Validation and Parameter Estimation of a Behaviour-Modification Model

by

Nikesh Parsotam

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2015

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this thesis, an attempt at validating and identifying parameters of a quantitative model of attitudes and behaviours is presented. The model, established in earlier work, describes the dynamics of a subject's attitude and behaviour when he or she is offered a sequence of rewards in exchange for producing some desired behaviour. The dynamics of attitude are governed by the theory of planned behaviour, cognitive dissonance theory, and overjustification theory. Validation is performed in two specific cases in which overjustification does not arise. These two cases are defined based on the sign of the attitude. In the first case, the subject's attitude is negative, while in the second case, the subject's attitude is positive and all offers are declined (this additional requirement is necessary to ensure overjustification effects are not present).

The parameters to be identified are the cognitive dissonance gain parameter, K_1, sensitivity to a reward, μ_1, and the attitude measurement proportionality constant, μ_2. These parameters are unique for each person. A review of standard parameter estimation approaches concludes that the extended Kalman filter (EKF) is best suited for parameter estimation in this model. Simulations show that in this particular application, the EKF cannot accurately estimate multiple parameters simultaneously; thus, only K_1 is estimated using the EKF. An experiment is designed to produce the required conditions for each case and elicit attitude and behaviour data from human participants at 11 instances. In the initial phase of the experiment, carefully phrased questions are used to estimate μ_1 and μ_2. In the remainder of the experiment, participants are iteratively offered a reward to listen (for a specified duration) to a sound they initially rated as unpleasant or pleasant (depending on whether data is being collected for the negative or positive attitude case). Following their response to an offer, the participants are asked to rate the sound. The details of each offer, the participant's response, and the rating of the sound following each offer are used to validate the model and estimate K_1.

The experimental data collected in the first case (negative attitude) shows that the model correctly captures experimental trends in 5 of 7 trials. Further, the EKF's estimates of K_1 were almost always positive and appear to be converging the majority of the time. These two observations support the validity and utility of the model when attitude is negative. The experimental data collected in the second case (positive attitude) is largely unusable. The unusable data is attributed to multiple deficiencies in the design of the experiment. The deficiencies result in the experiment not accurately producing the conditions required to excite the effects of cognitive dissonance when attitude is positive.

Acknowledgements

I would like to thank my advisor, Professor Daniel Davison, for his guidance, patience, feedback, and support. I would also like to thank my family for their never ending love, motivation, and support. Finally, thanks to my friends who helped me enjoy this journey.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

A main objective of psychology research is to understand and explain a wide range of human behaviours. Of particular interest in this thesis is the field of social psychology, which encompasses how people's thoughts, behaviours, and interactions are affected by other people. Social psychology researchers have proposed theories that capture the effects of various social influences. An intricate network of social influences surround every thought or action. Sometimes these influences are subtle and go unnoticed, while other times they are quite clear.

Examples of such theories include *group polarization* and *groupthink*, which describe changes in cognitions (i.e., attitudes, thoughts, or internal feelings) as a result of group influences. Group polarization theorizes that each person in a group of like minded people will have a stronger attitude towards a topic following a group discussion [22][23][28]. In a group sharing the same view point, discussion allows a person to reinforce existing ideas and gather new ideas, thus strengthening their attitude in the direction of their original inclination. Further support of group polarization is provided by Festinger's *social comparison theory*, which posits that a person has a desire to evaluate their abilities and opinions, and may compare themselves to others in order to evaluate themselves [10]. In a group of like-minded people, a person may find their opinions are shared and thus supported by other group members. Suppose a group of like-minded people must make a decision; they may not make the best decision since their decision may not be impartial owing to their similar viewpoints. A more general group may still fail to make good decision as a result of the groupthink phenomenon. Groupthink, first studied by Janis, describes the tendency to seek cohesiveness and minimize conflict within a group to the extent that the group makes suboptimal decisions [16]. The symptoms of groupthink fall under three broad categories: overestimation of the group's ability and authority, closedmindedness, and pressure towards uniformity. Analysis of group decisions in history with poor outcomes have shown that the symptoms of groupthink are often present (e.g., the Titanic deciding to go at full speed despite several warnings of icebergs [23]). Similar analysis of group decisions with good outcomes have shown that emphasis was placed on preventing groupthink symptoms [23].

Additional theories such as *social facilitation* and *social loafing* describe changes in a person's behaviour when in a group. Social facilitation states that a person will perform a familiar and simple individual task better in the presence of others. Zajonc suggests being observed and evaluated by others produces social arousal [23][32]. This arousal improves performance on familiar tasks but reduces performance in more difficult or unfamiliar tasks. The effect is strongest when the person believes others are evaluating their performance [6][15]. Social facilitation primarily applies when a person is performing an individual task; in cases where a group collaborates to achieve a common goal, social loafing may occur. Social loafing is the tendency for each person to exert less effort towards achieving a goal when they work in a group as opposed to working individually [23]. Social loafing is prevalent in situations where individual performance is difficult to evaluate, or when each person cannot be held accountable. Consider pulling a rope, the activity which Ringelmann used to identify that combined effort of a group was not the sum of the individual efforts [19]. Each person's effort is difficult to evaluate, and thus each person may not pull their hardest. Suppose a group is asked to make a decision, it is possible that both social loafing and groupthink may apply.

While these are only four of the many theories in social psychology, they demonstrate that social influences apply in a wide variety of situations. Further, these illustrate that multiple sources of influence can be present. A criticism of these and many other theories in psychology is that they only have a qualitative notion of cause and effect. Discussions focusing on the effects of a stimulus are often based on the data at one time before the stimulus and one time after. Rarely do the discussions or experiments involve longitudinal data (i.e., data captured over a larger period of time at regular intervals).

Researchers with backgrounds in psychology or systems and controls (or sometimes both) have worked for decades to combine both these fields to produce models that better explain and describe human behaviour. Such models inherently include the idea of longitudinal data and allow for more complicated dynamics than 'cause and effect'. An example is Powers' *perceptual control theory* (PCT) [26]. PCT posits that a person's behaviour is the direct result of attempting to maintain some internal perception; thus, behaviour is the output of a control system. The theory models a person's internal perceptions in a hierarchical manner. To illustrate this, suppose a person wishes to open a door. Their hierarchical perceptions would be: I am near the door, I am oriented such that I face the door knob, I have my hand on the knob, I have twisted the knob, and I have applied sufficient force in the correct direction to open the door. The theory assumes that each hierarchical perception has an associated control system; thus, the control systems also form a hierarchy. The control system at each level is responsible for achieving or maintaining the perception at that level. Together, these perceptions and control systems describe behaviour.

Other models aim to enable positive lifestyle changes such as regularly exercising, quitting smoking, or making healthy food choices. One such example is the adaptive intervention proposed by Rivera, Pew, and Collins in [27]. In their work, the authors describe a model and develop a method to apply dynamic interventions (be it drugs or treatment) to help participants achieve their lifestyle changes. Their approach differs from standard intervention methods which prescribe fixed interventions which may be too much or too little at a given time instance. In [27], the authors simulate examples using quantitative measures. Another example is the combination of the *theory of planned behaviour*, *theory of cognitive dissonance*, and *overjustification theory* explained by Ni, Kulic, and Davison in [24]. This combination presents a fascinating opportunity to modify human attitudes and thus behaviour by offering a sequence of rewards in exchange for completing an activity.

The work in [24] is the most recent in a line of research related to behaviour modification led by Davison (see [29], [30], and [8] for the evolution of the model, alternate applications, and alternate control strategies). This line of research incorporates a control engineering perspective; as such, attitudes, behaviours, and rewards are quantified. The authors of [24] propose a controller that modifies behaviour to a desired level. The controller manipulates each offer of reward to excite cognitive dissonance and overjustification effects in order to modify attitudes (the simulated example in [24] increases attitude). Additionally, the control effort at the end is reduced to zero; that is, rewards are no longer required to produce the desired behaviour.

This thesis builds on the work of [24]. There are several practical issues that must be addressed before control design can be seriously considered. First, the model must be validated. Validation of the model ensures that a controller can be designed using an accurate plant model. Second, methods to quantitatively measure human attitudes and behaviour must be devised. Without such measures, control objectives cannot be realized. Finally, assuming adaptive control techniques are not applied, model parameters specific to each person must be identified before any control activity begins. This last issue also presents the challenge of determining whether the parameters are reasonably correct.

This thesis addresses the issues described above. The fundamental goal of this thesis is to validate the model presented in [24]. The validation focuses on two specific cases in

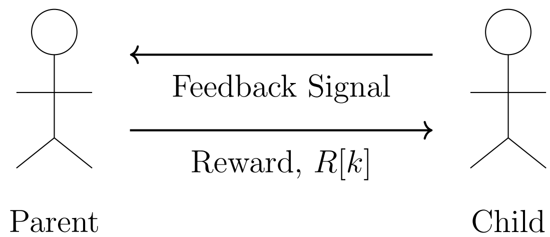

Figure 1.1: The basic setup in which a parent offers a reward $R[k]$ (with $R[k] \geq 0$) to their child to drive the child's behaviour to a desired level [24].

which overjustification effects are not present. The first case is when attitude is negative. The second case is when attitude is positive and offers are repeatedly declined. In parallel to validating the model, a feasible method to estimate relevant parameters is investigated.

As in [24], this thesis is built on the context that one person is attempting to modify another person's behaviour using a sequence of rewards. Figure 1.1 illustrates the context used in Chapter 2: a parent attempts to motivate their child to jog by offering them a reward. Only two people's interactions are considered in order to create the simplest scenario. Further, only the most basic set of cognitions are considered. External influences are assumed to be non-existent or negligible. As described earlier in this chapter, influences arise from many sources, and an attempt to account for all influences quickly increases the complexity of the model.

The following chapters guide the reader through the contributions and relevant details. More specifically, Chapter 2 presents the model used in this thesis and explains the two specific cases of the model. Chapter 3 reviews multiple parameter estimation approaches and discusses their suitability for use with the model. A subset of these approaches are selected for further study. Chapter 4 develops the necessary equations to apply the selected parameter estimation approaches to the model and provides simulation results. Chapter 5 contains the experiment design details and the results of the experiment. Chapter 6 draws conclusions based on the results of the experiment and summarizes the next steps for future work.

Chapter 2

The Behaviour Model

In this chapter, the models relevant to this thesis are described. The chapter begins with a review of the existing model and is followed by sections outlining simplifications that can be made under certain scenarios.

2.1 Review of Existing Model

The human behaviour model proposed in Ni, Kulic, and Davison's paper [24] is the model to be validated. This model evolved from the works of Vanderwater, Davison and Zhou's models in [29], [30], and [8]. The detailed description of this model in this section is similar to those in [24], [29], [30], and [8]. The equations described in this section are those presented in [24].

Preliminaries

The model is best described with the help of some context. The context used is that of a parent who wishes to encourage their child to perform an activity. More specifically, we will use the context of a father who would like to motivate his daughter to jog. Note that the model is a discrete-time system. Events are assumed to occur at discrete time instants denoted $k = 0, 1, 2, \cdots$.

The daughter's intrinsic attitude, denoted $A_{out}[k]$, is formed through a combination of outcomes regarding the activity. Examples of outcomes for jogging are exhaustion, perspiration, improved health, and less time to play video games. Each outcome produces a subcomponent of attitude. For simplicity, a single outcome is considered in this model. A question that arises when discussing attitudes is what units to use. Any arbitrary units may be used; however, this will influence the other parameters of the model. In the jogging example with the father and daughter, units of "minutes" are chosen;, that is, the attitude, $A_{out}[k]$, represents how many minutes the daughter would jog on day k without external motivation. Negative attitudes are allowed in the model despite lacking physical meaning. Negative attitudes are essential as to allow the model to capture different levels of dislike or resistance to an activity. A practical note is that the father may measure his daughter's attitude using an instrument that produces an output, $z[k]$, in arbitrary units. The instrument's output is assumed to follow a linear relationship with the daughter's attitude,

$$A_{out}[k] = \mu_2 z[k], \tag{2.1}$$

where $\mu_2 > 0$ is a proportionality constant that scales and normalizes the instrument's output to units of attitude (i.e., minutes).

Inducement to perform the activity comes in the form of an external reward, $R[k]$. The father may entice the daughter to jog for some fixed duration the next day by offering rewards of money or cookies. To be concrete, assume money is used as the reward and is measured in dollars. The duration set by the father represents a desired behaviour, $B_d[k+1]$, measured in minutes ($k+1$ is used as this is the desired behaviour for the following day). The units of $B_d[k+1]$ are the same as that of attitude. Accepting the reward means the daughter agrees to perform the desired behaviour the following day. Declining the reward means the daughter does not agree to produce the desired behaviour the following day and may do as she pleases.

Theory of Planned Behaviour

The *theory of planned behaviour* relates an individual's internal views and attitudes, and external influences, to the individual's behavioural intentions [1][2]. These intentions influence the actual behaviour of the individual. These relationships are captured in the model as follows:

$$A_{out}[k] = A_{out}[k-1] + \Delta A_{out}[k-1], \tag{2.2}$$

$$A_{rew}[k] = r_1 A_{rew}[k-1] + (1-r_1)\mu_1 R[k-1], \tag{2.3}$$

$$BI[k] = A_{out}[k] + A_{rew}[k], \tag{2.4}$$

$$B[k] = \begin{cases} B_d[k] & \text{if } BI[k] \geq B_d[k] \text{ and } A_{out}[k] \leq B_d[k] \\ A_{out}[k] & \text{if } (BI[k] < B_d[k] \text{ and } A_{out}[k] \geq 0) \\ & \text{or } A_{out}[k] > B_d[k] \\ 0 & \text{otherwise.} \end{cases} \tag{2.5}$$

Equation (2.2) captures, at a high level, the dynamics of the daughter's attitude. In words, (2.2) says the attitude at a given interval is the attitude at the previous interval plus the change in attitude at the previous time interval. The change in attitude is defined using cognitive dissonance theory and overjustification theory. It is described in greater detail below.

The attitude towards the reward, $A_{rew}[k]$, represents the external influences as per the theory of planned behaviour. The attitude to the reward is calculated in (2.3) and includes first-order mental processing dynamics. This introduces *memory* of earlier rewards allowing for previous attitudes to affect future attitudes. The influence of this memory is determined by the constant r_1 which may take values in $(0, 1)$. The reward, $R[k-1] \geq 0$, is scaled and converted by μ_1 (with $\mu_1 > 0$), a constant that captures the sensitivity of the daughter to the reward and normalizes the reward into units of attitude.

