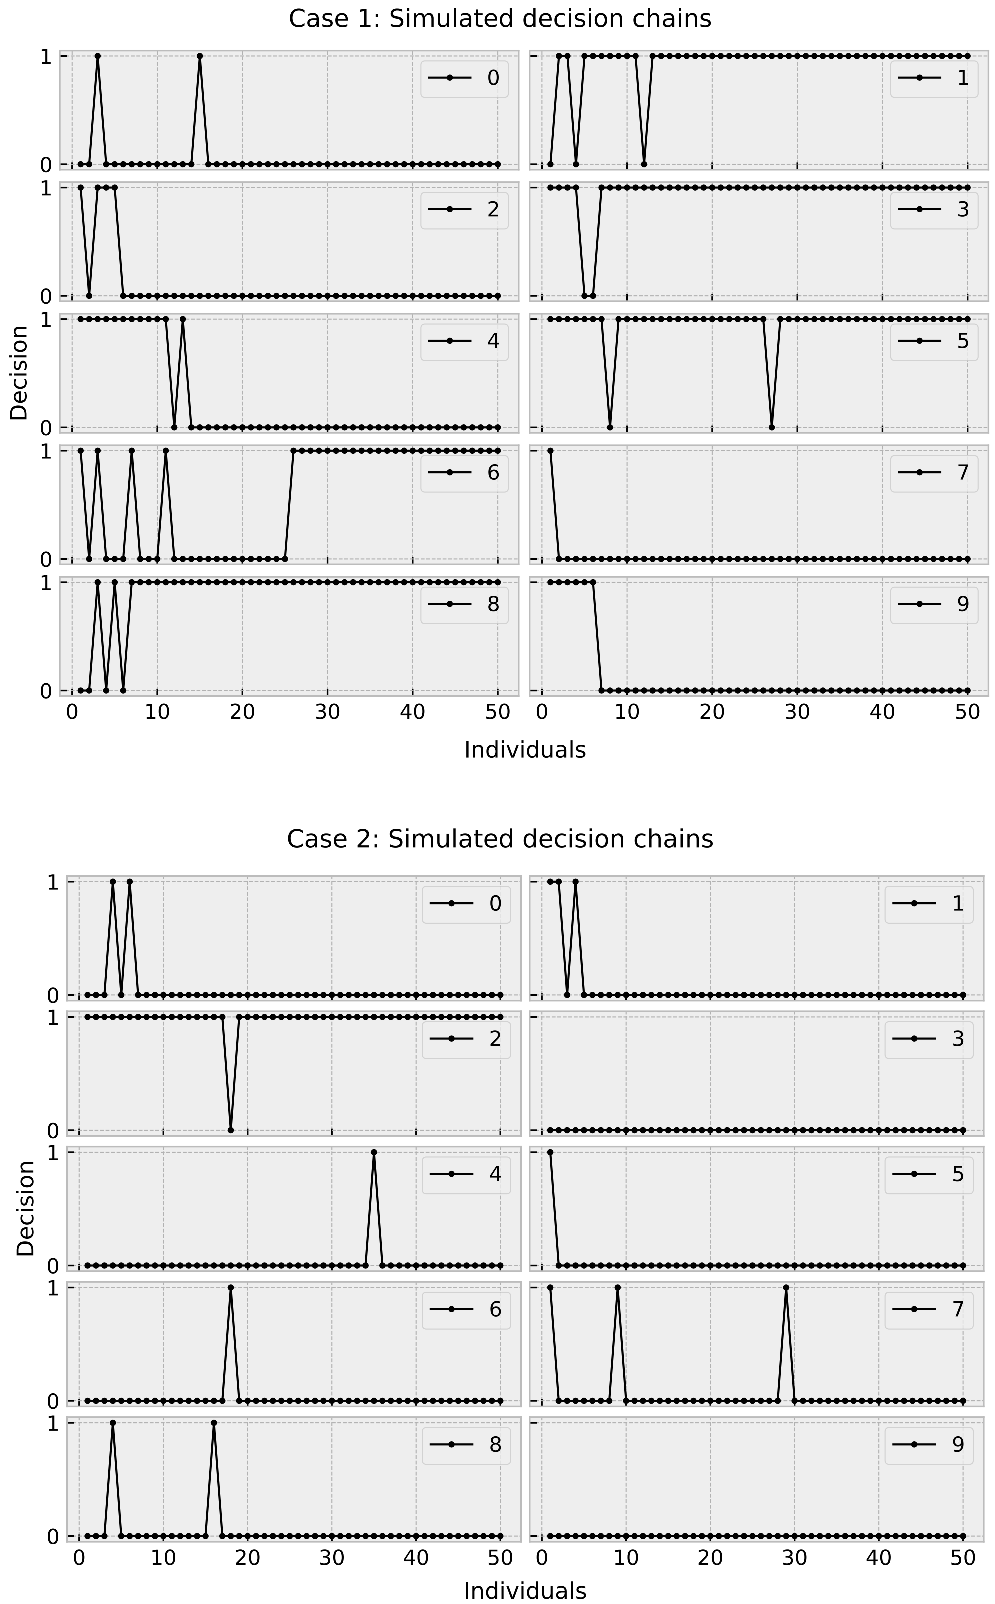

Figure 5.1: The 10 decision chains simulated from Case 1 and 2.

	Case 1		**Case 2**		**Case 2***	
	Prior $(a_{\theta_i}, b_{\theta_i})$	Proposal $(g_{\theta_i}, h_{\theta_i})$	Prior $(a_{\theta_i}, b_{\theta_i})$	Proposal $(g_{\theta_i}, h_{\theta_i})$	Prior $(a_{\theta_i}, b_{\theta_i})$	Proposal $(g_{\theta_i}, h_{\theta_i})$
τ	$(1.1, 0.5)$	$(10, 0.1)$	$(1.1, 0.5)$	$(10, 0.1)$	$(1.1, 0.5)$	$(10, 0.1)$
σ	$(5, 0.5)$	$(10, 0.1)$	$(5, 0.5)$	$(10, 0.1)$	$(5, 0.5)$	$(10, 0.1)$
α	$(4, 2)$	$(7, 0.15)$	$(4, 7)$	$(7, 0.15)$	$(9, 2)$	$(7, 0.15)$
β	$(11, 5)$	$(7, 0.15)$	$(4, 7)$	$(7, 0.15)$	$(9, 2)$	$(7, 0.15)$
p	–	$\nu_p = 0.4$	–	$\nu_p = 0.4$	–	$\nu_p = 0.4$
$\tilde{p}$	–	$\nu_{\tilde{p}} = 1$	–	$\nu_{\tilde{p}} = 1$	–	$\nu_{\tilde{p}} = 1$

Table 5.2: Parameter values for the prior and proposal distributions presented in Chapters 4.2 and 4.3 for the hyperparameters, p and $\tilde{p}$ in the numerical experiments. The initial value for $\tilde{p}$ is generated randomly by adding noise to the initial value of p as described in Chapter 4.3.

we know that α and β are approximately equal, and we have chosen the same prior for the two parameters. Based on the 5 and 95 percentiles for the priors, we expect that the overall competences lie within $(0.57, 0.92)$ with expected value 0.75, before observing any data. This is a larger interval than for Case 1, meaning the prior distributions for α and β in Case 2 are less informative than for Case 1. Lack of prior knowledge about the parameters is resembled by choosing quite vague prior distributions. In this way, the resulting posteriors are determined mainly by the observed data. However, if there is little information in the system, results will be affected by changes in the prior distribution. Because of this, we also introduce a third case, Case 2*. This case is similar to Case 2, except a change in the parameters of the prior distributions of α and β. We are interested in investigating the effect of adjusting the prior distributions of these parameters such that the variances are reasonably decreased. In particular, based on the 5 and 95 percentiles of the prior distribution of α and β, we go from the situation where we assume the competence in the population lies in the very broad interval (0.57, 0.92) to (0.62, 0.88). This corresponds to a more specific knowledge that most competences are neither very high nor very low.

When it comes to the choices of parameters for the prior distributions of τ and σ, we have considered the mathematical properties of the parameters based on Figures 3.2 and 3.4. It is considered reasonable to assume that τ should be less than 1, which for the decision in Figure 3.4 corresponds to higher than 60% probability of making the optimal choice. Similarly, it is considered reasonable to assume that σ should be between 0 and 6, which corresponds to a correlation between 1 and 0.1.

No systematic experiments for the tuning parameters for the proposal distributions are performed. The choice of these parameters are hence somewhat arbitrary, but are based on some preliminary runs with focus on monitoring the acceptance rates for the sampled parameters. Recall that acceptance rates can give us an indication of the mixing of the resulting Markov chains, and that acceptance rates should not be either too high or too low.

	$\tau^{(0)}$	$\sigma^{(0)}$	$\alpha^{(0)}$	$\beta^{(0)}$	$\mathbf{p}^{(0)}$
Run 1	10	10	10	10	0.9
Run 2	50	50	150	150	0.75
Run 3	1	0.01	50	1	0.51

Table 5.3: Initial values for each of the runs in Experiment 1. For each case, the hyperparameters summarized in Table 5.2 are used. The three different runs are performed for Case 1, Case 2 and Case 2*.

5.2.1 Experiment 1: Convergence properties

We begin by properly investigate the convergence and mixing of the sampled chains from the derived parameter estimation procedure. When doing inferences on output from MCMC-methods, we assume that the samples are from the target distribution. It is therefore crucial that the chain in fact has converged sufficiently. As mentioned in Chapter 2.2.1, one way to assess the convergence of an MCMC-method is to run multiple chains with different initial values. After the burn-in period, the chain should be independent of the initial value, and the different chains should behave similarly. Hence, this analysis will also give insight about the number of iterations to choose for the burn-in period of each of the sampled chains. We stress that since we are interested in estimating the full posterior of the system, we need to choose this number at an iteration where all of the parameters have converged.

In this experiment, we use all 10 decision-chains as input to the algorithm, as it is assumed that this gives the most stable results. We further assume that the last 20 decisions of each chain of decisions provide very little information as the decision-chains in most cases have cascaded at an early individual. Consequently, we run this analysis for decision-chains of length $n = 30$ individuals as input, both for Case 1, Case 2 and Case 2*, using the prior and proposal distributions presented in Table 5.2. We choose some sets of different and extreme initial values for each of the parameters, which is presented in Table 5.3. In this experiment we iterate through each input chain a total of 40000 times. Since we use 10 chains, this corresponds to 400000 updates of the hyperparameters and 10000 updates of the chain-specific parameters.

5.2.2 Experiment 2: Testing properties of the input decision-chains

Based on some preliminary runs of the implemented method it has been found that the method is not sufficiently stable with only one decision chain as input. We are therefore interested in quantifying the improvement of including multiple chains. Another aspect to consider is the payoff between including more chains and the increased usage of cpu-time. Theoretically, one can assume that if we had an infinite amount of cpu-time and data, for instance in terms of number of decision-chains, the method should converge perfectly. This is obviously not possible to carry out in real life. However, as observed in Figure 5.1, we see that after $n \approx 30$ individuals, the simulated decisions are mostly stable for both cases. Decisions will vary most during the first few decisions, before cascades occur. It is

	τ	σ	α	β	$\mathbf{p}$
Initial values	10	10	10	10	0.9

Table 5.4: The initial values used in Experiment 2.

	Case 1		**Case 2**	
	$n = 30$	$n = 50$	$n = 30$	$n = 50$
Run 1	1 chain	1 chain	1 chain	1 chain
Run 2	4 chains	4 chains	4 chains	4 chains
Run 3	10 chains	10 chains	10 chains	10 chains

Table 5.5: Number of input decision-chains for the runs in Experiment 2. For each case, the same set of prior distributions, proposal distributions and initial values is used.

therefore reasonable to assume that most information in the system is contained in the first decisions, and we will investigate if the performance of the method will be affected by using shorter chains, where the last sequences of mostly equal decisions are removed. If it turns out that the last parts of the chains are insignificant for the estimation of the posterior distribution of the system, we can save ourselves for some unnecessary usage of cpu-time.

In this experiment, we will for each of the two cases presented in Section 5.1 choose a set of initial values, and use the parameters for the prior and the proposal distributions listed in Table 5.2 for the M-H algorithm. The initial values used in this experiment are listed i Table 5.4. These are chosen to be reasonably close to the true value for both cases. We run the parameter estimation procedure using 1, 4 and 10 decision-chains as input for each case. For each of these, we will do one run with decision-chains of length $n = 30$ and one run with chains of length $n = 50$ individuals. Since we use the same initial values and the same tuning parameters for each run, these experiments will give us an idea of the number of decision-chains to include, in addition to the number of decisions in each chain to include in order to obtain stable results. The second experiment is summarized in Table 5.5.

5.2.3 Experiment 3: Properties of the sampled posterior distribution

After having established convergence of the MCMC method, in addition to having investigating stability of the method for different variations of the input data, it is natural to take a closer look at the resulting samples of the posterior distribution of the system. The aim of the third experiment is to assess the accuracy of the resulting estimates from some of the runs performed in Experiment 1. We are also interested in the differences in the results for the different cases.

The main objective of the numerical experiments in this chapter is to assess whether or not there is enough information in the observed chains of decisions alone to give reasonable parameter estimates, as the large amount of unobserved variables makes it reasonable to suspect the contrary. It is natural to assume that the accuracy of the resulting samples of the posterior distributions from our MCMC-method improves with more information available. Therefore, it is interesting to compare the resulting posterior distributions from the runs described above to the results from the situation where the private signals of each individual, $y_1, \ldots, y_n$, are assumed to be observed additionally to the decisions. For this reason, we will repeat the runs from Experiment 1 for Case 1, Case 2 and Case 2*, but with the change that the private signals are assumed to be known. The resulting posterior distributions will be compared to the corresponding runs of Experiment 1, where the private signals were unobserved.

5.3 Results and discussion

In this section, we present the results of the experiments described above.

5.3.1 Experiment 1

We begin by presenting the results from Experiment 1, where we are interested in assessing the convergence of the method. Figure 5.2 shows the trace-plots for the first iterations of the first experiment for Case 1. We have zoomed in on the first 150000 iterations, since we in this particular analysis are interested in the convergence of the chain. We observe that the method seems to converge towards the same steady-state for all of the parameters for the different starting values. The parameters τ and σ converges relatively fast in all of the runs. We are however interested in the full posterior, and the system has not converged until all parameters seem to reach some equilibrium. The parameters α and β seem to converge more slowly. The bottom row in the figure shows the trace plots for the competences of two arbitrary individuals, number 5 and 15. We use rather infrequent updates of these parameters, which are specific to each of the input chains. The figure visualizes the competences corresponding to the first out of the ten chains. In order to find the number of iterations before the whole system has reached a steady-state, we need to assess all the sampled competences, in addition to the other chain specific variables $y, \tilde{p}$ and x for all chains. It seems to take about 2000-2500 iterations before all parameters have converged. Since we have used 10 chains as input, this corresponds to 80000-100000 iterations of the hyperparameters before convergence, which is about 20%-25% of the whole chain. Overall, the sampled chains seem to behave equally after the burn-in period, which is an indicator of sufficient convergence for the chain.