The behavioural intent, $BI[k]$, is given in (2.4). In accordance with the theory of planned behaviour, the behavioural intent is simply the sum of the internal and external attitudes. The behavioural intent plays a significant role in determining the daughter's actual behaviour, $B[k]$. Both behavioural intent and behaviour have the same units as attitude.

Equation (2.5) presents the three possible resultant behaviours at any time. The first possibility is that the daughter's behavioural intent is greater than or equal to the desired behaviour while her intrinsic attitude is less than the desired behaviour requested in the offer. In this scenario the daughter's attitude related to the reward is large enough to increase her behaviour intent such that it is equal to or greater than the desired behaviour. As the daughter does not intrinsically wish to produce the desired behaviour, she will do the minimum required, that is $B_d[k]$, such that she may receive the reward. This holds true regardless of the whether the daughter's attitude is positive or negative.

The second possibility in (2.5) is that the daughter's behavioural intent is less than the desired behaviour but her intrinsic attitude is positive. The reward offered is too small and fails to motivate the daughter to produce the desired behaviour. Her behaviour would follow her intrinsic attitude. Suppose the daughter wished to jog for half the time requested by the father and that the father had offered only a tiny reward. The daughter would decline the offer as the reward is not valuable enough to cause her to change her behaviour and jog for half the time as that is all she desires to jog.

This second possibility in (2.5) also encompasses the scenario in which the daughter's intrinsic attitude is greater than the desired behaviour. In this scenario it is clear that the daughter's behavioural intent will be larger than the desired behaviour as the reward cannot be negative (therefore her attitude towards the reward cannot be negative). The daughter would accept the reward and produce behaviour that follows from her intrinsic attitude. Suppose the daughter enjoys jogging and wishes to jog for twice the time requested by the father. The daughter would have jogged regardless of whether the reward was offered or not and thus accepts the reward for jogging that would have been done in the absence of any external motivation. The daughter would then jog for twice the time requested.

The third and final possibility in (2.5) is that the daughter's behavioural intent is less than the desired behaviour and her intrinsic attitude is negative. In this scenario, the reward has failed to raise the daughter's behavioural intent to at least the desired behaviour. The daughter declines the reward and does nothing.

Absent from the model (2.2)-(2.5) is the notion of *perceived behavioural control* [1][2]. Perceived behavioural control embodies an individual's intrinsic belief of whether they are inherently capable of the activity. It also factors in any obstacles that may prevent the activity. In the theory of planned behaviour, perceived behavioural control is coupled with the intrinsic and external attitudes to form the behavioural intent. In the model presented in this thesis it is assumed that there are no obstacles to the desired activity or behaviour and that the individual has no concerns about being able to complete the activity. Referring to the jogging example once again, the assumptions imply that the daughter is among other things in good health and the weather is favourable. Further, the daughter does not question her ability to jog for the desired length of time.

This concludes the portion of the model related to the theory of planned behaviour. Reviewing the model it is clear that attitude is dependent on the change in attitude, $\Delta A_{out}[k]$. The change in attitude is determined in accordance with cognitive dissonance theory and overjustification theory. These theories are the subject of the remainder of this section.

Cognitive Dissonance Theory

Cognitive dissonance theory relates attitudes and behaviour. More specifically, it states that when a person's behaviour is inconsistent with their beliefs relating to the activity, an internal pressure called dissonance pressure arises [11]. The dissonance pressure is proportional to the magnitude of the inconsistent cognitions present in the individual. The individual will seek to reduce dissonance pressure as the pressure lingers and causes feelings of discomfort. Dissonance pressure can be reduced using approaches grouped into three broad categories. The first group of approaches is identified by changes in attitude [11]. A change in attitude will reduce inconsistencies or bolster existing consistencies, thus reducing dissonance pressure. The second group is characterized by the gathering of new information that creates additional consistencies [11]. The third and final group of dissonance reduction approaches is characterized by changes in the significance of existing cognitions which reduces the magnitude of the consistencies or inconsistencies present [11]. The individual may achieve this by rationalizing actions (e.g., saying that "the money is not worth it").

Multiple approaches can be used simultaneously to reduce dissonance pressure. However, to maintain simplicity in the model, it is assumed that dissonance pressure is reduced only by changing one's attitude. Gathering of new information is difficult to model as it may occur sporadically and is difficult to detect. Similarly, changes in significance of cognitions are also difficult to observe and measure; as such, the significance of cognitions are assumed to remain constant.

Referring to the father and daughter jogging example, suppose the father offers his daughter a sum of money to jog for some time and that jogging is an activity that the daughter dislikes. Dissonance pressure arises whether the daughter accepts or declines the offer. In the case of accepting the offer, the daughter has agreed to jog but this is inconsistent with her attitude towards jogging. This inconsistency produces a dissonance pressure. If the daughter declines the offer, the reward is declined, which is inconsistent with her desire for the reward. Once again, a dissonance pressure would arise. In an effort to reduce the pressure in either scenario, the daughter could change her attitude to create greater consistency. In the first case, offer accepted, greater consistency is possible only by reducing her dislike towards jogging. This also reduces the inconsistency. In the second case, offer declined, greater consistency is created by increasing her dislike towards jogging.

To keep track of the behaviour, a pair of indicator variables is used. The first, B_{sgn}, keeps track of whether any behaviour was produced:

$$B_{sgn}[k] = \begin{cases} 0 & \text{if } B_d[k] = 0 \\ +1 & \text{if } B[k] \geq B_d[k] > 0 \text{ or } A_{out}[k] \geq 0 \\ -1 & \text{otherwise.} \end{cases} \tag{2.6}$$

In the example of the father and daughter, $B_{sgn}[k]$ is assigned 0 if nothing is requested of or offered to the daughter (i.e., $B_d[k] = 0$). This null case ensures that no change in attitude arises from disingenuous requests. $B_{sgn}[k]$ is assigned +1 if the daughter jogs on day k for any length of time. $B_{sgn}[k]$ takes value -1 in all other conditions, indicating the daughter did not jog (even in the presence of an offer containing nonzero reward and nonzero desired behaviour).

The second indicator variable, $B_{rel}[k]$, indicates whether the desired behaviour was produced or not produced by taking values $+1$ and -1 respectively. Alternatively, $B_{rel}[k]$ can be viewed as an indicator of whether the reward was accepted or declined:

$$B_{rel}[k] = \begin{cases} +1 & \text{if } B[k] \geq B_d[k] \\ -1 & \text{otherwise.} \end{cases} \tag{2.7}$$

Referring to the father and daughter jogging example, $B_{rel}[k]$ is assigned a value of $+1$ if the daughter accepts the offer and jogs the number of minutes requested (or more) by the father. If the daughter declines the offer, $B_{rel}[k]$ is assigned a value of -1 regardless of how much she jogs.

The two indicator variables allow the raw dissonance pressure, $P_{raw}^{CD}[k]$, to be calculated as follows:

$$M_{incon}^{1}[k] = \begin{cases} |A_{rew}[k]| & \text{if } \mathrm{sgn}(A_{rew}[k]) \neq B_{rel}[k] \\ 0 & \text{otherwise,} \end{cases} \tag{2.8}$$

$$M_{incon}^{2}[k] = \begin{cases} |A_{out}[k]| & \text{if } \mathrm{sgn}(A_{out}[k]) \neq B_{sgn}[k] \\ 0 & \text{otherwise,} \end{cases} \tag{2.9}$$

$$M_{con}^{1}[k] = \begin{cases} |A_{rew}[k]| & \text{if } \mathrm{sgn}(A_{rew}[k]) = B_{rel}[k] \\ 0 & \text{otherwise,} \end{cases} \tag{2.10}$$

$$M_{con}^{2}[k] = \begin{cases} |A_{out}[k]| & \text{if } \mathrm{sgn}(A_{out}[k]) = B_{sgn}[k] \\ 0 & \text{otherwise,} \end{cases} \tag{2.11}$$

$$M_{incon}[k] = \sum_{i=1}^{2} M_{incon}^{i}[k], \tag{2.12}$$

$$M_{con}[k] = \sum_{i=1}^{2} M_{con}^{i}[k], \tag{2.13}$$

$$P_{raw}^{CD}[k] = \begin{cases} B_{sgn}[k]\dfrac{M_{incon}[k]}{M_{incon}[k] + M_{con}[k]} & \text{if } M_{incon}[k] + M_{con}[k] > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{2.14}$$

Equations (2.8) through (2.11) quantify the basic consistencies and inconsistencies between behaviours and attitudes. More specifically, (2.8) represents the first inconsistency which arises when the desired behaviour is not produced. Referring to the jogging example, the daughter always finds a reward desirable; thus, $sgn(A_{rew}[k]) = +1$. If the daughter's behaviour is insufficient to attain the reward, then $B_{rel} = -1$, and an inconsistency is present between the daughter's desire for the reward and her behaviour. The magnitude of the inconsistency is proportional to the magnitude of her attitude towards the reward. Consider the following scenario: the father offers his daughter \$10 to jog 30 minutes but she declines the offer and instead jogs for 10 minutes. In this scenario the offer undoubtedly results in $A_{rew} > 0$ and her behaviour results in, referring back to (2.7), $B_{rel} = -1$. The daughter would internally recognize the inconsistency of declining a reward that is desirable; the magnitude of the inconsistency would depend primarily on her attitude to the reward.

Equation (2.9) captures the second type of inconsistency that arises when the daughter's attitude towards an activity is inconsistent with her behaviour. This occurs when the daughter enjoys an activity but does not perform the activity, or dislikes the activity yet does it anyway. In the model presented, the inconsistency arises only in the case when $A_{out}[k] < 0$ due to the careful definitions in (2.6) (which states that when $A_{out}[k] \geq 0$, $B_{sgn}[k] = +1$). The magnitude of the inconsistency is proportional to the magnitude of the daughter's attitude towards the activity.

Equation (2.10) captures the consistency that occurs when an offer and associated reward are accepted. This consistency is proportional to the magnitude of the daughter's attitude towards a reward.

Equation (2.11) represents the consistency between the daughter's view on the activity and her behaviour. This consistency arises in two scenarios: first, when she dislikes an activity and does not partake in the activity and, second, when she enjoys an activity and performs the activity. This consistency is dependent on the magnitude of her attitude towards the activity. Thus when the daughter's behaviour is consistent with attitude, stronger attitudes result in larger feelings of consistency.

The model allows for other consistencies and inconsistencies to be modelled. Additional equations and indicator variables would be needed to do accommodate this; nevertheless, it is possible. In this thesis only the basic consistencies and inconsistencies are considered.

The raw dissonance pressure is calculated as a fraction of the sum of the inconsistencies over the sum of all consistencies and inconsistencies as shown in (2.14). $B_{sgn}[k]$ determines whether the raw dissonance pressure drives attitude up or down.

Note that we assume the only method used to reduce dissonance pressure is by changing one's attitude towards the behaviour. Referring to the father and daughter jogging example, the following scenarios demonstrate that changes in attitude towards the behaviour reduce dissonance pressure:

- Suppose the daughter dislikes jogging and declines an offer. Dissonance pressure would arise due to the inconsistency between the daughter's behaviour and her desire for the reward. To reduce this pressure, greater consistency is required amongst consistent cognitions. The only consistent cognition present is that between the dislike of jogging and refusing to jog. Thus the daughter can increase consistency by further disliking jogging (i.e., lowering her attitude towards jogging).

- Suppose, once again, that the daughter dislikes jogging, but accepts an offer rather than declining it. The inconsistency causing dissonance in this scenario is due to the daughter jogging despite her negative attitude towards jogging. The only method for the daughter to decrease dissonance would be by changing her attitude towards jogging such that it is smaller in magnitude. Given her initial negative attitude, attitude must increase.

- Alternatively, suppose the daughter enjoys jogging but not to degree that is requested in an offer, nor is the offer large enough to convince her to jog. In this case she would likely decline the offer. The daughter would experience dissonance pressure stemming from the inconsistency between her desire for the reward and declining it. As the daughter will jog an amount dictated by her internal attitude, there is consistency between her attitude and behaviour. Increasing this consistency, by raising her attitude, reduces the dissonance pressure.

In addition to the three scenarios mentioned above, there is a special case in which attitudes may not change as described. In the event that only a small positive attitude is present towards an activity and a large request for behaviour is made with a small associated reward, the daughter may decline the reward due to *attitude reversal* [24]. Rather than the normal response, in which dissonance is reduced by increasing attitude, the daughter may revert to a negative attitude to reduce dissonance. In this manner, her attitude would be driven lower and lower. Attitude reversal can be viewed as simply "giving up" on the activity. Attitude reversal is captured by the indicator $r[k]$ defined below:

$$r[k] = \begin{cases} +1 & \text{if } B_d[k] - BI[k] > \alpha_{rev} A_{out}[k], A_{out}[k] \geq 0, \\ & K_1 P^{CD}[k] > 2A_{out}[k] \text{ and } A_{rew}[k] > 0, \\ -1 & \text{otherwise.} \end{cases} \tag{2.15}$$

Reviewing (2.15), $r[k]$ is set to $+1$ if four conditions are simultaneously satisfied. The first condition, $B_d[k] - BI[k] > \alpha_{rev} A_{out}[k]$, ensures that the requested behaviour, $B_d[k]$, is large and that the attitude towards the activity, $A_{out}[k]$, is relatively small. The parameter α_{rev} captures how likely the daughter is to "give up" and takes values larger than 1. The second condition is simply to ensure $A_{out}[k]$ is positive. The third condition implies that the dissonance pressure must be large enough to cause the daughter's new attitude to be larger in magnitude than the existing attitude. This is necessary to ensure that the change in attitude results in decreased dissonance pressure. The final condition implies that a nonzero reward must be present by requiring that $A_{rew}[k]$ be nonzero.

Similar to $A_{rew}[k]$, first-order mental processing dynamics are included when calculating the change in attitude due to dissonance pressure, $P^{CD}[k]$. The time constant r_2 determines the significance of the processing dynamics:

$$P^{CD}[k] = \begin{cases} (1 - r_2) P^{CD}_{raw}[k] & \text{if } r[k-1] = 1 \\ r_2 P^{CD}[k-1] + (1 - r_2) P^{CD}_{raw}[k] & \text{otherwise.} \end{cases} \tag{2.16}$$

In (2.16) it can be seen that the first-order processing dynamics are reset when reversal occurs. The logic behind this is that the drastic change in psychological state will cause any memories of past dissonance pressure to be ignored.