Turning to the mixing properties of the chain, there does not seem to be any special trends in the trace plots for the hyperparameters in Figure 5.2. The trace plots for the competences p shows that after convergence, the chain often stays in the same states for many iterations, and only very small steps are accepted. This cause bad exploration of the target distribution. The acceptance rates of the parameters are shown in Figure 5.3. These are overall quite high, but are particularly high for σ and τ. This may contribute to a slow exploration, and is an indication of high autocorrelation. This is not ideal. By construction of MCMC-samplers, two successively drawn values are dependent. However, when the correlation between sample $\theta_i^{(t)}$ and $\theta_i^{(t+k)}$ decays very slowly as k increases,

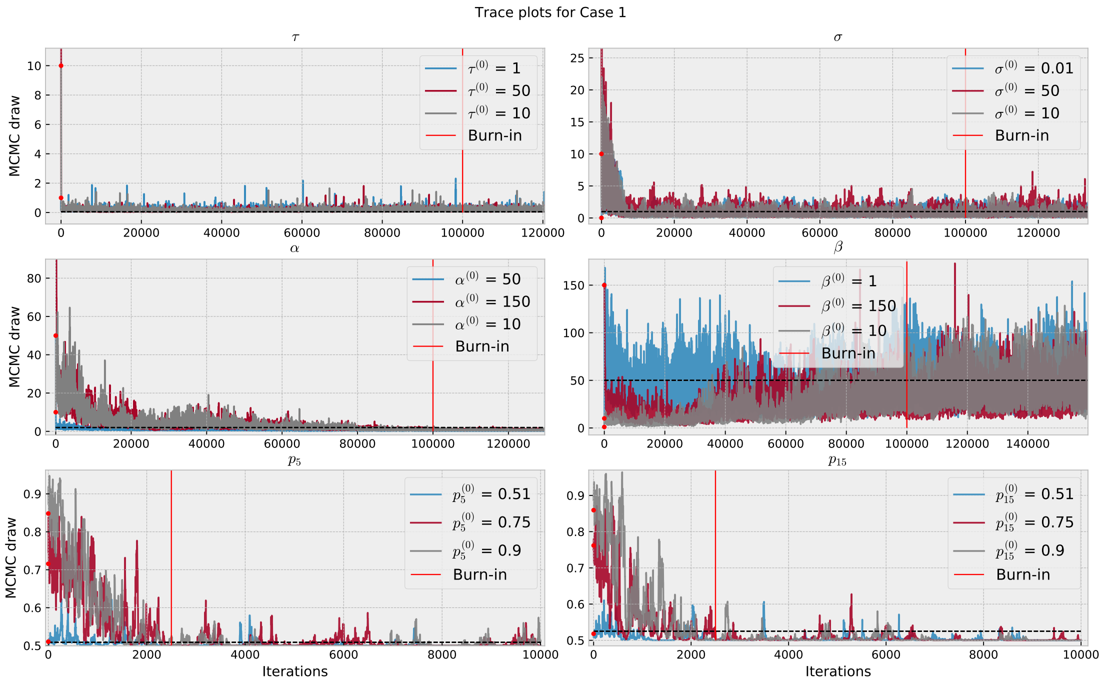

Figure 5.2: Trace plots for the first iterations of Experiment 1 for Case 1. Note that we have zoomed into the first few iterations to visualize the convergence properties of the chains. Initial values are indicated in the labels, and for some of the parameters these are outside of the plot.

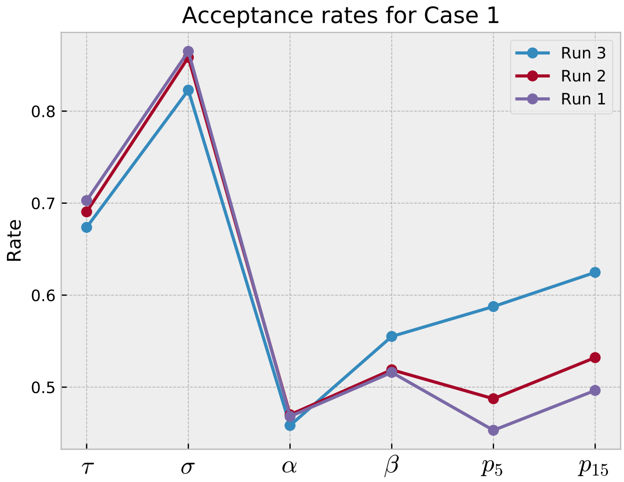

Figure 5.3: Acceptance rates for the runs in Experiment 1 for case 1.

we need many iterations of the algorithm to obtain a representative sample of the target distribution. For this reason, one could consider to adjust tuning parameters in order to make the method propose larger steps.

The first experiment discovered that the chains for different initial values in Case 2 did not seem to converge towards the same posterior distribution for the chosen number of iterations. The trace plots from this experiment is depicted in Figure 5.4. Estimates for τ and σ, depicted in the upper row, seems reasonable. For α and β, it is clear that the different chains are dependent on the different starting values. Since the competences are governed by the parameters α and β, the tendency becomes particularly evident in the trace plots for p_5 and p_{15} plotted in the lowest row, where we see that the chains quickly becomes 'glued' to either 1 or 0.5, and stays in the same states for many consecutive iterations. The trace plots for Case 2* is shown in Figure 5.5. Compared to the trace plots from Case 2, the figure shows that some properties of the sampled chains improve by assuming more prior knowledge about the parameters α and β. The mixing for the trace plots of the personal competences seems to be better for Case 2*. However, there is still apparent that the sampled chains are dependent on the different starting values. This is particularly evident for Run 1, depicted as the grey curves, whereas the red and blue curves seem to behave somewhat equally. The true values for α and β also seem to be within the variation range of the sampled chains. It is possible that the starting values for Run 1 is too 'far away' from the limiting distribution. In combination with a slow convergence rate, this may cause the chain to require many more iterations before sufficient convergence is reached.

Summing up, the runs for Case 1 and Case 2* seem to have essentially converged. However, convergence is slow. For certain poor initial values one will need to run the

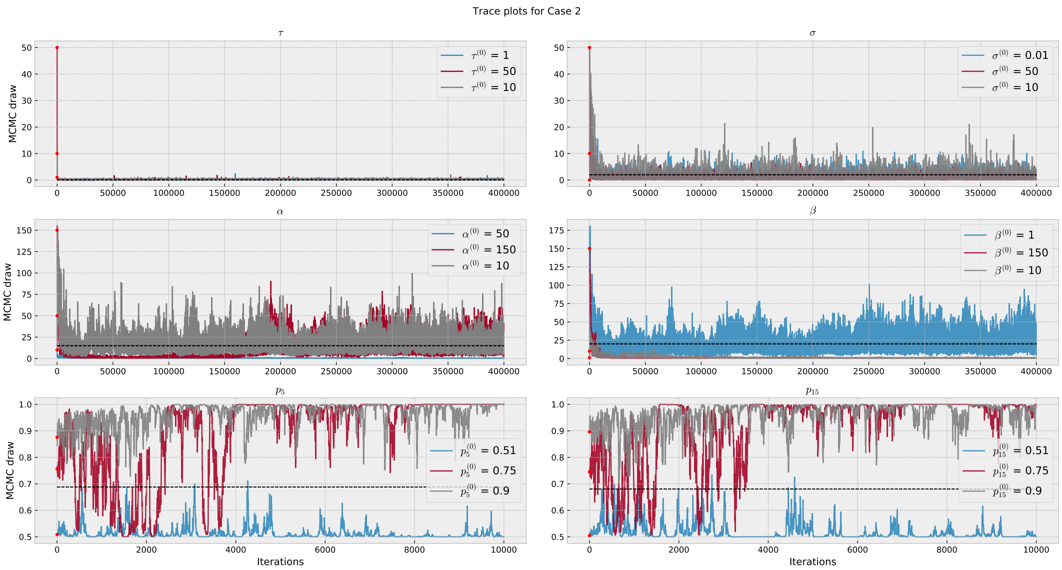

Figure 5.4: Trace plots for Case 2 in Experiment 1.