The final change in attitude, $\Delta A^{CD}_{out}[k]$, due to the effects of cognitive dissonance is proportional to dissonance pressure as calculated in (2.16). The strictly positive proportionality constant, K_1, with the same units as attitude (minutes in the example used), relates the two parameters:

$$\Delta A^{CD}_{out}[k] = \begin{cases} -K_1 P^{CD}[k] & \text{if } r[k] = 1 \\ +K_1 P^{CD}[k] & \text{otherwise.} \end{cases} \tag{2.17}$$

Overjustification Theory

Overjustification theory states that external rewards given to a person to perform an activity will reduce their intrinsic motivation or desire to perform that activity [9]. This is of great importance in our model. The minimum level that attitude may decrease to due to overjustification effects is denoted by $B_t[k]$. In the model, we assume this attitude level is proportional to $B_d[k]$:

$$B_t[k] = \alpha_{B_d} B_d[k]. \tag{2.18}$$

The proportionality constant, α_{B_d}, seen in (2.18), takes values in $(0, 1)$. Overjustification effects will not cause attitudes to become less than $B_t[k]$; thus, attitude may only decrease due to overjustification effects if $A_{out}[k] > B_t[k]$. If this condition is not satisfied, then, referring to the father and daughter jogging example, the daughter's attitude is already below the minimum level dictated by $B_t[k]$ and attitude cannot further decrease. To capture this effect we introduce attitude relative to the level $B_t[k]$:

$$A_{out}^{rel}[k] = \max\{0, A_{out}[k] - B_t[k]\}. \tag{2.19}$$

Overjustification pressure is modelled as the product of the attitude to the reward and the relative attitude:

$$P_{raw}^{OJ}[k] = \begin{cases} A_{out}^{rel}[k]A_{rew}[k] & \text{if } A_{out}^{rel}[k] > 0 \text{ and } A_{rew}[k] > 0 \text{ and } B[k] \geq B_d[k] \\ 0 & \text{otherwise.} \end{cases} \tag{2.20}$$

The raw overjustification pressure is filtered by the same first-order mental processing dynamics presented earlier in the model. The constant $r_3 \in [0, 1)$ captures the effects of the mental processing dynamics:

$$P^{OJ}[k] = r_3 P^{OJ}[k-1] + (1 - r_3)P_{raw}^{OJ}[k]. \tag{2.21}$$

The change in attitude due to overjustification effects, $\Delta A_{out}^{OJ}[k]$, is proportional to the overjustification pressure calculated in (2.21). The proportionality constant, $K_2 > 0$, captures how significant overjustification effects may be for the daughter:

$$\Delta A_{out}^{OJ}[k] = \begin{cases} -K_2 P^{OJ}[k] & \text{if } K_2 P^{OJ}[k] \leq A_{out}^{rel}[k] \\ -A_{out}^{rel}[k] & \text{otherwise.} \end{cases} \tag{2.22}$$

The condition $K_2 P^{OJ}[k] \leq A_{out}^{rel}[k]$ in (2.22) ensures that attitude is not driven below $B_t[k]$, thus remaining consistent with the theory.

Combining the Effects of Cognitive Dissonance and Overjustification

The attitude change, $\Delta A_{out}[k]$, first introduced in (2.2) is the sum of the change in attitude due to cognitive dissonance and overjustfication effects:

$$\Delta A_{out}[k] = \Delta A_{out}^{CD}[k] + \Delta A_{out}^{OJ}[k]. \tag{2.23}$$

The model is now complete.

In summary, the model is built on the theory of planned behaviour, cognitive dissonance theory, and overjustification theory. The theory of planned behaviour models high-level attitude changes at all instances. Cognitive dissonance theory applies primarily when attitude is negative and to a lesser degree when attitudes are positive. Overjustification theory's importance arises only when attitudes are positive and sufficiently large.

2.2 Model Simplifications when Attitude is Negative

It is convenient to assume that attitude remains either negative or positive when working with the model presented in the previous section. In this section we assume attitude begins and remains negative and simplify the model accordingly. Once again, the example of a father attempting to motivate his daughter to jog is used to illustrate aspects of the model. In this example, the daughter does not find jogging pleasant at first; thus, her intrinsic attitude, $A_{out}[k]$, towards jogging is negative at time $k = 0$.

We begin by stating our initial assumptions, primarily that attitude is negative. We also require the desired behaviour and reward to be strictly positive. This eliminates the need to handle spurious offers in our analysis. It should be noted that the full model is capable of handling such offers. Reiterating these assumptions, we have

$$\begin{aligned} A_{out}[k] &< 0 \\ B_d[k] &> 0 \\ R[k] &> 0 \\ &\forall k = 0, 1, 2, \ldots, N, \end{aligned}$$

where N is the number of offers to be made.

The theory of planned behaviour in the full model contains various first-order mental processing dynamics. These allow the memory of prior attitudes and pressures to influence future attitudes. To simplify the identification of key parameters in the model, it is assumed that ample time has passed between changes in attitude and offers such that the effect of changes and rewards at time k are fully experienced before time $k + 1$. This assumption implies:

$$\begin{aligned} r_1 &= 0 \\ r_2 &= 0 \\ r_3 &= 0. \end{aligned}$$

Thus (2.2) to (2.5) become:

$$A_{out}[k] = A_{out}[k-1] + \Delta A_{out}[k-1] \tag{2.24}$$
$$A_{rew}[k] = \mu_1 R[k-1] \tag{2.25}$$
$$BI[k] = A_{out}[k] + A_{rew}[k] \tag{2.26}$$

$$B[k] = \begin{cases} B_d[k] & \text{if } BI[k] \geq B_d[k] \text{ and } A_{out}[k] \leq B_d[k] \\ 0 & \text{otherwise.} \end{cases} \tag{2.27}$$

Comparing (2.25) to (2.3) we see a simpler calculation of $A_{rew}[k]$. More specifically, $A_{rew}[k]$ depends on only the most recently offered reward.

The resultant behaviour, $B[k]$, in (2.27), is reduced from three to two possible outcomes. The scenario in which the daughter produces the behaviour dictated by her intrinsic attitude cannot occur since negative attitude implies she does not wish to perform the activity and will do nothing if there is no external motivation. The first outcome in (2.27) describes the behaviour when an offer is accepted and the second describes the behaviour when the offer is declined.

The indicator variables remain unchanged and are restated below. The requirement that $B_d[k] > 0$ implies that $B_{sgn}[k] \neq 0$.

$$B_{sgn}[k] = \begin{cases} 0 & \text{if } B_d[k] = 0 \\ +1 & \text{if } B[k] \geq B_d[k] > 0 \text{ or } A_{out}[k] \geq 0 \\ -1 & \text{otherwise} \end{cases} \tag{2.28}$$

$$B_{rel}[k] = \begin{cases} +1 & \text{if } B[k] \geq B_d[k] \\ -1 & \text{otherwise} \end{cases} \tag{2.29}$$

The raw dissonance pressure, $P^{CD}_{raw}[k]$, value depends on the consistencies, inconsistencies, and $B_{sgn}[k]$. The consistencies and inconsistencies are determined using the sign of the daughter's attitude towards the activity, the sign of her attitude towards the reward, and the sign of the two indicator variables. The assumptions and requirements stated earlier imply:

$$\operatorname{sgn}(A_{out}[k]) = -1$$
$$\operatorname{sgn}(A_{rew}[k]) = +1.$$

This leaves the signs of the indicator values to be determined. This is best done by examining the offer accepted and offer declined scenarios separately. Focusing first on

Table 2.1: Summary of signs of attitudes and indicator variables when attitude is negative.

	Offer Accepted	Offer Declined
$\text{sgn}(A_{out}[k])$	-1	-1
$\text{sgn}(A_{rew}[k])$	$+1$	$+1$
$\text{sgn}(B_{sgn}[k])$	$+1$	-1
$\text{sgn}(B_{rel}[k])$	$+1$	-1

scenarios in which the offer is accepted, $B_{sgn}[k]$ would take a value of $+1$ indicating that the daughter performs the behaviour (i.e., jogs). As the offer is accepted, the daughter agrees to perform the desired behaviour and thus $B_{rel}[k]$ also takes a value of $+1$.

Turning to scenarios in which the offer is declined, the daughter does not produce any amount of the desired behaviour (as her attitude is negative). Thus $B_{sgn}[k]$ is set to -1 and $B_{rel}[k]$ is also set to -1. Table 2.1 summarizes the signs of attitudes and indicators when attitude is negative and rewards are accepted or declined.

The consistencies and inconsistencies in both the offer accepted and offer declined scenarios are determined using the values in Table 2.1 and (2.8)-(2.11). Table 2.2 conveniently summarizes the value of each consistency and inconsistency. The raw dissonance pressure can be calculated by substituting the summed consistency and inconsistencies in Table 2.2, and the value of $B_{sgn}[k]$, into (2.14). In doing so, we find that the dissonance pressure is simply the raw dissonance pressure:

$$P_{raw}^{CD}[k] = \begin{cases} +\dfrac{|A_{out}[k]|}{|A_{out}[k]| + |A_{rew}[k]|} & \text{if offer accepted at time } k \\ -\dfrac{|A_{rew}[k]|}{|A_{rew}[k]| + |A_{out}[k]|} & \text{if offer declined at time } k. \end{cases} \tag{2.30}$$

The original formulation of dissonance pressure (refer to (2.16)) considers attitude reversal and first-order mental processing dynamics. When attitude is negative, attitude reversal cannot occur. As such we remove it from the model. Additionally, in our initial assumptions we assume that the first-order mental processing parameter, r_2, is zero. Applying these simplifications, we obtain:

$$P^{CD}[k] = P_{raw}^{CD}[k]. \tag{2.31}$$

With the removal of attitude reversal, the change in attitude due to cognitive dissonance

Table 2.2: Summary of consistencies and inconsistencies when attitude is negative.

	Offer Accepted	Offer Declined
$M^1_{incon}[k]$	0	$\lvert A_{rew}[k]\rvert$
$M^2_{incon}[k]$	$\lvert A_{out}[k]\rvert$	0
$M^1_{con}[k]$	$\lvert A_{rew}[k]\rvert$	0
$M^2_{con}[k]$	0	$\lvert A_{out}[k]\rvert$
$M_{incon}[k]$	$\lvert A_{out}[k]\rvert$	$\lvert A_{rew}[k]\rvert$
$M_{con}[k]$	$\lvert A_{rew}[k]\rvert$	$\lvert A_{out}[k]\rvert$

effects can now be described by the two cases shown in (2.32):

$$\Delta A^{CD}_{out}[k] = \begin{cases} +K_1 \dfrac{|A_{out}[k]|}{|A_{out}[k]| + |A_{rew}[k]|} & \text{if offer accepted at time } k \\ -K_1 \dfrac{|A_{rew}[k]|}{|A_{rew}[k]| + |A_{out}[k]|} & \text{if offer declined at time } k. \end{cases} \tag{2.32}$$

Overjustification does not apply when attitude is negative. Overjustification requires an intrinsic motivation or desire to perform an activity, in our model a negative attitude denotes resentment or reluctance towards an activity. As a result, the change in attitude due to overjustification effects, $\Delta A^{OJ}_{out}[k]$, is zero. Therefore (2.23) becomes:

$$\Delta A_{out}[k] = \Delta A^{CD}_{out}[k]. \tag{2.33}$$

Combining (2.24) and (2.33), we are left with a single equation describing the dynamics of attitude:

$$A_{out}[k] = \begin{cases} A_{out}[k-1] + K_1 \dfrac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |A_{rew}[k-1]|} & \text{if offer accepted at time } k-1 \\ A_{out}[k-1] - K_1 \dfrac{|A_{rew}[k-1]|}{|A_{rew}[k-1]| + |A_{out}[k-1]|} & \text{if offer declined at time } k-1. \end{cases} \tag{2.34}$$

By assuming the attitude begins negative and remains negative we are left with a manageable system of equations that describe the dynamics of attitude given the daughter's (or more generally, anyone's) behaviour. The form of (2.34) is convenient for analysis and allows for various parameter identification techniques to be applied. The parameters of interest are μ_1, μ_2, and K_1.

2.3 Model Simplifications when Attitude is Positive

In this section we assume that the attitude begins positive and remains positive for the entire interval. The full model described in Section 2.1 is simplified under this assumption. Similar to earlier sections, aspects of the simplified model are described using the example of a father who wishes to motivate his daughter to jog.

Our initial assumption is that attitude is positive during the interval of interest. Additionally, offers contain interesting (i.e., non-zero) rewards and desired behaviour. Thus we have

$$\begin{aligned} A_{out}[k] &> 0, \\ B_d[k] &> 0, \\ R[k] &> 0, \\ &\forall k = 0, 1, 2, \ldots, N, \end{aligned}$$

where N is the number of offers to be made.

Next, we focus on simplifying the first-order mental processing dynamics. We assume that sufficient time passes between changes in attitude and subsequent offers. This means that changes in attitude at time k are fully experienced, causing attitude to settle before time $k+1$. With this assumption we can state:

$$\begin{aligned} r_1 &= 0 \\ r_2 &= 0 \\ r_3 &= 0. \end{aligned}$$

Incorporating these assumptions into (2.2)-(2.5) reduces them to:

$$A_{out}[k] = A_{out}[k-1] + \Delta A_{out}[k-1] \tag{2.35}$$

$$A_{rew}[k] = \mu_1 R[k-1] \tag{2.36}$$

$$BI[k] = A_{out}[k] + A_{rew}[k] \tag{2.37}$$

$$B[k] = \begin{cases} B_d[k] & \text{if } BI[k] \geq B_d[k] \text{ and } A_{out}[k] \leq B_d[k] \\ A_{out}[k] & \text{if } (BI[k] < B_d[k] \text{ and } A_{out}[k] \geq 0) \\ & \quad \text{or } A_{out}[k] > B_d[k]. \end{cases} \tag{2.38}$$

There are two notable changes in (2.35)-(2.38). First, $A_{rew}[k]$ in (2.36) now depends only on the most recent reward rather than previous attitudes as in (2.3). Second, the resultant behaviour, $B[k]$, in (2.38) is always non-zero. In the full model, it was possible that the daughter would not jog at all (e.g., when her attitude is highly negative, and the reward offered is small). By assuming the daughter has a positive attitude, the model dictates that she will invariably jog. This follows from the idea that a positive attitude towards an activity represents an intrinsic motivation to perform the activity.

Building on the certainty that the daughter will always jog for some duration, the indicator variable, $B_{sgn}[k]$, that tracks if any behaviour is produced, is always +1. The indicator variable, $B_{rel}[k]$, remains unchanged:

$$B_{sgn}[k] = +1 \tag{2.39}$$

$$B_{rel}[k] = \begin{cases} +1 & \text{if } B[k] \geq B_d[k] \\ -1 & \text{otherwise.} \end{cases} \tag{2.40}$$

The sign of attitude towards the activity, attitude towards the reward, and indicator variables are needed to quantify the inconsistencies and consistencies present. The consistencies and inconsistencies are then used to determine the raw dissonance pressure, $P_{raw}^{CD}[k]$. Recalling our earlier assumptions and statements (namely that attitudes are positive, rewards are non-zero, and (2.39)), we can state:

$$\begin{aligned} \operatorname{sgn}(A_{out}[k]) &= +1 \\ \operatorname{sgn}(A_{rew}[k]) &= +1 \\ \operatorname{sgn}(B_{sgn}[k]) &= +1. \end{aligned}$$

The sign of $B_{rel}[k]$ is easily determined by examining each response to an offer. In scenarios where the daughter accepts an offer, she jogs enough to receive the reward and $B_{rel} = +1$. On the other hand, when she declines an offer, she still jogs but not enough to receive the reward. In this alternate scenario, B_{rel} is set to -1. The sign of $B_{rel}[k]$ is coincidentally the same as its value. Table 2.3 summarizes the relevant signs.

Applying (2.8)-(2.11) in conjunction with the values in Table 2.3 produces the consistencies and inconsistencies present in each scenario. Table 2.4 lists the consistencies and inconsistencies present in either response to an offer. The summed values in the final two rows of Table 2.4 are a result of applying (2.12) and (2.13). It is interesting to note that no inconsistencies arise when an offer is accepted. Reflecting on this, it is clear that when attitude is positive and a desirable reward is offered, a person who accepts the offer would experience no inconsistent cognitions. Substituting these summed values and $B_{sgn}[k]$ into

Table 2.3: Summary of signs of attitudes and indicator variables when attitude is positive.

	Offer Accepted	Offer Declined
$\text{sgn}(A_{out}[k])$	$+1$	$+1$
$\text{sgn}(A_{rew}[k])$	$+1$	$+1$
$\text{sgn}(B_{sgn}[k])$	$+1$	$+1$
$\text{sgn}(B_{rel}[k])$	$+1$	-1

Table 2.4: Summary of consistencies and inconsistencies when attitude is positive.

	Offer Accepted	Offer Declined
$M^1_{incon}[k]$	0	$\lvert A_{rew}[k]\rvert$
$M^2_{incon}[k]$	0	0
$M^1_{con}[k]$	$\lvert A_{rew}[k]\rvert$	0
$M^2_{con}[k]$	$\lvert A_{out}[k]\rvert$	$\lvert A_{out}[k]\rvert$
$M_{incon}[k]$	0	$\lvert A_{rew}[k]\rvert$
$M_{con}[k]$	$\lvert A_{rew}[k]\rvert + \lvert A_{out}[k]\rvert$	$\lvert A_{out}[k]\rvert$

(2.14), the raw dissonance pressure is calculated to be:

$$P^{CD}_{raw}[k] = \begin{cases} 0 & \text{if offer accepted at time } k \\ +\dfrac{|A_{rew}[k]|}{|A_{rew}[k]| + |A_{out}[k]|} & \text{if offer declined at time } k. \end{cases} \tag{2.41}$$

Attitude reversal captures the dissonance reduction achieved by switching a positive intrinsic attitude to a highly negative attitude. Attitude reversal arises in a limited number of cases; namely, when attitude is small and positive. It may be avoided by ensuring that attitude is sufficiently large. As the primary focus of our simplifications is to analyze the dynamics of attitude change when attitude is positive, we assume attitude is sufficiently large such that reversal does not occur. In doing so, we implicitly assume at least one of the following conditions:

$$B_d[k] - BI[k] \leq \alpha_{rev} A_{out}[k],$$
$$K_1 P^{CD}[k] \leq 2A_{out}[k],$$

is met such that $r[k] = -1$.

With attitude reversal absent and the first-order mental processing parameters set to zero, the dissonance pressure is simply:

$$P^{CD}[k] = P^{CD}_{out}[k]. \tag{2.42}$$

Using (2.42) we can write the change in attitude as a result of cognitive dissonance to be:

$$\Delta A^{CD}_{out}[k] = \begin{cases} 0 & \text{if offer accepted at time } k \\ +K_1 \dfrac{|A_{rew}[k]|}{|A_{rew}[k]| + |A_{out}[k]|} & \text{if offer declined at time } k. \end{cases} \tag{2.43}$$

Reviewing (2.43), cognitive dissonance will increase or maintain attitude in scenarios which uphold the assumptions made. Referring to the example, declining her father's offer means that the daughter does not agree to jog for the duration requested, $B_d[k+1]$. We know the offer contains a desired behaviour that is larger than the daughter's behavioural intent, $BI[k+1]$. Given our assumption that attitudes are positive, the model asserts that the daughter will jog for a duration that is proportional to her current intrinsic attitude. Dissonance pressure will arise as a result of refusing or not receiving the desirable reward. The daughter reduces this pressure by increasing her intrinsic attitude. In the alternate scenario, accepting an offer means the daughter receives the reward and performs an activity she enjoys (where the activity is jogging). Her attitude and behaviour are entirely consistent; thus, dissonance pressure does not arise. It follows that (2.43) is consistent with our theory.

This is not a complete picture of how attitude changes, however. When attitude is positive, overjustification theory applies; thus, we must consider its effects. Recall that overjustification effects may only drive attitude to a minimum positive level denoted $B_t[k]$. The daughter's attitude relative to this minimum threshold is denoted $A^{rel}_{out}[k]$. The calculation of these values remains unchanged from the original model:

$$B_t[k] = \alpha_{B_d} B_d[k],$$
$$A^{rel}_{out}[k] = \max\{0, A_{out}[k] - B_t[k]\}.$$

The restriction that rewards are non-zero imply that the condition $A_{rew}[k] > 0$ in (2.20) is always satisfied. Thus (2.20) is rewritten as:

$$P^{OJ}_{raw}[k] = \begin{cases} A^{rel}_{out}[k] A_{rew}[k] & \text{if } A^{rel}_{out}[k] > 0 \text{ and } B[k] \geq B_d[k] \\ 0 & \text{otherwise.} \end{cases} \tag{2.44}$$

Note that the condition $B[k] \geq B_d[k]$ in (2.44) is equivalent to the condition that the offer is accepted. Equation (2.44) can be further streamlined by noting that when $A_{out}^{rel} \ngtr 0$, it is exactly 0. Combining these two properties, overjustification pressure may potentially arise when an offer is accepted and will never arise when an offer is declined. Equation (2.45) captures this:

$$P_{raw}^{OJ}[k] = \begin{cases} A_{out}^{rel}[k]A_{rew}[k] & \text{if offer accepted at time } k \\ 0 & \text{if offer declined at time } k. \end{cases} \tag{2.45}$$

With the first-order mental processing dynamics in (2.21) set to zero, we find that:

$$P^{OJ}[k] = P_{raw}^{OJ}[k]. \tag{2.46}$$

The final change in attitude due to overjustfication effects is rewritten as a case statement conditioned on the response to an offer:

$$\begin{aligned} \Delta A_{out}^{OJ}[k] &= \begin{cases} -K_2 P^{OJ}[k] & \text{if } K_2 P^{OJ}[k] \leq A_{out}^{rel}[k] \\ -A_{out}^{rel}[k] & \text{otherwise} \end{cases} \\ &= \max\{-K_2 P^{OJ}[k], -A_{out}^{rel}[k]\} \\ &= \begin{cases} \max\{-K_2 A_{out}^{rel}[k]A_{rew}[k], -A_{out}^{rel}[k]\} & \text{if offer accepted at time } k \\ 0 & \text{if offer declined at time } k. \end{cases} \end{aligned} \tag{2.47}$$

Examining (2.47), it is clear that overjustification only applies when the daughter accepts an offer. As K_2 is modelled to be positive, overjustification will only reduce her attitude. Further, if her attitude is not above the minimum threshold (i.e., $A_{out}^{rel}[k] = 0$), then overjustification will not lower her attitude. This is consistent with the principles of overjustfication theory.

Substituting (2.43) and (2.47) into (2.23) produces:

$$\Delta A_{out}[k] = \begin{cases} \max\{-K_2 A_{out}^{rel}[k]A_{rew}[k], -A_{out}^{rel}[k]\} & \text{if offer accepted at time } k \\ +K_1 \dfrac{|A_{rew}[k]|}{|A_{rew}[k]| + |A_{out}[k]|} & \text{if offer declined at time } k. \end{cases} \tag{2.48}$$

Finally, combining (2.35) and (2.48), we are left with a single equation describing the dynamics of attitude when attitude is positive:

$$A_{out}[k] = \begin{cases} A_{out}[k-1] & \text{if offer accepted} \\ \quad + \max\{-K_2 A_{out}^{rel}[k-1] A_{rew}[k-1], -A_{out}^{rel}[k-1]\} & \text{at time } k-1 \\ A_{out}[k-1] + K_1 \dfrac{|A_{rew}[k-1]|}{|A_{rew}[k-1]| + |A_{out}[k-1]|} & \text{if offer declined} \\ & \text{at time } k-1. \end{cases} \tag{2.49}$$

This reduced model simplifies analysis when attitude is positive. The parameters of interest in this model are μ_1, μ_2, and K_1. The following chapter presents methods to estimate these parameters.

Chapter 3

Review of Parameter Estimation Approaches

In this chapter, four well known parameter estimation approaches are considered and described. These four approaches are: least squares [13][17][20], extended Kalman filtering [3][18][14], maximum likelihood estimation [13][20][21][25], and maximum a posteriori estimation [13][25]. The properties of each approach are analysed. It is expected that any data collected will be subject to measurement noise; thus, noise is considered in the analysis of the approaches. The analysis in this chapter is used to select the approaches that are best suited to be applied to the models presented in Chapter 2.

3.1 Least Squares

The least squares estimate is the parameter estimate that produces the minimum sum of squared errors between all observations and all measurement estimates. Least squares was developed in the early 19th century. Gauss and Legendre independently published papers outlining the the basics of least squares within the first decade of the 1800s. Both publications were motivated by applications in astronomy. Today it is used in a wide variety of applications. It requires no prior knowledge of the unknown parameters which eliminates the risk of supplying incorrect or suboptimal initial conditions.

Two variants of the least squares method are considered. Section 3.1.1 presents the first method, standard least squares, while Section 3.1.2 presents the second, modified least squares.

3.1.1 Standard Least Squares

In this section the classic least squares method is presented. This method is widely used as it has many desirable properties; these properties are explained below. The derivation of least squares is commonly found in books on system identification, system modelling, and filtering. Examples of such books are [13], [17], and [20].

The standard least squares problem for an unknown scalar parameter θ is described as follows: given a set of n output measurements

$$Y = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_n \end{bmatrix}, \tag{3.1}$$

the corresponding n inputs

$$X = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix}, \tag{3.2}$$

and knowledge that the outputs follow the linear relationship

$$y_i = x_i\theta, \tag{3.3}$$

estimate θ. The estimate of θ is denoted $\hat{\theta}$. The error, e_i, is defined

$$e_i = y_i - x_i\hat{\theta} \tag{3.4}$$

$$= y_i - \hat{y}_i, \tag{3.5}$$

where

$$\hat{y}_i = x_i\hat{\theta}. \tag{3.6}$$

The error in vector form, E, is defined

$$E = Y - X\hat{\theta} \tag{3.7}$$

$$= Y - \hat{Y}, \tag{3.8}$$

where

$$\hat{Y} = X\hat{\theta}. \tag{3.9}$$

The least squares estimator produces, as its name suggests, the estimate with the smallest sum of squared errors:

$$\hat{\theta}_{LS} = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^{n} (y_i - x_i\theta)^2. \tag{3.10}$$

Equation (3.10) may be written in vector form:

$$\hat{\theta}_{LS} = \underset{\theta}{\operatorname{argmin}} (Y - X\theta)^T (Y - X\theta) \tag{3.11}$$

$$= \underset{\theta}{\operatorname{argmin}} (Y - \hat{Y})^T (Y - \hat{Y}) \tag{3.12}$$

$$= \underset{\theta}{\operatorname{argmin}}\, E^T E. \tag{3.13}$$

Equation (3.11)-(3.13) is a standard optimization problem for which the solution is known to be

$$\hat{\theta}_{LS} = (X^T X)^{-1} X^T Y, \tag{3.14}$$

when $(X^T X)^{-1}$ exists.

Properties

The least squares solution is a closed form solution. In the standard problem stated above, the estimate is unbiased; that is, the expected value of the least squares estimate is equal to the true value of the parameter (i.e., $E[\hat{\theta}_{LS}] = \theta_0$ where θ_0 is the true value of the parameter). It is easy to compute and produces a good fit [17]. Further, as the number of measurements increases, it can be shown that the variance of estimates decreases [17][20].

It is uncommon to have perfect, noise free measurements of the output. In the presence of zero mean measurement noise, v where $E[v] = 0$, that is uncorrelated to the inputs (i.e., $E[X^T v] = 0$), the least squares estimate remains unbiased. This is seen by replacing Y in the least squares solution with the noisy measurement, $\tilde{Y} = Y + v$, yielding:

$$\hat{\theta}_v = (X^T X)^{-1} X^T \tilde{Y}. \tag{3.15}$$

Taking the expectation

$$\begin{aligned}
E[\hat{\theta}_v] &= E[(X^T X)^{-1} X^T \tilde{Y}] \\
&= E[(X^T X)^{-1} X^T (Y + v)] \\
&= E[(X^T X)^{-1} X^T Y + (X^T X)^{-1} X^T v] \\
&= E[(X^T X)^{-1} X^T Y] + E[(X^T X)^{-1} X^T v] \\
&= E[(X^T X)^{-1} X^T Y] + 0 \\
&= E[(X^T X)^{-1} X^T Y] \\
&= E[\hat{\theta}_{LS}] \\
&= \theta_0
\end{aligned} \tag{3.16}$$

reveals that the estimate in the presence of zero mean uncorrelated measure noise does not affect the expected value of the estimate.

When the output, Y, depends on both the input, X, and previous outputs, the effective input is no longer uncorrelated with the noise. To illustrate this, suppose we have a dynamic system whose output can be written as

$$Y = \Phi\theta, \tag{3.17}$$

where Φ is used to indicate a dependence on both inputs and outputs. In a noiselesss situation, the least squares estimate of θ is:

$$\hat{\theta}_{LS} = (\Phi^T \Phi)^{-1} \Phi^T Y. \tag{3.18}$$

When measurement noise, v, is considered, the noisy output measurements appear in Φ; this noisy matrix is denoted Φ_v. Taking the expected value of the parameter estimate, $\hat{\theta}_v$, we obtain:

$$\begin{aligned}
E[\hat{\theta}_v] &= E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T \tilde{Y}] \\
&= E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T (Y + v)] \\
&= E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T Y + (\Phi_v^T \Phi_v)^{-1} \Phi_v^T v] \\
&= E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T Y] + E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T v].
\end{aligned} \tag{3.19}$$

Reviewing (3.19), we note that $E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T Y] \neq E[(\Phi^T \Phi)^{-1} \Phi^T Y]$ due to the noise in Φ_v. Further, $E[(\Phi_v^T \Phi_v)^{-1} \Phi_v^T v] \neq 0$ as Φ_v is now correlated with v. We conclude that $E[\hat{\theta}_v] \neq \theta_0$; that is, the estimate is biased.