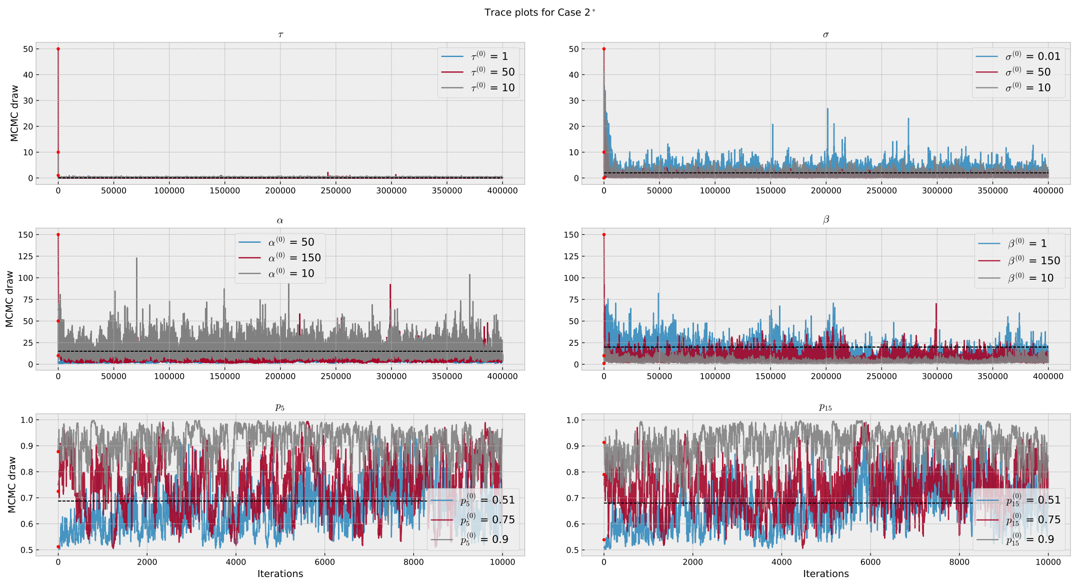

Figure 5.5: Trace plots for Case 2* in Experiment 1.

algorithm for many more iterations than performed in this experiment in order to obtain representative samples of the limiting distribution. Some modifications of the implemented algorithm may improve the rate of convergence. As discussed in Chapter 2.2.1, the choice of proposal distribution affects the convergence, and changing these by for example update several parameters at the time may improve the algorithm.

5.3.2 Experiment 2

We now present some results form the second experiment. As we conducted many rather similar runs, we do not present all of them here. The overall aim of this experiment is to investigate the stability of the method with regard to different variations of the observed decisions used as input, and we present and summarize the key findings. Since the overall tendencies was the same for both of the cases presented in Chapter 5.1, we present mostly results for Case 1, and in this section only briefly discuss Case 2.

The upper and lower plots in Figure 5.6 show the trace plots from Case 1 for decision-chains of length $n = 30$ and $n = 50$, respectively. In both figures, there is used 10 input decision-chains. The corresponding histograms after discarding the burn-in period is shown in the upper and lower plots in Figure 5.7. For $n = 30$, the true parameter values seems to mostly lie within the variation range of the sampled posteriors. However, the true parameter value of τ is far out in the tail. Compared to the lower plot, there seems to be little difference. The parameters τ, σ and β seem to have equally good mixing. Convergence happens approximately equally fast for the situations depicted in the figures, though possibly a bit later for the longer decision-chains in the four lower plots. The variance of the sampled τ is somewhat smaller when using $n = 50$, and the samples from the posterior of α seem to converge towards 0. This will cause the sampled competences to converge towards 0.5 for $n = 50$. This may, however, be reasonable, as the true distribution for p has its spike close to 0.5. Based on these results, there seem to be no reason to include the last decisions of chains where cascades happen early.

Much of the above-mentioned also holds for Case 2. The runs for this case gave reasonable estimates for σ and τ, and there was negligible differences between the runs with decision-chains of length $n = 30$ and $n = 50$. However, for this case, the estimated posterior samples for α and β were not satisfying. The resulting posterior distributions was different for the runs with chains of $n = 30$ and $n = 50$, but both gave estimates unreasonably far away from the true parameter values. This leaves us with the conclusion that there seems to be little difference between chains of length $n = 30$ and $n = 50$. It is unnecessary to include long parts of decision chains that has already cascaded. This will only increase the cpu-time without substantially contributing to the resulting estimates.

We now focus on the number of decision-chains to use as input in the method. Figure 5.8 shows the trace plots of the parameters σ and p_{15}. The latter denotes the sampled posterior competence of the arbitrarily chosen individual number 15 for the first decision-chain used as input, for a differing number of input decision-chains, each of length $n = 50$. It is clear from the plots that the method is unstable for one and four input decision-chains, as the sampled values for σ suddenly spikes. The same tendency is observed for the other runs in this experiment. The rightmost column shows the sampled values

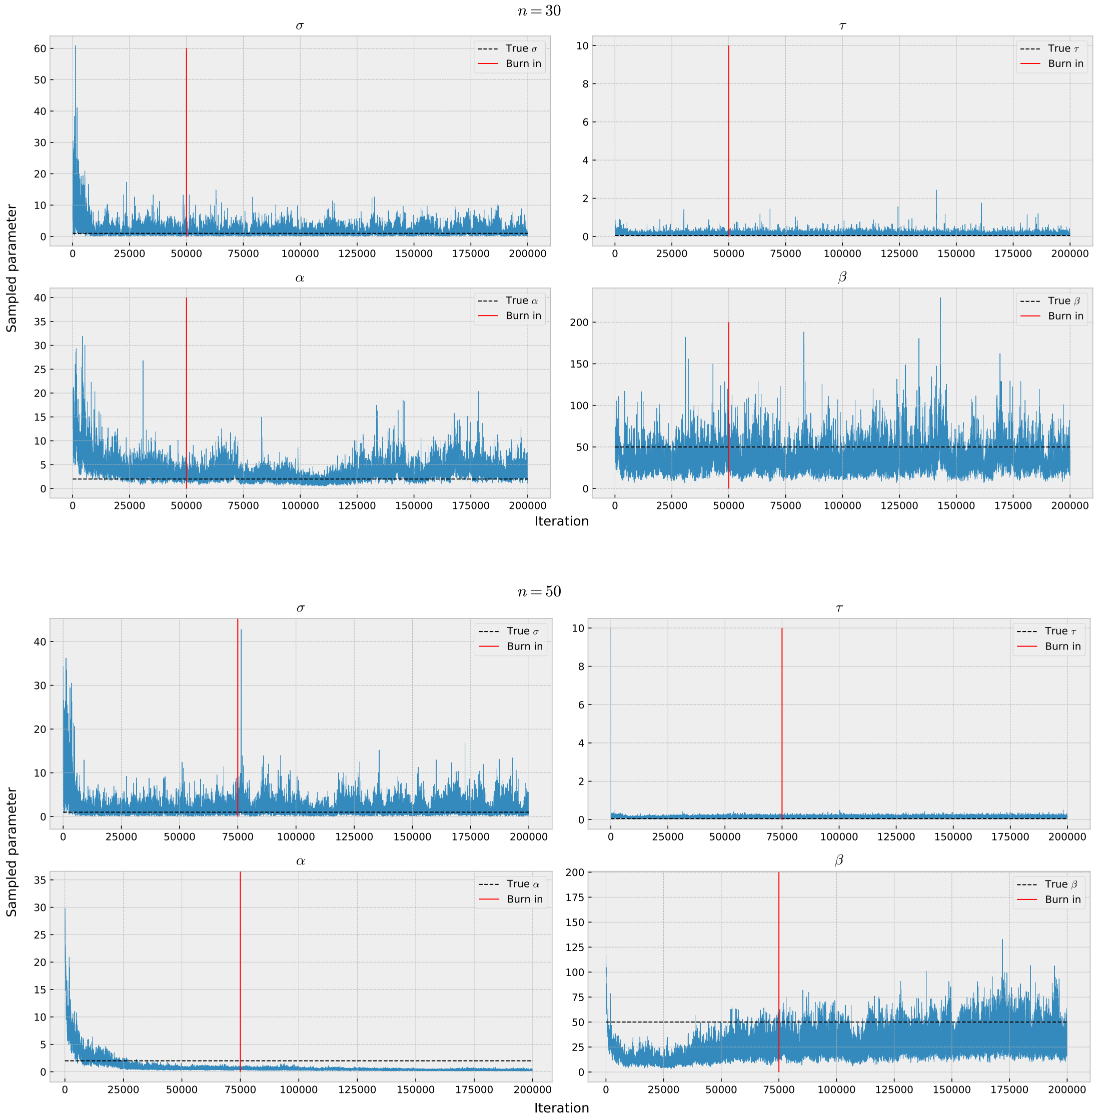

Figure 5.6: Trace plots for the hyperparameters in Case 1 for decision-chains consisting of $n = 30$ (upper plot) and $n = 50$ individuals (lower plot). 10 decision-chains are used as input.

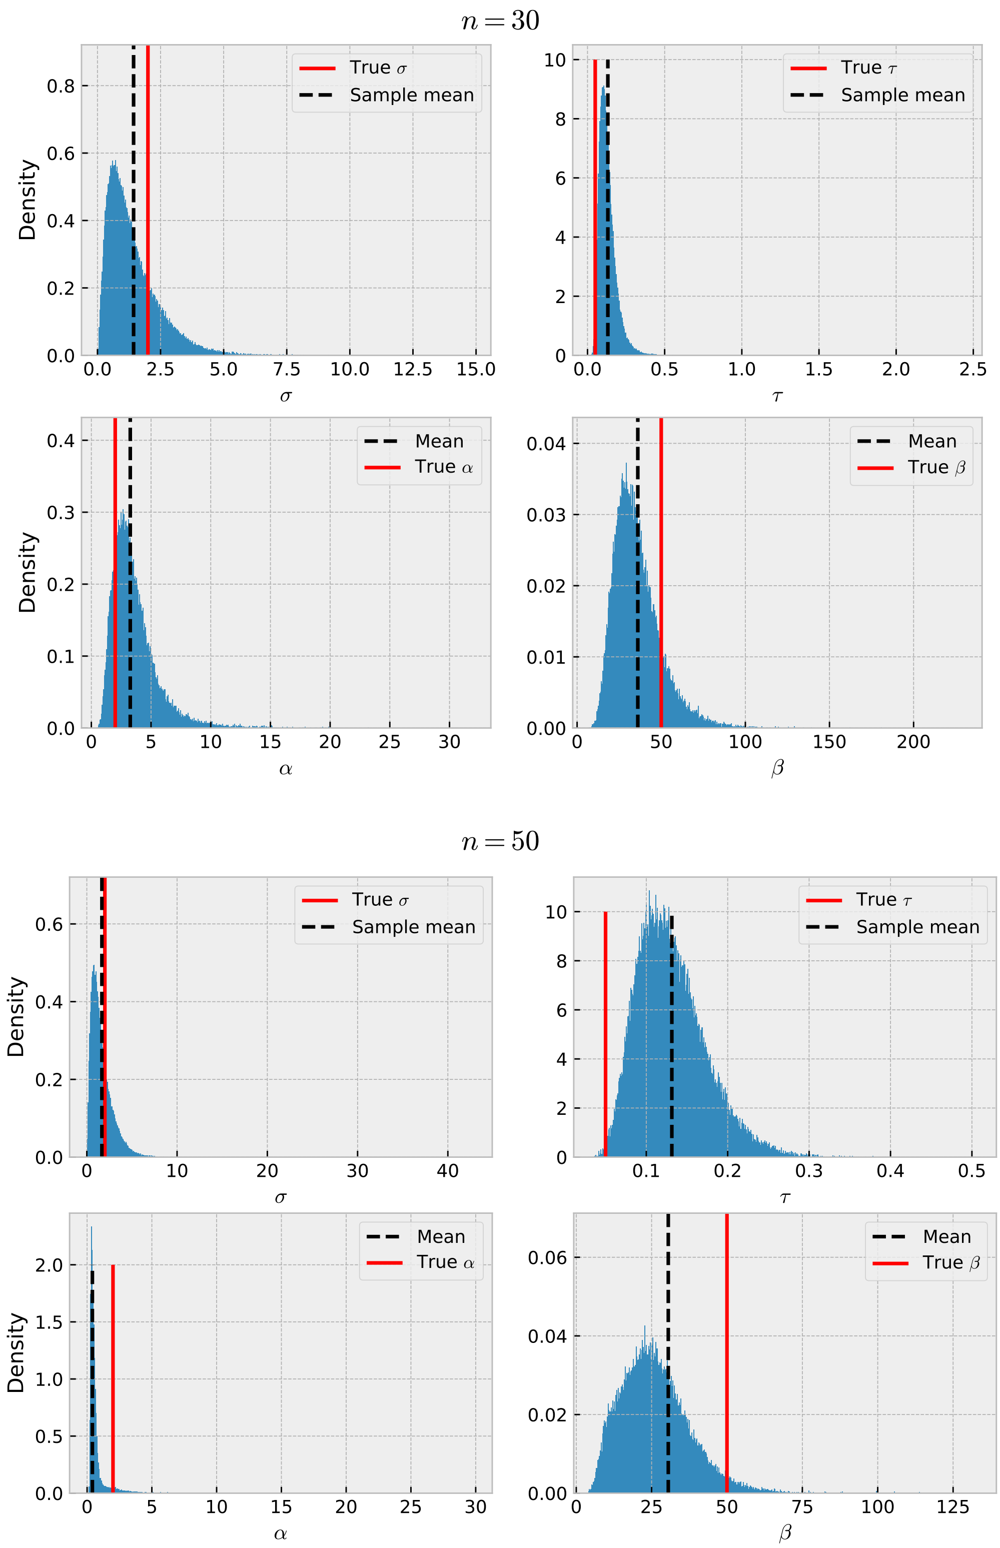

Figure 5.7: The corresponding histograms after discarding the burn-in period for the trace plots presented in Figure 5.6.