3.1.2 Modified Least Squares

Modified least squares builds on the original least squares framework. It is a method used to reduce the bias in parameter estimates when noisy data is collected from dynamic systems. More information on modified least squares is available in [7] and [31].

Suppose we have a discrete dynamic system with the proper transfer function

$$\frac{Y[z]}{X[z]} = \frac{b_m z^m + b_{m-1} z^{m-1} + \cdots + b_1 z + b_0}{z^n + a_{n-1} z^{n-1} + \cdots + a_1 z + a_0} \tag{3.20}$$

where X is the input and Y is the output. In the time domain, this system is described by the difference equation:

$$\begin{aligned} y[k] = b_m x[k+m-n] + b_{m-1} x[k+m-1-n] + \cdots + b_1 x[k+1-n] + b_0 x[k-n] \\ - a_{n-1} y[k-1] - a_{n-2} y[k-2] - \cdots - a_1 y[k-n+1] - a_0 y[k-n]. \end{aligned} \tag{3.21}$$

The parameters b_i and $-a_j$ where $i = 0, 1, 2, \cdots, m$ and $j = 0, 1, 2, \cdots, n$ can be estimated using least squares. The problem is of the form

$$Y = \Phi\theta, \tag{3.22}$$

where Φ contains both the inputs and measurements of past outputs. We know from Section 3.1.1 that the least squares estimator applied to a dynamic system will produce a biased estimate when zero-mean measurement noise is present (as Φ contains noisy output measurements). Modified least squares addresses this by requiring both the inputs and output measurements to be filtered using the system characteristic dynamics prior to applying least squares [31]. More specifically, the required filter is

$$\frac{1}{A[z]} \tag{3.23}$$

where

$$A[z] = z^n + a_{n-1} z^{n-1} + \cdots + a_1 z + a_0. \tag{3.24}$$

The filtered inputs and outputs are defined:

$$\bar{Y}[z] = \frac{Y[z]}{A[z]}, \tag{3.25}$$

$$\bar{U}[z] = \frac{U[z]}{A[z]}. \tag{3.26}$$

The filtered signals are generated in the time domain as follows:

$$\bar{y}[k] = y[k-n] - a_{n-1}\bar{y}[k-1] - \cdots - a_1\bar{y}[k-n+1] - a_0\bar{y}[k-n], \tag{3.27}$$
$$\bar{u}[k] = u[k-n] - a_{n-1}\bar{u}[k-1] - \cdots - a_1\bar{u}[k-n+1] - a_0\bar{u}[k-n]. \tag{3.28}$$

The filtered inputs and outputs replace the unfiltered elements in Y and Φ to produce $\bar{Y}$ and $\bar{\Phi}$ respectively. The modified least square estimate is:

$$\hat{\theta}_{MLS} = (\bar{\Phi}^T\bar{\Phi})^{-1}\bar{\Phi}^T\bar{Y}. \tag{3.29}$$

A natural question that arises is how do we choose the filter coefficients, a_j's in (3.27)-(3.28). Given that the filter coefficients are the parameters we are trying to estimate, it follows that we do not know their value. A suitable approach is to begin by using standard least squares to produce an initial estimate of the parameters and thus the filter coefficients. After this, we filter the data using the most recently estimated filter coefficients, estimate the parameters, and come up with an improved parameter estimate. This process is iterated until the parameter estimates converge to the desired level. Convergence is not guaranteed but if it is achieved, the parameter estimate is unbiased [31].

Properties

Modified least square's most notable property is the ability to filter zero mean noise in dynamic systems to produce unbiased estimates.

3.2 Extended Kalman Filter

The Kalman filter is a well known tool for estimating the state of linear dynamic systems given noisy measurements. It is named after one of its primary developers, Rudolf E. Kalman. Kalman's original paper, [18], provides a complete derivation and many applications. Further, due to its ubiquity, there are many books exclusively on the Kalman filter and its uses (e.g., [14], and [33]). In short, the Kalman filter is known for being a recursive, computationally simple, optimal estimator. Despite its fame and success, the Kalman filter is applicable to only linear systems; thus, simultaneous state and parameter estimation is not possible when the state is multiplied by the parameter.

The Extended Kalman Filter (EKF) is an extension of the Kalman filter onto nonlinear systems. It enables nonlinear state and parameter estimation, the latter achieved by treating parameters as states. The EKF relies on a linear approximation at each time instant to approximate the dynamics of the system and then follows the standard Kalman filter equations. The remainder of this section describes the relevant equations and details required to apply the EKF. Further details and derivations are available in [3] and [14].

Suppose we are given a nonlinear dynamic system with state transition function and output function

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, u_k) + w_k, \tag{3.30}$$

$$y_k = h(\mathbf{x}_k) + v_k, \tag{3.31}$$

where the functions f and h are sufficiently smooth, $\mathbf{x}_k$ represents the state at time k, u_k represents the input at time k, y_k represents the measured output at time k, w_k represents the process noise at time k, and v_k represents the measurement noise at time k. Additionally, the process noise and measurements noise are known to be zero-mean, Gaussian, have covariance $\mathbf{Q}_k$ and $\mathbf{R}_k$ respectively, and be independent; that is:

$$E[w_k] = 0, \qquad E[w_k w_k^T] = \mathbf{Q}_k, \tag{3.32}$$

$$E[v_k] = 0, \qquad E[v_k v_k^T] = \mathbf{R}_k, \tag{3.33}$$

$$E[w_k v_k^T] = 0. \tag{3.34}$$

As functions f and h are nonlinear, they cannot be used to update the covariance predictions and estimates as in the Kalman filter. The partial derivatives of these functions evaluated at each time allow the dynamics to be approximated when the functions are sufficiently smooth. In the EKF, the Jacobian of f and h are computed at each time k and used in place of the linear state transition matrix (i.e., the matrix defining the linear relationship of a linear model):

$$\mathbf{F}_{k-1} = \left.\frac{\partial f}{\partial x}\right|_{\hat{\mathbf{x}}_{k-1|k-1}, u_{k-1}}, \tag{3.35}$$

$$\mathbf{H}_k = \left.\frac{\partial h}{\partial x}\right|_{\hat{\mathbf{x}}_{k|k-1}}. \tag{3.36}$$

The EKF recursively predicts the state at time k using the measurements available at time $k-1$, denoted $\hat{\mathbf{x}}_{k|k-1}$, and updates the prediction to an estimate when the output is measured at time k, denoted $\hat{\mathbf{x}}_{k|k}$. Similarly, the error covariance is predicted and estimated, these are denoted $\mathbf{P}_{k|k-1}$ and $\mathbf{P}_{k|k}$ respectively. The complete EKF prediction and update equations are as follows:

$$\hat{\mathbf{x}}_{k|k-1} = f(\hat{\mathbf{x}}_{k-1|k-1}, u_{k-1}) \tag{3.37}$$
$$\mathbf{P}_{k|k-1} = \mathbf{F}_{k-1}\mathbf{P}_{k-1|k-1}\mathbf{F}_{k-1}^T + \mathbf{Q}_{k-1} \tag{3.38}$$
$$\tilde{y}_k = y_k - h(\hat{\mathbf{x}}_{k|k-1}) \tag{3.39}$$
$$\mathbf{S}_k = \mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^T + \mathbf{R}_k \tag{3.40}$$
$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}_k^T\mathbf{S}_K^{-1} \tag{3.41}$$
$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k\tilde{y}_k \tag{3.42}$$
$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1}. \tag{3.43}$$

The predictions are generated using (3.37)-(3.38). Equation (3.37) shows the state prediction, $\hat{x}_{k|k-1}$, is obtained by simply substituting the most recent estimate, $\hat{x}_{k-1|k-1}$, into the state transition function (3.30). Equation (3.38) shows the predicted covariance calculation.

In order to update the state prediction to a state estimate, the EKF must first compute the innovation, the innovation covariance, and the Kalman gain. The innovation is denoted $\tilde{y}_k$, and is the difference between the measured output and the output prediction (refer to (3.39)). The innovation captures the information missed by the prediction and is later used to update the state estimate. Note that the predictions can only be updated to estimates after the measurement is made. The innovation covariance, $\mathbf{S}_k$, is calculated in (3.40). The Kalman gain, $\mathbf{K}_k$, is calculated in (3.41).

The final state estimate, $\hat{\mathbf{x}}_{k|k}$, following an output measurement is calculated in (3.42). The updated covariance estimate, $\mathbf{P}_{k|k}$, calculation is shown in (3.43).

An initial state estimate, $\hat{\mathbf{x}}_{0|0}$, and covariance, $\mathbf{P}_{0|0}$, must be supplied before the EKF can begin recursive estimation.

Properties

As the nonlinearity of the system is only described using a simple approximation, the EKF may not always work well. Performance is best when the error arising from linearization is smaller than other sources of uncertainty in the system (i.e., process or measurement noise) [14]. Higher-order approximations may reduce the error arising from linearization and thus improve performance.

The EKF is not an optimal estimator. Unlike linear systems, the transformation of Gaussian distributions is complicated and is not simply determined by means and covariances in nonlinear systems [14]. It is recommended that performance of the EKF be evaluated in simulation to ensure acceptable performance is achieved.

The selection of $\hat{\mathbf{x}}_{0|0}$, $\mathbf{P}_{0|0}$, $\mathbf{Q}_k$, and $\mathbf{R}_k$ greatly affect the behaviour and performance of the EKF. Existing knowledge of the system can be incorporated by selecting these parameters accordingly. Using existing knowledge can improve performance of the EKF (e.g., supplying an accurate initial state estimate and covariance matrix reduces time to convergence). The parameters may also be adjusted to improve numerical properties or to reduce the likelihood of divergence [3][14][33].

3.3 Maximum Likelihood Estimate

In this section, we describe maximum likelihood estimation (MLE). Additional details are available in [13], [20], [21], and [25].

The MLE approach assumes that the parameter, θ, being estimated from a set Θ specifies a set of probability distribution functions, $\{f(\theta)|\theta \in \Theta\}$. Given a set of observations $Y = \{y_1, y_2, \cdots, y_n\}$, the maximum likelihood estimate, $\hat{\theta}_{MLE}$, is the parameter that specifies the probability distribution function that maximizes the probability of producing the observations Y. More formally, this is written

$$\hat{\theta}_{MLE} = \underset{\theta \in \Theta}{\operatorname{argmax}}\, p(Y|\theta) \tag{3.44}$$

where $p(Y|\theta)$ is the called the likelihood function.

Often the log of the likelihood function is used in place of the likelihood function in (3.44). The log-likelihood function achieves its maximum at the same point as the likelihood function. Usually calculations are simpler when working with $\log p(Y|\theta)$.

Properties

In linear models of the form

$$Y = X\theta + v \tag{3.45}$$

where v is Gaussian noise of some known covariance, the maximum likelihood estimate can be shown to be equivalent to the least squares estimate [12].

MLE is a consistent estimator [25]; thus, as the number of observations increases, the probability that the estimate converges to the true value approaches one. Despite this property, it is still possible for MLE to produce a biased estimate.

3.4 Maximum A Posteriori Estimate

In this section the maximum a posteriori (MAP) estimation method is described. Further details are available in [13] and [25].

The MAP estimator utilizes observations, $Y = \{y_1, y_2, \cdots, y_n\}$, and prior knowledge of the distribution of the parameter, θ, to form its estimate. The prior distribution, denoted $f(\theta)$, is used to form the posterior distribution:

$$p(\theta|Y) = \frac{p(Y|\theta)f(\theta)}{\int_{\vartheta\in\Theta} p(Y|\vartheta)f(\vartheta)d\vartheta}. \tag{3.46}$$

As the name suggests, the MAP estimator produces an estimate, $\hat{\theta}_{MAP}$, that maximizes the posterior distribution:

$$\hat{\theta}_{MAP} = \underset{\theta\in\Theta}{\operatorname{argmax}}\, p(Y|\theta)f(\theta). \tag{3.47}$$

The denominator of (3.46) is independent of θ once evaluated and does not affect the maximization; consequently, it is not required in (3.47).

Maximizing the log of the right-hand side of (3.47) also produces the MAP estimate. This technique is often used as it simplifies the calculation of the maximum.

Properties

When the prior distribution is uniform, the MAP estimate is equivalent to to the maximum likelihood estimate [13][25].

3.5 Applicability of Approaches

In this section we consider each of the approaches described in Sections 3.1 through 3.4 for use with the models presented in Chapter 2. Ultimately, we wish to select the most promising parameter estimation approaches and study them in simulation.

Recall that our goal is to validate the model of cognitive dissonance and identify the parameters μ_1, μ_2 and K_1. These parameters vary from person to person. There are two recurring considerations when considering each approach: first, existing knowledge of the parameters, and second, the structure of the model.

Least Squares

We have minimal knowledge of the parameters μ_1, μ_2 and K_1. All that is known is that each parameter should be strictly positive. This is not an issue when using least squares as no initial estimate or prior distribution is required.

Equation (2.34) of Section 2.2 captures the effects of cognitive dissonance when attitude is negative. Reviewing (2.34) reveals that the relationship between attitude and the past attitudes, inputs (i.e., rewards), and parameters is not linear. The parameter K_1 is multiplied by the inverse of attitude plus the parameter μ_1 multiplied by a constant (i.e., $(A_{out}[k-1] + \mu_1 A_{rew}[k-2])^{-1}$). The system is best described as a dynamic non-linear system. The same can be said for the offer declined case of (2.49) the positive attitude model in Section 2.3. This situation is not ideal for least squares which works on linear systems. It is possible to rearrange some models such that the unknown parameters appear linearly. Equations (2.34) and the offer declined case of (2.49) are two such models; therefore, least squares is considered in the following chapter.

Extended Kalman Filter

The EKF requires an initial state estimate, error covariance estimate, process noise covariance, and measurement noise covariance. We have little knowledge of what to select as the initial state estimate and error covariance; fortunately we know that the initial error covariance can be made large to account for this. A drawback to bad initial estimates is the EKF takes more observations (i.e., longer) to approach the actual values. The process and measurement noise covariance should be sized with our expectation of their relative size.

The EKF can potentially produce good results when the initial estimates and noise covariances are selected appropriately. Further, the nonlinearities of (2.34) and the offer declined case of (2.49) are free of discontinuities and should be sufficiently smooth. The EKF looks promising and is considered in the next chapter.

Maximum Likelihood Estimate

In order to apply MLE, we must compute the likelihood function. Computing the likelihood function would be extremely difficult due to the nonlinear terms in the model. MLE is not considered in the following chapter.

Maximum A Posteriori Estimate

Reiterating the observations from above, we have minimal knowledge of the parameters μ_1, μ_2 and K_1. We do not have knowledge of the distribution of these parameters or even their actual value for a person. Assuming a uniform prior distribution of parameters, the most logical assumption when little is known, the MAP estimator is equivalent to MLE. As the MAP estimator provides no additional benefit over MLE, it will not be considered in the following chapter.

Chapter 4

Parameter Estimation Applied to the Behaviour Model

The goal of this chapter is to examine the behaviour of the parameter estimates provided by least squares and extended Kalman filtering. The unknown parameters in the model are μ_1, μ_2, and K_1. Equations are developed such that least squares and EKF can be applied to the models presented in Chapter 2. Following this, the error of the estimate over many iterations is studied as the number of samples, noise, and model parameters are varied in simulation. The outcomes of this chapter provide evidence to support choosing the EKF for use on experimentally collected data.

4.1 Least Squares

In this section we begin by describing how least squares may be applied to the models presented in Chapter 2. Following this, we summarize the simulation results of least squares applied to the models.

4.1.1 Development of Equations

Recall that (2.34), repeated below, describes the dynamics of attitude when attitude is negative:

$$A_{out}[k] = \begin{cases} A_{out}[k-1] + K_1 \dfrac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |A_{rew}[k-1]|} & \text{if offer accepted at time } k-1 \\ A_{out}[k-1] - K_1 \dfrac{|A_{rew}[k-1]|}{|A_{out}[k-1]| + |A_{rew}[k-1]|} & \text{if offer declined at time } k-1. \end{cases} \tag{2.34}$$

Note the similarity in structure between the offer declined case in (2.34) the offer declined case in (2.49) (repeated below):

$$A_{out}[k] = \begin{cases} A_{out}[k-1] & \text{if offer accepted} \\ \quad + \max\{-K_2 A_{out}^{rel}[k-1] A_{rew}[k-1], -A_{out}^{rel}[k-1]\} & \text{at time } k-1 \\ A_{out}[k-1] + K_1 \dfrac{|A_{rew}[k-1]|}{|A_{rew}[k-1]| + |A_{out}[k-1]|} & \begin{array}{l}\text{if offer declined} \\ \text{at time } k-1.\end{array} \end{cases} \tag{2.49}$$

Given the similarity in structure of the offer declined case in both (2.34) and (2.49), we focus our discussion (and simulation) on the negative model (2.34). In the negative model, the structure differs based on the response to an offer; as such, these are treated as separate least squares problems. We explore four ways in which least squares may be applied to (2.34): rearrange the model then apply standard least squares; rearrange the model then apply modified least squares; rearrange the model, apply ad hoc filtering, and then apply standard standard least squares; and linearize the model then apply standard least squares. Each of these is discussed below.

In order to simplify the analysis, we assume that K_1 and μ_1 are unknown and that μ_2 is known precisely.

Rearrange The Model Then Apply Standard Least Squares

The first least squares approach is to rearrange (2.34) so the parameters enter linearly. In the offer accepted case, the parameter K_1 is divided by a term including parameter μ_1:

$$A_{out}[k] = A_{out}[k-1] + K_1 \frac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \text{if offer accepted at time } k-1. \tag{4.1}$$

We may rearrange the model by multiplying both sides of (4.1) by $|A_{out}[k-1]|+|\mu_1 R[k-2]|$:

$$\begin{aligned} A_{out}[k](|A_{out}[k-1]| + |\mu_1 R[k-2]|) = A_{out}[k-1](|A_{out}[k-1]| + |\mu_1 R[k-2]|) \\ + K_1|A_{out}[k-1]| \end{aligned} \tag{4.2}$$

$$\begin{aligned} A_{out}[k]|A_{out}[k-1]| + A_{out}[k]|\mu_1 R[k-2]| &= A_{out}[k-1]|A_{out}[k-1]| \\ &\quad + A_{out}[k-1]|\mu_1 R[k-2]| \\ &\quad + K_1|A_{out}[k-1]| \end{aligned} \tag{4.3}$$

$$\begin{aligned} A_{out}[k]|A_{out}[k-1]| - A_{out}[k-1]|A_{out}[k-1]| &= A_{out}[k-1]|\mu_1 R[k-2]| \\ &\quad - A_{out}[k]|\mu_1 R[k-2]| \\ &\quad + K_1|A_{out}[k-1]| \end{aligned} \tag{4.4}$$

$$\begin{aligned} (A_{out}[k] - A_{out}[k-1])|A_{out}[k-1]| &= |\mu_1|(A_{out}[k-1] - A_{out}[k])|R[k-2]| \\ &\quad + K_1|A_{out}[k-1]| \end{aligned} \tag{4.5}$$

Equation (4.5) has the form of a least squares problem:

$$\begin{aligned} \underbrace{(A_{out}[k] - A_{out}[k-1])|A_{out}[k-1]|}_{y_k} &= \underbrace{|\mu_1|}_{\theta_1} \underbrace{(A_{out}[k-1] - A_{out}[k])|R[k-2]|}_{\phi_{k,1}} \\ &\quad + \underbrace{K_1}_{\theta_2} \underbrace{|A_{out}[k-1]|}_{\phi_{k,2}} \end{aligned} \tag{4.6}$$

where attitude measurements are made at $k = 0, 1, 2, \cdots, n$, rewards are collected at $k = 1, 2, \cdots, n$, and

$$Y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}, \tag{4.7}$$

$$\Phi = \begin{bmatrix} \phi_{1,1} & \phi_{1,2} \\ \phi_{2,1} & \phi_{2,2} \\ \vdots & \vdots \\ \phi_{n,1} & \phi_{n,2} \end{bmatrix}, \tag{4.8}$$

$$\Theta = \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix}. \tag{4.9}$$

When an offer is declined, the model uses the following dynamics:

$$A_{out}[k] = A_{out}[k-1] - K_1 \frac{|\mu_1 R[k-2]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \text{if offer declined at time } k-1 \quad (4.10)$$

which is rearranged as follows:

$$\begin{aligned} A_{out}[k](|A_{out}[k-1]| + |\mu_1 R[k-2]|) &= A_{out}[k-1](|A_{out}[k-1]| + |\mu_1 R[k-2]|) \\ &\quad - K_1 |\mu_1 R[k-2]| \end{aligned} \quad (4.11)$$

$$\begin{aligned} A_{out}[k]|A_{out}[k-1]| + A_{out}[k]|\mu_1 R[k-2]| &= A_{out}[k-1]|A_{out}[k-1]| \\ &\quad + A_{out}[k-1]|\mu_1 R[k-2]| \\ &\quad - K_1 |\mu_1 R[k-2]| \end{aligned} \quad (4.12)$$

$$\begin{aligned} A_{out}[k]|A_{out}[k-1]| - A_{out}[k-1]|A_{out}[k-1]| &= A_{out}[k-1]|\mu_1 R[k-2]| \\ &\quad - A_{out}[k]|\mu_1 R[k-2]| \\ &\quad - K_1 |\mu_1 R[k-2]| \end{aligned} \quad (4.13)$$

$$\begin{aligned} (A_{out}[k] - A_{out}[k-1])|A_{out}[k-1]| &= |\mu_1|(A_{out}[k-1] - A_{out}[k])|R[k-2]| \\ &\quad - K_1 |\mu_1||R[k-2]| \end{aligned} \quad (4.14)$$

If $K_1|\mu_1|$ is treated as a single parameter, (4.14) has the structure of a standard least squares problem:

$$\begin{aligned} \underbrace{(A_{out}[k] - A_{out}[k-1])|A_{out}[k-1]|}_{y_k} &= \underbrace{|\mu_1|}_{\theta_1} \underbrace{(A_{out}[k-1] - A_{out}[k])|R[k-2]|}_{\phi_{k,1}} \\ &\quad + \underbrace{K_1|\mu_1|}_{\theta_2} \underbrace{(-|R[k-2]|)}_{\phi_{k,2}} \end{aligned} \quad (4.15)$$

where attitude measurements are made at $k = 0, 1, 2, \cdots, n$, rewards are collected at

$k = 1, 2, \cdots, n$, and

$$Y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}, \tag{4.16}$$

$$\Phi = \begin{bmatrix} \phi_{1,1} & \phi_{1,2} \\ \phi_{2,1} & \phi_{2,2} \\ \vdots & \vdots \\ \phi_{n,1} & \phi_{n,2} \end{bmatrix}, \tag{4.17}$$

$$\Theta = \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix}. \tag{4.18}$$

As K_1 does not directly map to its own parameter, an additional step is required to obtain its estimate:

$$\hat{K}_1 = \frac{\hat{\theta}_2}{\hat{\theta}_1} \tag{4.19}$$

The resemblance of (4.5) and (4.14) to least squares problems is far from ideal for a number of reasons. First, many of the elements forming matrices Y and Φ, as shown in (4.6) and (4.15), are nonlinear functions of multiple outputs and inputs. Second, because of the rearrangement of the model, noise enters matrices Y and Φ in a nonlinear manner, noise terms are multiplied or added to other noise terms. Third, the output, Y, is no longer attitude, $A_{out}[k]$. As a result, the error being minimized does not ensure the smallest attitude error. Nonetheless, least squares is known to sometimes work well even under non-ideal conditions; as such, it is still considered a potential candidate.

Rearrange The Model Then Apply Modified Least Squares

Modified least squares is of interest when dealing with (4.6) and (4.15) in the presence of measurement noise. Recall from Sections 3.1.1–3.1.2 that least squares applied to a noisy dynamic system produces a biased estimate, and that the bias can be removed using modified least squares.

To apply modified least squares, we must select the filter coefficients that describe the system dynamics. Unfortunately, (4.15) does not have the same structure as (3.21); namely, the elements ϕ_i are based on more than just prior output measurements y_j and the elements are nonlinear. Given these challenges, we choose not to further pursue MLS.

Rearrange The Model, Apply Ad Hoc Filtering, Then Apply Standard Least Squares

An alternative to modified least squares is to apply an ad hoc low-pass filter to the input and output data. This technique often works well and is quite simple. It is commonly used in applications in which high frequency measurement noise is known to be present and filter coefficients are often chosen by intuition.

This approach is not without its own challenges. Selecting filter coefficients is troublesome in our application as little to no intuition is available. Another challenge is that the order of the filter impacts the number of measurements available to form matrices Y and Φ. This is a serious concern in situations where the number of data samples is small. Finally, it is unclear whether to apply the filter to the actual inputs and outputs (i.e., $A_{out}[k]$ and $R[k]$) or to the elements forming the Y and Φ matrices. Given these challenges, we do not consider low-pass filtering in the following chapter[1].

Linearize The Model Then Apply Standard Least Squares

The next method to apply least squares is to linearize about an initial attitude, A_0, and an initial reward, R_0, and use changes about these initial values to describe the dynamics. These changes are defined:

$$\Delta A_{out}[k-1] = A_{out}[k-1] - A_0 \tag{4.20}$$

$$\Delta R[k-2] = R[k-2] - R_0. \tag{4.21}$$

Equation (2.34) states that when offers are accepted:

$$A_{out}[k] = A_{out}[k-1] + K_1 \frac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \text{if offer accepted at time } k-1. \tag{4.22}$$

[1]This approach was tried in preliminary simulations with disappointing results. Even in the best simulations, parameter convergence never occurred unless the measurement noise was chosen to be zero.

Linearizing (4.22) about A_0 and R_0 and examining the changes in attitude about these initial values produces:

$$\begin{aligned}\Delta A_{out}[k] &= \Delta A_{out}[k-1] \\ &\quad + K_1 \frac{\partial}{\partial A_{out}[k-1]} \left(\frac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \right) \Bigg|_{A_0,R_0} \Delta A_{out}[k-1] \\ &\quad + K_1 \frac{\partial}{\partial R[k-2]} \left(\frac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \right) \Bigg|_{A_0,R_0} \Delta R[k-2]\end{aligned} \tag{4.23}$$

where:

$$\begin{aligned}\frac{\partial}{\partial A_{out}[k-1]} \left(\frac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \right) &= \frac{1}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \cdot \frac{A_{out}[k-1]}{|A_{out}[k-1]|} \\ &\quad - \frac{A_{out}[k-1]}{(|A_{out}[k-1]| + |\mu_1 R[k-2]|)^2}\end{aligned} \tag{4.24}$$

$$\frac{\partial}{\partial R[k-2]} \left(\frac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |\mu_1 R[k-2]|} \right) = \frac{-|A_{out}[k-1]}{(|A_{out}[k-1]| + |\mu_1 R[k-2]|)^2} \cdot \frac{\mu_1^2 R[k-2]}{|\mu_1 R[k-2]|} \tag{4.25}$$

Equation (4.23) can be written:

$$\begin{aligned}\Delta A_{out}[k] &= \Delta A_{out}[k-1] \\ &\quad + \underbrace{K_1 \left(\frac{1}{|A_0| + |\mu_1 R_0|} \cdot \frac{A_0}{|A_0|} - \frac{A_0}{(|A_0| + |\mu_1 R_0|)^2} \right)}_{\theta_1} \Delta A_{out}[k-1] \\ &\quad + \underbrace{K_1 \left(\frac{-|A_0|}{(|A_0| + |\mu_1 R_0|)^2} \cdot \frac{\mu_1^2 R_0}{|\mu_1 R_0|} \right)}_{\theta_2} \Delta R[k-2]\end{aligned} \tag{4.26}$$

where θ_1 and θ_2 are the parameters to be identified. Once these parameters are identified, we may solve for K_1 and μ_1 using the definitions of θ_1 and θ_2 shown in (4.26).

This approach is expected to work well only when small changes in attitude and reward are present. The error arising from the linearization increases the further we move away from the initial values A_0 and R_0. This error can be reduced by reevaluating the linearization at each time, similar to the EKF, or when the change in attitude is large. Given the lack of literature on this method and the similarity to the EKF, which is well studied, we opt not to simulate this formation.

4.1.2 Simulation Results

In this section we outline the simulation setup and parameters. Following this, least squares is applied to the rearranged model in simulation and the results of a single simulation are presented. Finally, Monte Carlo analysis is performed over 1,000,000 trials and the error is used as an indicator of the least squares estimator's performance.

Simulation Environment

As in Section 4.1.1, we treat the offer accepted case and offer declined case separately. As a result, a set of simulations is presented with the assumption that all offers are accepted while a separate set of simulations show the behaviour of the estimator when all offers are declined.