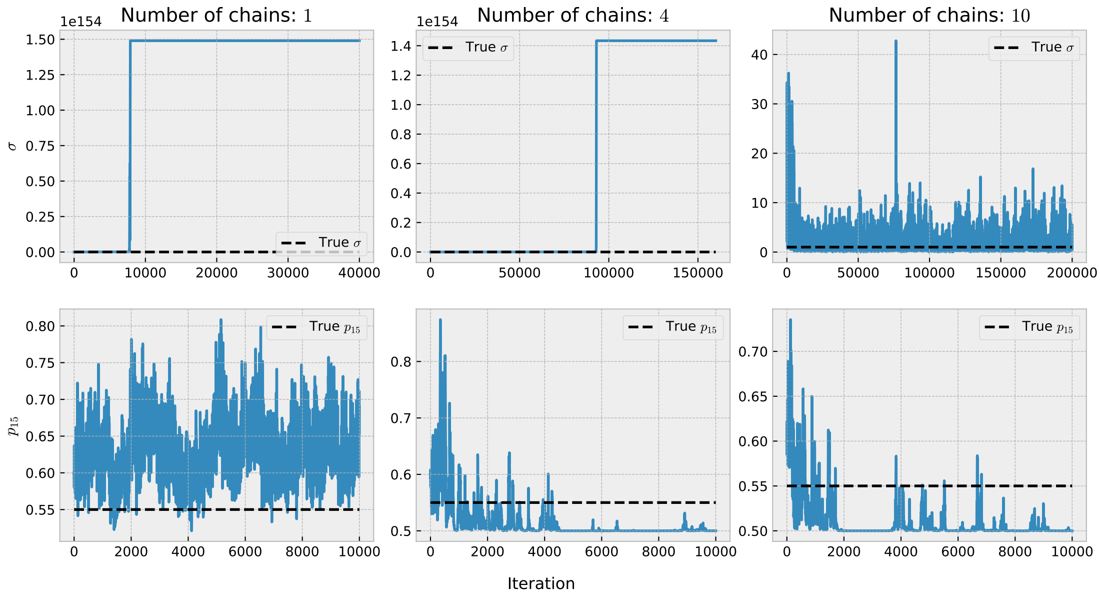

Figure 5.8: Trace plots of simulated chains for Case 1 with a differing number of input decision chains. In all runs, it has been used chains of length $n = 50$. It is clear that the derived method is becomes more stable for an increasing number of input decision chains.

where ten input decision-chains are used, and the computations here seems to be stable, although there is a small spike for the sampled σ-values around iteration 75000. This may indicate that calculations are somewhat unstable even when using ten chains as input. As previously, mixing is clearly poor for the sampled competences. The true value of the competences is rather close to 0.5 (the true mean value for $p_i = 0.52$), and the method seems to reach a steady-state at 0.5, meaning it detects the overall low competences in the population. Also for the posterior of the competences, we see that only one decision chain as input to the method carries too little information to give reasonable estimates, as the true value is mostly outside the variation range of the samples.

Table 5.6 contains summary statistics for the run with 10 input decision chains. The table shows an overview of sample means along with empirical 90% credible intervals from the marginal posterior distribution of each parameter. From the table it is clear that the derived method does not give satisfying results for Case 2, for these particular choices of parameters. We will discuss this further in the next experiment. Overall, the same tendencies are however observed for both of the cases in Experiment 2. The method is more unstable when using fewer decision-chains as input. Based on the results presented in this subsection, four or fewer chains give numerically unstable results, whereas for ten chains the results seem stable. Similarly, using long chains where cascades have occurred at an early individual gives little additional information to the system compared to using

		Case 1			Case 2	
	True	$n = 30$	$n = 50$	True	$n = 30$	$n = 50$
σ	1.00	1.44 $(0.28, 3.49)$	1.60 $(0.31, 3.91)$	2.00	1.48 $(0.29, 3.66)$	1.3 $(0.18, 4.04)$
τ	0.05	0.13 $(0.066, 0.24)$	0.13 $(0.07, 0.21)$	0.10	0.09 $(0.04, 0.18)$	0.15 $(0.07, 0.25)$
α	2.00	3.27 $(1.30, 6.40)$	0.41 $(0.24, 0.65)$	15.00	2.06 $(0.88, 4.05)$	0.22 $(0.05, 0.35)$
β	50.00	36.18 $(18.24, 63.55)$	30.52 $(16.31, 49.86)$	20.00	13.60 $(5.59, 25.37)$	0.27 $(0.11, 0.51)$

Table 5.6: Posterior means together with values for the endpoints of the 90% credibility intervals from each of the estimated marginal posterior distributions for Experiment 2 using 10 decision-chains as input.

shorter chains where only the most informative parts are included.

5.3.3 Experiment 3

Figure 5.9 shows the corresponding histograms of Run 1 for Case 1, after discarding the burn-in period of $m = 100000$ iterations for the hyperparameters. These histograms represents the posterior marginal distributions resulting from the MCMC procedure. The red line in the plots shows the true parameter value used to simulate the decisions, while the black line shows the sample mean. Overall, the histograms show varying results. For the parameter τ, the true value is far out in the tail of the distribution. This suggests that the method is not able to provide a reasonably good estimate for this parameter. The same is the case for the parameter α, which again have an effect on the sampled competences p_i. The sampled posterior distribution looks similar for both individual 5 and 15, even though the true competence of individual 15 is slightly higher than the average in the distribution. This might suggests that there is not enough information in the decisions for the method to be able to distinguish the different competences between different individuals. The parameter estimates for σ and β seems reasonable for this particular case. However, we note that the variation of both posteriors seem very large as many values are covered in the variation range.

As previously mentioned, there did not seem to be enough information in Case 2 for the samples to properly converge. Hence, the assumption that the resulting chains constitute samples from the posterior distribution does not hold, and the samples cannot be used for inference as they do not represent the distribution of interest. The differences in performance of the method was anticipated, and is probably caused by the fact that there is less information contained in the decisions from Case 2, in combination with the vague prior knowledge. We therefore move directly to Case 2*. Figure 5.10 shows the histograms corresponding to the third run in Experiment 1 for Case 2*. This run has favourable initial values, so the sampled Markov chains seem to converge rather early, as seen in the trace plots in Figure 5.5. The overall tendency is the same as for Case 1. For many of the parameters, the true parameter value lies within the variation range of the sampled posterior. However, the sampled distribution for the parameter α has a long tail, and the true value lies far out in this tail. The sampled posteriors of the competences plotted in the lowest row seem reasonable, although the variances of the