In order to apply least squares we must measure and collect data from the model. The nominal model is simulated over N time steps using randomly selected parameters, and the input and outputs generated in simulation are stored. Following this, noise is added to the recorded outputs to mimic noisy output measurements. These artificially noisy output measurements are used to compute the least squares estimate. This approach allows us to vary the measurement noise and examine the estimate's behaviour with respect to the magnitude of noise and number of iterations of the model N. The parameters are randomly selected from ranges we expect to see in the experiment (refer to Chapter 5); these ranges are listed in Table 4.1. In addition, Table 4.1 also shows the units assumed for each parameter.

In an experimental environment we expect to go through 10 iterations in a 1 hour study session. Assuming the initial attitude (i.e., the output) is known, we will have data from 11 time instances. The quantity of data available in our application is much less than many other estimation problems (which typically allow for thousands of measurements to be made with ease). The properties of least squares (see Section 3.1.1) state that error and variance of the parameter estimate decrease as the number of samples increases; the extent of this reduction can be explored in simulation. We calculate the estimate using the data from 5, 10, and 20 iterations and compare the accuracy of each estimate.

Table 4.1: Parameter values used in least squares simulation. *Note: The reward is chosen such that it is large enough for the offer to be accepted or small enough to be declined as required.

Parameter	Range of Values	Units
N	$\{5, 10, 20\}$	–
K_1	$[0.01, 100.01]$	s
μ_1	$[4.5, 100]$	s/\$
μ_2	1	s/mm
$A_{out}[0]$	$[-80, -40]$	s
$R[k]$	$[0, 20]^* \, \forall k$	\$
$\lvert\text{noise}\rvert$	$[-5, 5]$	s
$B_d[k]$	$10 \forall k$	s

Attitude measurements are expected to be noisy in experimental environments. All measurement methods are affected by a participant's cognitions and by the accuracy of the measurement method. For the analysis in this section, the simulated noise includes both sources and does not differentiate between them. The measurement noise used is uniformly distributed with zero mean. The maximum magnitude of the noise is selected uniformly from the range listed in Table 4.1.

Negative Attitude

We begin with the offer accepted case and simulate the dynamics defined by (4.1). We use the following parameters: $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, and $B_d[k] = 10$. Applying the standard least squares formation from (4.6) produces the parameter estimates shown in Table 4.2. As noise increases in magnitude the parameter estimates suffer immensely. The least squares parameter estimates are far from the actual values and do not even have the correct sign when the magnitude of the noise is set to 1. This amount of noise equates to less than 2% noise in the measurement.

Next, we simulate the offer declined case of the negative attitude model; the dynamics are defined in (4.10). With the exception of the reward $R[k]$, the same parameters as above are used. The reward is reduced to $R[k] = 2$ to ensure that offers are declined. Using the standard least squares formation in (4.15) produces the parameter estimates shown in Table 4.3. The estimates are not even close when tiny amounts of noise are added (e.g.,

Table 4.2: Least squares parameter estimates when all offers are accepted with $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, and $B_d[k] = 10$.

N	Parameter	$\lvert$noise$\rvert$			
		0	0.01	0.1	1
5	$\hat{K}_1$	2.000	1.0379	0.0818	-0.0004
	$\hat{\mu}_1$	7.000	1.0111	-4.9456	-5.4329
10	$\hat{K}_1$	2.000	1.6400	0.1225	-0.0058
	$\hat{\mu}_1$	7.000	4.8069	-4.4856	-5.2535
20	$\hat{K}_1$	2.000	1.8230	0.3890	0.0290
	$\hat{\mu}_1$	7.000	5.9515	-2.5458	-4.5983

$|\text{noise}| = 0.01$). Many estimates are negative when noise is added. Increasing N from 5 to 20 does not consistently improve the estimates.

The values presented in Table 4.2 and Table 4.3 show the least squares estimate for only two possible sets of parameters. The results are not necessarily representative of least squares' performance for all combinations of parameters. We turn to Monte Carlo analysis to better understand the behaviour of the standard least squares estimate over a wide range of parameters. Parameters are randomly selected from the ranges listed in Table 4.1 and the least squares estimate is calculated; this process is iterated 1,000,000 times. During each iteration, the error relative to the true values of K_1 and μ_1 is calculated as a percentage and stored. After all iterations are complete, the error distributions shown in Figures 4.1, 4.2, 4.3, and 4.4 are generated.

Figures 4.1 and 4.2 present the error distribution for $\hat{K}_1$ and $\hat{\mu}_1$ respectively when all offers are accepted. The first subplot of each figure shows the error distribution when the data from the first 5 time instances is used, the second subplot shows the error distribution when the data from the first 10 time instances is used, and the third subplot shows the error distribution when all the data is used. The title of each subplot also shows the mean percent error, and the proportion of iterations with less than 10% error. Reviewing Figure 4.1 we see that $\hat{K}_1$ is typically smaller than K_1. Further, the majority of observations are in the bin with edges $[-120\%, -96\%]$; reducing the bin size shows that the majority of estimates are centered around 0. It is clear that the least squares estimate $\hat{K}_1$ is biased. As the number of measurements, N, increases, the bias of the least squares estimate fails to decrease; rather, the bias increases slightly. Further, the fraction of estimates with less

Table 4.3: Least squares parameter estimates when all offers are declined with $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 2$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, and $B_d[k] = 10$.

N	Parameter	|noise|			
		0	0.01	0.1	1
5	$\hat{K}_1$	2.000	-0.1660	-0.0087	0.0168
	$\hat{\mu}_1$	7.000	-18.6620	-25.3700	-25.8163
10	$\hat{K}_1$	2.000	-0.6022	-0.0116	0.0076
	$\hat{\mu}_1$	7.000	-10.8581	-25.7042	-26.1015
20	$\hat{K}_1$	2.000	-2.5878	-0.0248	0.0149
	$\hat{\mu}_1$	7.000	-3.7234	-25.7854	-27.3074

than 10% error decreases as N increases. The observations made from Figure 4.2 regarding $\hat{\mu}_1$ are identical.

Figures 4.3 and 4.4 present the error distribution for $\hat{K}_1$ and $\hat{\mu}_1$ respectively when all offers are declined. Reviewing Figure 4.3 reveals two peaks, one around 0% error and another at -100% error. As N increases the mean error (noted in the title of each subplot) changes; however, the error distributions do not reflect this. Additionally, as N increases, the fraction of estimates that end with less than 10% error does not increase. The bins at the outer most edges indicate that a number of estimates have an absolute error of 480% or more.

Figure 4.4 shows that almost all estimates of μ_1 are negative. This is troublesome as the data is generated using positive values of μ_1. The distribution is not symmetric and has a one-sided long tail. The mean error worsens as N increases (i.e., the bias is increasing).

Overall, standard least squares applied to the rearranged model does not produce good estimates when noise is present. The error distributions shown in Figures 4.1–4.4 are simply unacceptable. Recall that we began the development of equations under the assumption μ_2 was known. It is unlikely performance will be acceptable if another parameter requires identification. Further, the set of parameters are not identical in the cases when offers are accepted and declined ($K_1|\mu_1|$ is a parameter the offer declined case). Given the poor performance with data sets in which all offers are accepted or declined, it is unlikely estimates would be better when using data sets with mixed responses. Standard least squares is not a suitable approach to identifying parameters of our model.

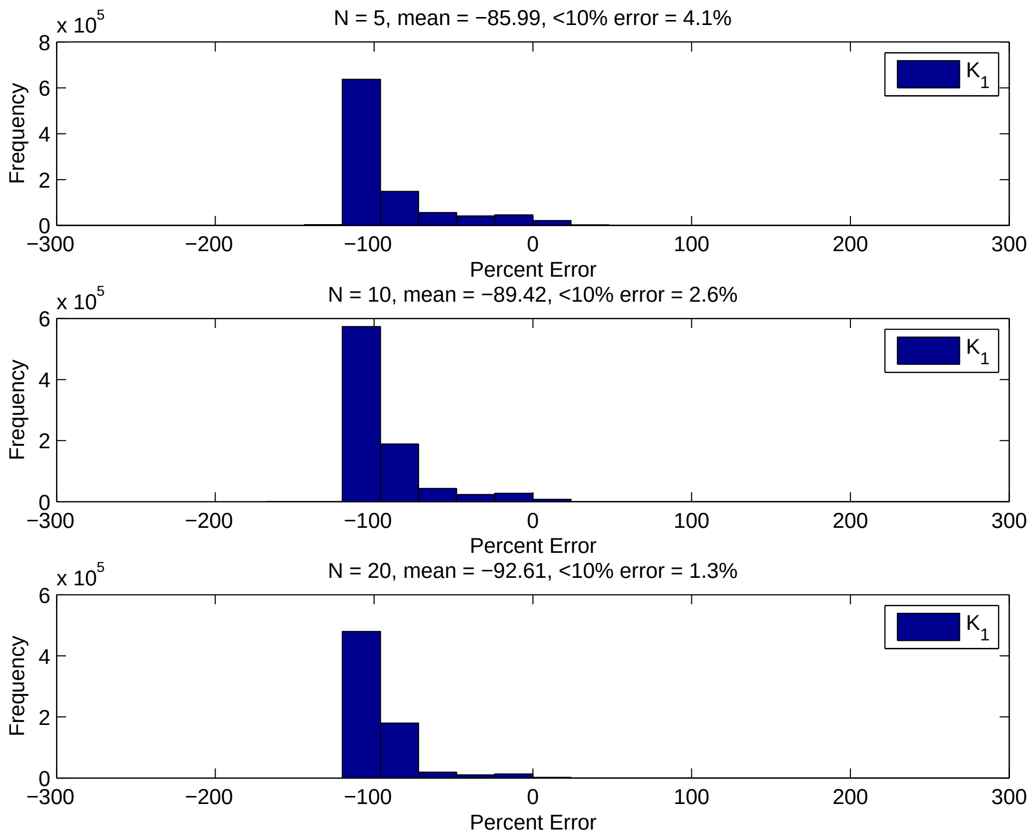

Figure 4.1: Distribution of errors of least squares K_1 estimate when offers are accepted. Simulation performed with 1 million iterations.

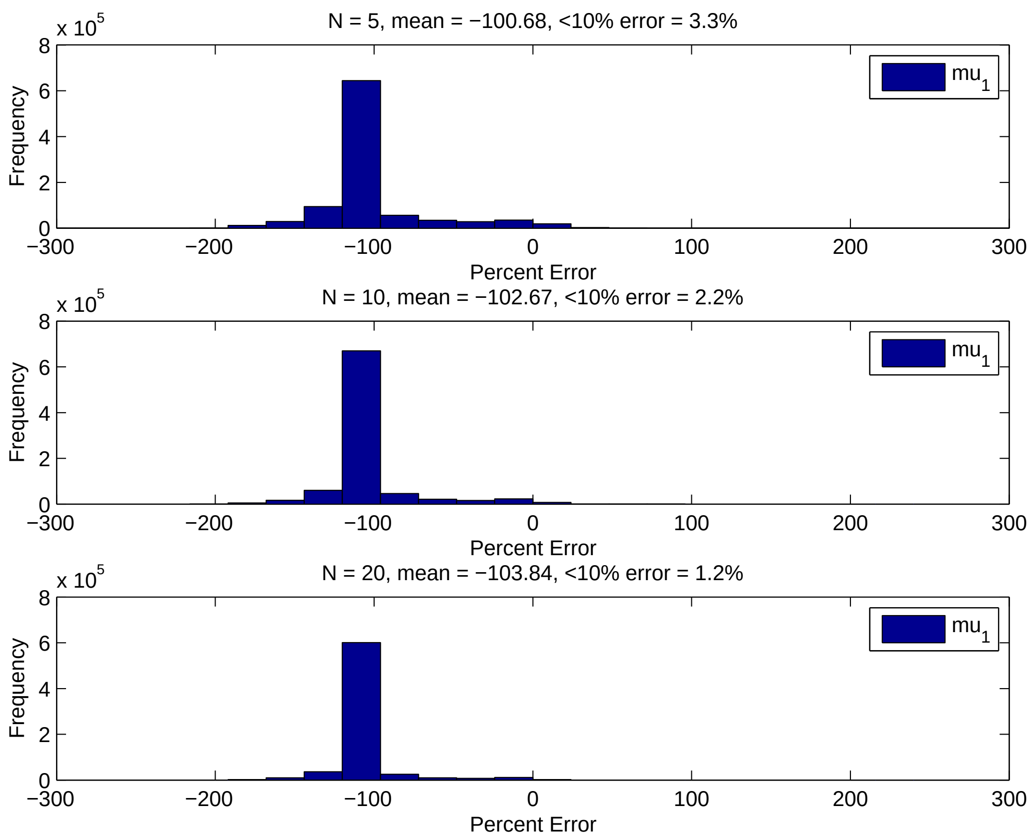

Figure 4.2: Distribution of errors of least squares μ_1 estimate when offers are accepted. Simulation performed with 1 million iterations.

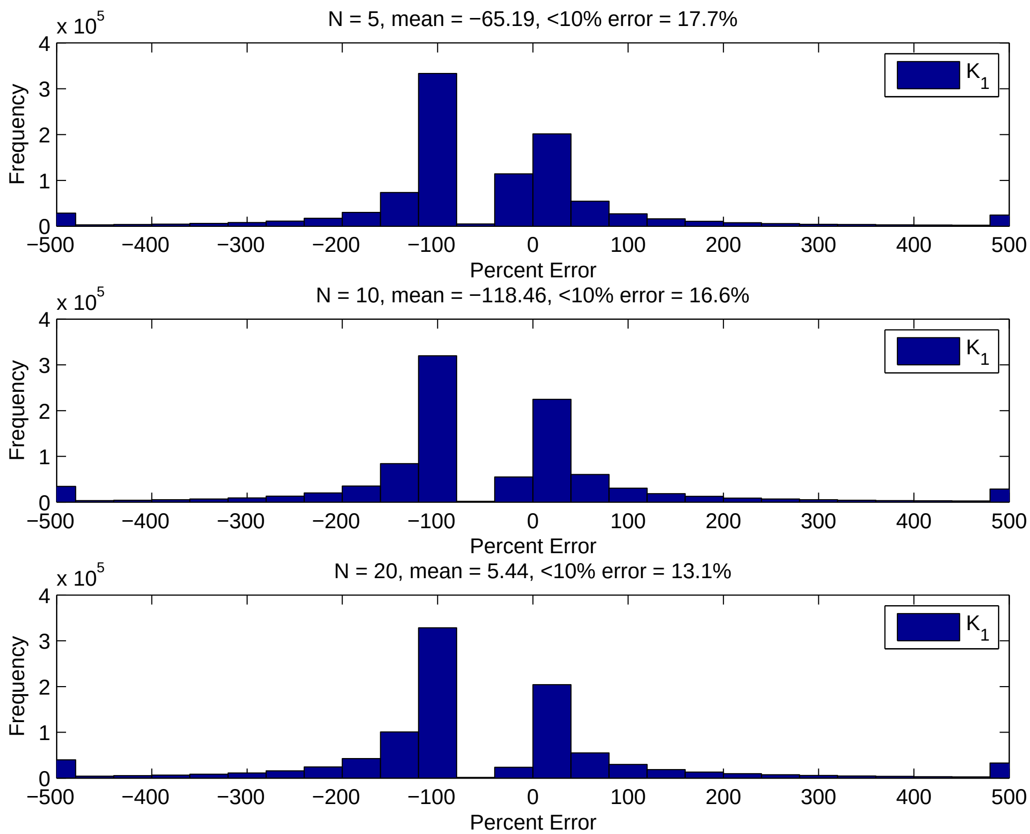

Figure 4.3: Distribution of errors of least squares K_1 estimate when offers are declined. Simulation performed with 1 million iterations.