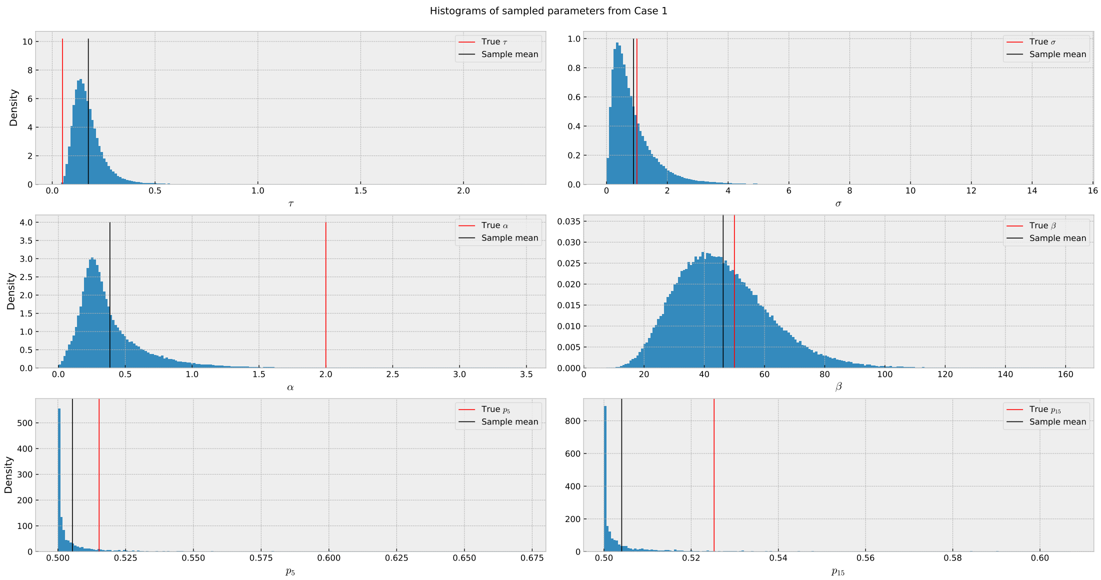

Figure 5.9: Histograms of the sampled marginal posteriors of the hyperparameters and competences for individuals 5 and 15, corresponding to Experiment 1, Run 3 for Case 1.

distributions are large.

Overall, the amount of information in the decision-chains alone seems to be insufficient for the method to give reasonably good estimates for all of the system parameters, as the sampled joint posterior distributions are not sufficiently accurate. In Case 1, we are for example able to detect that the competences in the population are low, but not exactly *how low*. The method accepts mostly competences close to 0.5, which does not entirely resemble the true distribution, which has its peak at a slightly higher competence. We further need a rather specific prior knowledge about the overall competences in the population in order to obtain reasonable estimates in Case 2. This is not necessarily favourable, as it may be unrealistic to assume that this amount of prior knowledge is available. Additionally, sensitivity for different priors indicates that there is little contribution in terms of information from the decision-chains alone.

We now compare the results discussed above to the same experiments for the situation where we assume that the private signals $\mathbf{y}$ are observed together with the decisions. As previously discussed, the results for Case 2 up to this point has been poor, as the chains does not properly converge for different starting values. The same tendency is also evident when we condition on the private signals as well as the decisions. Since there is little improvement by also assuming that the private signals are known, the conclusion is that the combination of little informative decision-chains, together with vague prior knowledge is not sufficient to obtain reasonable samples from the posterior distribution of the system for this particular case.

The histograms in Figure 5.11 visualises the differences between the sampled marginal

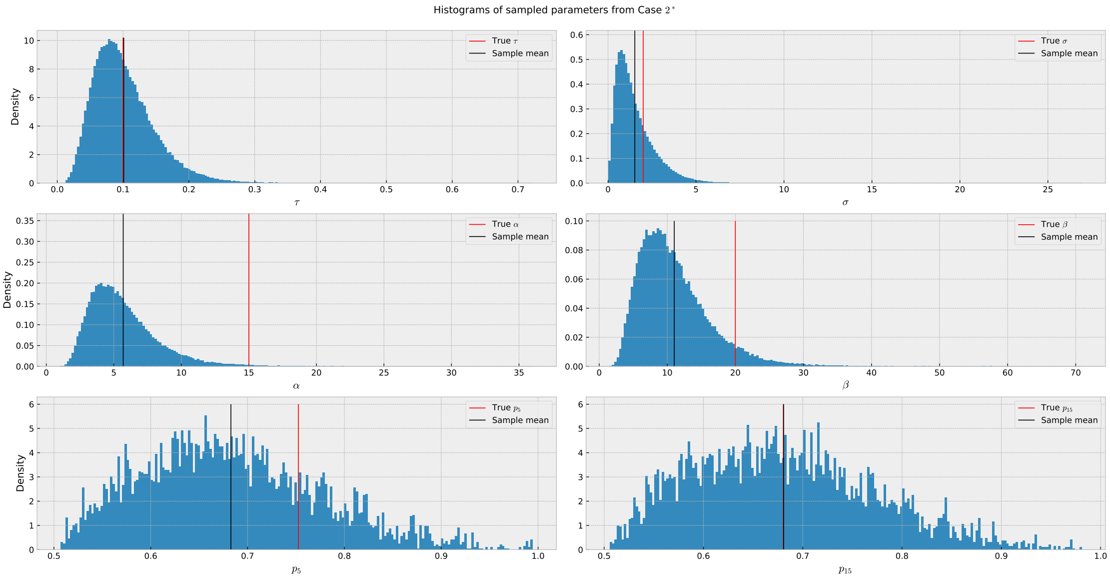

Figure 5.10: Histograms of the sampled marginal posteriors of the hyperparameters and competences for individuals 5 and 15, corresponding to Experiment 1, run 3 for Case 2*.

posteriors in Case 2* when the private signals are assumed to be known versus unknown. Conditioning on more information do not seem give any pronounced improvement on the sampled posteriors of the system for Case 2*, as the marginal posteriors seem quite similar. The same tendency is observed for Case 1. The tables 5.7-5.9 shows the mean value and the endpoints of the equal-tailed 90% credible intervals from the sampled marginal distributions for each case. In the tables, the left column shows summary statistics where the private signals are assumed to be unknown, and the right columns shows summary statistics where we assume private signals are observed. Hence, the tables shows the differences between the situation where the private signals are observed versus unobserved for all cases for a given run. Overall, the gain of assuming that the private signals known seem to be small. For Case 1, slight improvements of the estimate of the parameters σ, τ and β can be observed in terms of a sample mean that is closer to the true value, and more accurate credible intervals for the marginal posterior distributions. For Case 2, there is no improvement. Additionally, the sampled posterior for σ seems to have a particularly long tail as some very large values are accepted. For Case 2*, the different samples seem to give equally accurate estimates.

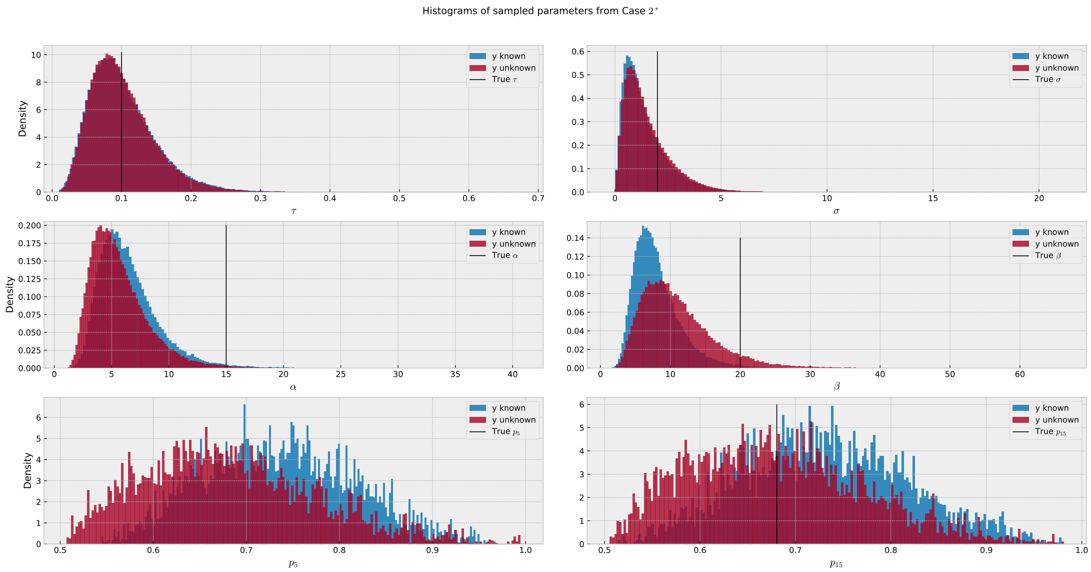

Figure 5.11: The figure illustrates differences of the sampled marginal posteriors of the hyperparameters and competences for individuals 5 and 15 for the situations where the private signals are assumed to be known and unknown. The runs correspond to Run 3 in Experiment 1.

The results presented in this subsection suggest that the amount of information in the system is insufficient. It is difficult to obtain sufficiently accurate estimates for the posterior distribution of the parameters of interest using the implemented method. In general, better estimates can be obtained by for example conditioning on even more data. Further improvements can also be obtained by using even more and longer decision-chains as input. However, this may be unrealistic to achieve when applying the method in practice, as the amount of available decision-data may be limited. Additionally, this will increase the required cpu-time and it may be necessary to improve the algorithm in terms of efficiency.

	True	Case 1 $\mathbf{y}$ unknown	 $\mathbf{y}$ known
σ	1.00	1.44 (0.30, 3.48)	1.29 (0.24, 3.20)
τ	0.05	0.17 (0.09, 0.31)	0.11 (0.04, 0.20)
α	2.00	0.18 (0.01, 0.44)	0.16 (0.01, 0.39)
β	50.00	47.10 (22.60, 81.78)	49.26 (24.08, 84.59)

Table 5.7: Summary statistics for Case 1 in Experiment 3. The setup corresponds to Run 3 in Experiment 1. The table shows posterior means along with values for the endpoints of the 90% credibility intervals in parentheses from each of the estimated marginal posterior distributions for the hyperparameters.

	True	Case 2 $\mathbf{y}$ unknown	 $\mathbf{y}$ known
σ	2.00	1.20 (0.24, 2.93)	14.02 (1.81, 41.17)
τ	0.10	0.13 (0.06, 0.23)	0.14 (0.07, 0.25)
α	15.00	0.36 (0.21, 0.56)	3.77 (1.81, 6.85)
β	20.00	19.63 (6.98, 41.38)	28.78 (10.95, 63.02)

Table 5.8: Summary statistics for Case 2 in Experiment 1.The setup corresponds to Run 3 in Experiment 1. The table shows posterior means along with values for the endpoints of the 90% credibility intervals from each of the estimated marginal posterior distributions for the hyperparameters.

	True	Case 2* $\mathbf{y}$ unknown	 $\mathbf{y}$ known
σ	2.00	1.52 (0.30, 3.67)	1.48 (0.29, 3.61)
τ	0.10	0.10 (0.04, 0.19)	0.14 (0.04, 0.19)
α	15.00	5.73 (2.60, 10.62)	6.53 (3.30, 11.49)
β	20.00	10.89 (3.93, 23.96)	8.05 (3.54, 16.57)

Table 5.9: Summary statistics for Case 2* in Experiment 3. The setup corresponds to Run 3 in Experiment 1. The table shows posterior means along with values for the endpoints of the 90% credibility intervals from each of the estimated marginal posterior distributions for the hyperparameters.

Chapter 6

Closing remarks

In this report, we have introduced and defined a statistical model for sequential decision-making where each individual have uncertain estimates about other individuals' competences. We derived a general expression for the probability of each possible decision for each decision-maker, and this expression was in turn implemented numerically. For different choices of parameter-values, we are able to simulate chains of decisions from the proposed model. We further derived an MCMC-method for parameter estimation based on the observed decisions, and performed a simulation study to assess the performance of this method.

The results from the simulation study suggest that the sampled Markov chains essentially converges for Case 1 and 2*. However, the convergence is slow, and assessing the convergence of an MCMC-procedure is in general difficult. The varying accuracy of the estimated parameters from the different marginal posterior distribution suggests that it is difficult to obtain sufficiently accurate samples from the joint posterior distribution of the system parameters based on the decision-chains alone. The method relies on specific knowledge about the parameters α and β. Based on this, we conclude that the information in the observed decisions alone is sparse. Allowing more information to enter the system by assuming that the private signals $\mathbf{y}$ are observed as well did only give a marginal improvement of the results. Most of the information lies in the first parts of the decision-chains, until the cascades occur. Taking advantage of this speed up computations. However, several chains of decisions are required to obtain stable results.

When it comes to the properties of the sampled posterior distribution, it is reasonable to assume that parameter estimates will improve by allowing for more information to enter the system. In particular and based on the results presented in Chapter 5, this involves using considerably more than ten decision-chains as input, or longer and more informative decision-chains. However, in addition to the increased demand of cpu-time, the necessary amount of decision-chains of the same nature may be difficult to acquire in practical situations. If it is achievable, another possibility may be to assume that one or more of the other system variables are known, and condition on this data as well.

In future work, we suggest addressing the convergence issues of the method. In general, this can be accomplished by running the algorithm for more iterations than what was done in the numerical experiments presented here. Further improvements in order to make the algorithm more efficient involve using a more sophisticated proposal distribution. This can be done by for example updating multiple parameters at each iteration.

References

Anderson, L. R. and Holt, C. A. (1997). "Information cascades in the laboratory." *The American economic review*, 87, 847–862.

Banerjee, A. V. (1992). "A simple model of herd behavior." *The quarterly journal of economics*, 107, 797–817.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy*, 100, 992–1026.

— (1998). "Learning from the behavior of others: Conformity, fads, and informational cascades." *Journal of economic perspectives*, 12, 151–170.

Dawid, A. P. (1979). "Conditional independence in statistical theory." *Journal of the Royal Statistical Society: Series B (Methodological)*, 41, 1–15.

Duan, W., Gu, B., and Whinston, A. B. (2009). "Informational cascades and software adoption on the internet: an empirical investigation." *MIS quarterly*, 33, 23–48.

Falnes, I. M. (2019). "Modelling observational learning and informational cascades under different information about other individuals' competence." *Unpublished*, Specialization project, Department of Mathematical Sciences, NTNU.

Frederic, P. and Lad, F. (2008). "Two moments of the logitnormal distribution." *Communications in Statistics—Simulation and Computation*, 37, 1263–1269.

Gamerman, D. and Lopes, H. F. (2006). *Markov Chain Monte Carlo. Stochastic Simulation for Bayesian Inference*. Chapman & Hall/CRC. Taylor & Francis Group.

Givens, G. H. and Hoeting, J. H. (2013). *Computational Statistics*. John Wiley & Sons, Ltd.

Hastings, W. K. (1970). "Monte Carlo sampling methods using Markov chains and their applications."

Højsgaard, S., Edwards, D., and Lauritzen, S. (2012). *Graphical models with R*. Springer Science & Business Media.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). "Equation of state calculations by fast computing machines." *The journal of chemical physics*, 21, 1087–1092.

Roberts, G. O. and Smith, A. F. (1994). "Simple conditions for the convergence of the Gibbs sampler and Metropolis-Hastings algorithms." *Stochastic processes and their applications*, 49, 207–216.

Rosas, F., Hsiao, J.-H., and Chen, K.-C. (2017). "A technological perspective on information cascades via social learning." *IEEE Access*, 5, 22605–22633.

Walpole, R. E., Myers, R. H., Myers, S. L., and Ye, K. (2012). *Probability & Statistics for Engineers & Scientists*. Pearson.

Zentall, T. R. (2006). "Imitation: definitions, evidence, and mechanisms." *Animal cognition*, 9, 335–353.

NTNU
Norwegian University of
Science and Technology