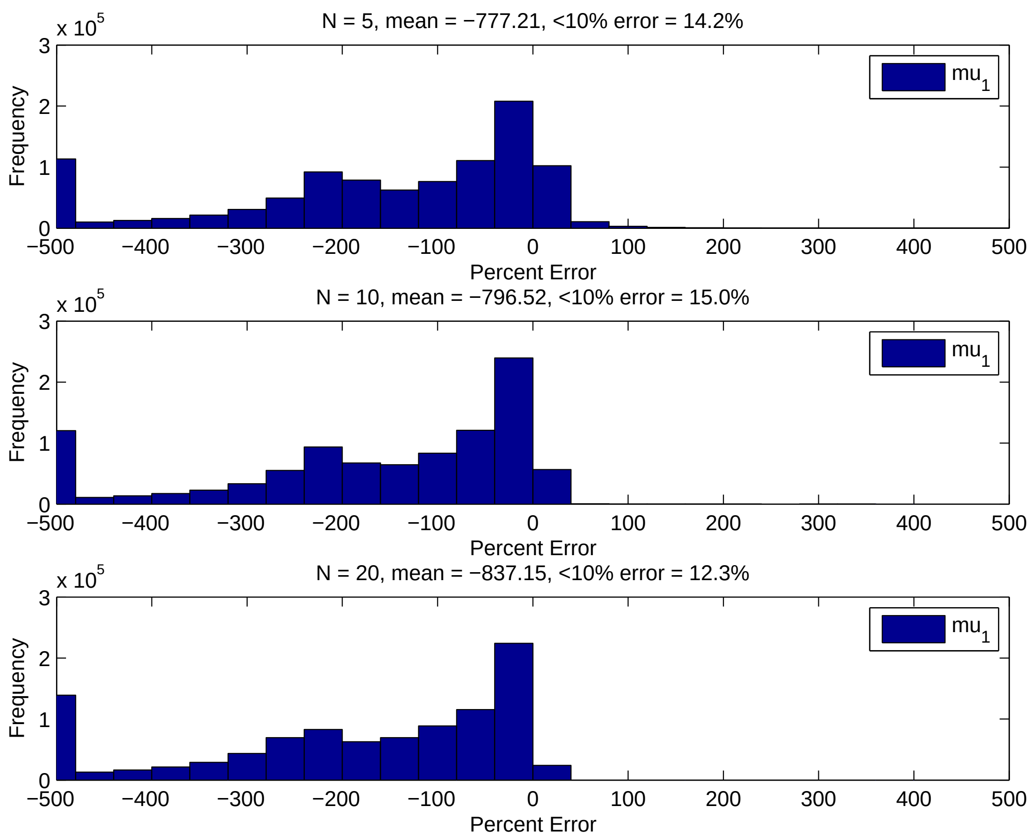

Figure 4.4: Distribution of errors of least squares μ_1 estimate when offers are declined. Simulation performed with 1 million iterations.

4.2 Extended Kalman Filter

In this section, we develop the equations such that the EKF may be applied to the simplified models presented in Chapter 2. Following this, the EKF is applied in simulation and the results are studied.

4.2.1 Development of Equations

Separate EKF equations are developed for the simplified negative and positive attitude models described in Section 2.2 and Section 2.3. The relevant equations describing the attitude dynamics in each model, (2.34) and (2.49), are repeated below for convenience.

$$A_{out}[k] = \begin{cases} A_{out}[k-1] + K_1 \dfrac{|A_{out}[k-1]|}{|A_{out}[k-1]| + |A_{rew}[k-1]|} & \text{if offer accepted at time } k-1 \\ A_{out}[k-1] - K_1 \dfrac{|A_{rew}[k-1]|}{|A_{out}[k-1]| + |A_{rew}[k-1]|} & \text{if offer declined at time } k-1. \end{cases} \tag{2.34}$$

$$A_{out}[k] = \begin{cases} A_{out}[k-1] & \text{if offer accepted} \\ \quad + \max\{-K_2 A_{out}^{rel}[k-1] A_{rew}[k-1], -A_{out}^{rel}[k-1]\} & \text{at time } k-1 \\ A_{out}[k-1] + K_1 \dfrac{|A_{rew}[k-1]|}{|A_{rew}[k-1]| + |A_{out}[k-1]|} & \text{if offer declined at time } k-1. \end{cases} \tag{2.49}$$

Common in both simplified models are the input, output, and states. The input, $u[k]$, is the reward, $R[k-1]$. The output, $y[k]$, is the output of an attitude measuring instrument, $z[k]$. Recall that attitude may be measured on some scale with arbitrary units and then scaled and converted according to (2.1)

$$A_{out}[k] = \mu_2 z[k], \tag{2.1}$$

where μ_2 is the conversion constant and $z[k]$ is the measurement output. The natural choice of state, $\mathbf{x}[k]$, is composed of attitude, the parameters we wish to estimate, and the input. In summary we have:

$$u[k] = R[k-1], \tag{4.27}$$

$$y[k] = \frac{A_{out}[k]}{\mu_2}, \tag{4.28}$$

$$\mathbf{x}[k] = \begin{bmatrix} x_1[k] \\ x_2[k] \\ x_3[k] \\ x_4[k] \\ x_5[k] \end{bmatrix} = \begin{bmatrix} A_{out}[k] \\ K_1 \\ \mu_1 \\ \mu_2 \\ u[k] \end{bmatrix}. \tag{4.29}$$

In simulation the EKF fails to produce good estimates of K_1, μ_1, and μ_2. The EKF is able to drive $\hat{y}_k \to y[k]$, but this does not necessarily imply that $\hat{\mathbf{x}}_{k|k} \to \mathbf{x}[k]$. The states $x_2[k]$, $x_3[k]$, and $x_4[k]$, which represent $\hat{K}_1$, $\hat{\mu}_1$, and $\hat{\mu}_2$ respectively, almost always diverge from the correct values. Obviously diverging estimates are unacceptable. The root cause of this behaviour stems from the lack of observability of linearized dynamics. Computing the observability matrix at each time reveals that the observability matrix is frequently less than full rank; consequently, not all states are uniquely determinable [14]. The observability matrix is a function of $\mathbf{F}_{k-1}$ and $\mathbf{H}_k$; thus, varying the state definition or the choice of outputs can improve the situation. Beyond attitude, no other outputs are available; further, all attempts to define alternate outputs based on attitude failed to make the system observable. Simulation results and attempts to define additional outputs are presented in Appendix A.

Given the lack of success with the original definitions (4.27)-(4.29), we opt to reduce the number of parameters being concurrently estimated. It is possible for μ_1 and μ_2 to be estimated through a series of questions posed to a participant in an experimental environment. These questions may be posed before the need for the values arises in any type of validation or control activity. Further details of how these parameters are estimated are presented in Section 5.1. For the remainder of this section we assume that μ_1 and μ_2 are estimated through alternate means and denote their estimates $\hat{\mu}_1$ and $\hat{\mu}_2$. The output remains the same as (4.28) while the state is redefined as follows:

$$\mathbf{x}[k] = \begin{bmatrix} x_1[k] \\ x_2[k] \\ x_3[k] \end{bmatrix} = \begin{bmatrix} A_{out}[k] \\ K_1 \\ u[k] \end{bmatrix}. \tag{4.30}$$

The linearized dynamics equations differ for each simplified model and are explained separately below.

Negative Attitude

When attitudes are negative, the state transition function, f, is

$$f\left(\mathbf{x}[k], u[k]\right) = \begin{bmatrix} x_1[k] + a[k] \\ x_2[k] \\ u[k] \end{bmatrix} \tag{4.31}$$

where $a[k]$ depends on the response to an offer:

$$a[k] = \begin{cases} +x_2[k] \dfrac{|x_1[k]|}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} & \text{if offer accepted at time } k \\ -x_2[k] \dfrac{|\hat{\mu}_1 x_3[k]|}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} & \text{if offer declined at time } k. \end{cases} \tag{4.32}$$

Clearly, the response, accepted or declined, to each offer must be known in order to select the correct dynamics. The response is easily recorded in both simulation and in experimental environments. The EKF is able to handle any sequence of responses provided attitude remains negative. This property is extremely desirable as it simplifies the design of the experiment and does not restrict the data sets which we may use. As K_1 is a constant, its dynamics in (4.31) state that the next value is simply the previous value. The output function, h, is defined as follows:

$$h\left(\mathbf{x}[k]\right) = \frac{x_1[k]}{\hat{\mu}_2}. \tag{4.33}$$

The Jacobian of the state transition function depends on the response to the offer. The superscripts A and D are used to denote the Jacobian of the accepted and declined dynamics respectively. For the accepted case,

$$\mathbf{F}_{k-1}^{A} = \left. \begin{bmatrix} \frac{\partial f_1^A}{\partial x_1[k]} & \frac{\partial f_1^A}{\partial x_2[k]} & \frac{\partial f_1^A}{\partial x_3[k]} \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \right|_{\substack{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1} \\ u[k]=u[k-1]}} \tag{4.34}$$

where:

$$\frac{\partial f_1^A}{\partial x_1[k]} = 1 + \frac{x_2[k]}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} \cdot \frac{x_1[k]}{|x_1[k]|} - \frac{x_1[k]x_2[k]}{(|x_1[k]| + |\hat{\mu}_1 x_3[k]|)^2} \tag{4.35}$$

$$\frac{\partial f_1^A}{\partial x_2[k]} = \frac{|x_1[k]|}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} \tag{4.36}$$

$$\frac{\partial f_1^A}{\partial x_3[k]} = \frac{-|x_1[k]|x_2[k]}{(|x_1[k]| + |\hat{\mu}_1 x_3[k]|)^2} \cdot \frac{\hat{\mu}_1^2 x_3[k]}{|\hat{\mu}_1 x_3[k]|} \tag{4.37}$$

and for the declined case,

$$\mathbf{F}_{k-1}^D = \left.\begin{bmatrix} \frac{\partial f_1^D}{\partial x_1[k]} & \frac{\partial f_1^D}{\partial x_2[k]} & \frac{\partial f_1^D}{\partial x_3[k]} \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right|_{\substack{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1} \\ u[k]=u[k-1]}} \tag{4.38}$$

where:

$$\frac{\partial f_1^D}{\partial x_1[k]} = 1 + \frac{x_2[k]|\hat{\mu}_1 x_3[k]|}{(|x_1[k]| + |\hat{\mu}_1 x_3[k]|)^2} \cdot \frac{x_1[k]}{|x_1[k]|} \tag{4.39}$$

$$\frac{\partial f_1^D}{\partial x_2[k]} = \frac{-|\hat{\mu}_1 x_3[k]|}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} \tag{4.40}$$

$$\frac{\partial f_1^D}{\partial x_3[k]} = \frac{-x_2[k]\hat{\mu}_1^2 x_3[k]}{(|x_1[k]| + |\hat{\mu}_1 x_3[k]|)(|\hat{\mu}_1 x_3[k]|)} + \frac{x_2[k]\hat{\mu}_1^2 x_3[k]}{(|x_1[k]| + |\hat{\mu}_1 x_3[k]|)^2}. \tag{4.41}$$

In both cases the Jacobian of the output function is:

$$\mathbf{H}_k = \left.\begin{bmatrix} \frac{1}{\hat{\mu}_2} & 0 & 0 \end{bmatrix}\right|_{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1}}. \tag{4.42}$$

Positive Attitude

When attitude is positive, only the offer declined cases invoke the effects of cognitive dissonance. Accepting an offer provides no additional information regarding K_1; as such, all offers must be declined in order to apply the EKF when attitudes are positive. The experiment must be designed such that offers are declined to ensure usable data is collected. We assume that all offers are declined in the equations that follow.

The state transition function is defined as follows when attitude is positive:

$$f\left(\mathbf{x}[k], u[k]\right) = \begin{bmatrix} x_1[k] + x_2[k]\dfrac{|\hat{\mu}_1 x_3[k]|}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} \\ x_2[k] \\ u[k] \end{bmatrix}. \tag{4.43}$$

The output function is identical to that of the negative attitude model:

$$h\left(\mathbf{x}[k]\right) = \frac{x_1[k]}{\hat{\mu}_2}. \tag{4.33}$$

The Jacobian of the state transition function is

$$\mathbf{F}_{k-1} = \left.\begin{bmatrix} \frac{\partial f_1}{\partial x_1[k]} & \frac{\partial f_1}{\partial x_2[k]} & \frac{\partial f_1}{\partial x_3[k]} \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right|_{\substack{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1} \\ u[k]=u[k-1]}} \tag{4.44}$$

where:

$$\frac{\partial f_1}{\partial x_1[k]} = 1 - \frac{x_2[k]|\hat{\mu}_1 x_3[k]|}{\left(|x_1[k]| + |\hat{\mu}_1 x_3[k]|\right)^2} \cdot \frac{x_1[k]}{|x_1[k]|} \tag{4.45}$$

$$\frac{\partial f_1}{\partial x_2[k]} = \frac{|\hat{\mu}_1 x_3[k]|}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} \tag{4.46}$$

$$\frac{\partial f_1}{\partial x_3[k]} = \frac{x_2[k]\hat{\mu}_1^2 x_3[k]}{|x_1[k]| + |\hat{\mu}_1 x_3[k]|} - \frac{x_2[k]\hat{\mu}_1^2 x_3[k]}{\left(|x_1[k]| + |\hat{\mu}_1 x_3[k]|\right)^2} \tag{4.47}$$

The Jacobian of the output function is also identical to that of the negative attitude model:

$$\mathbf{H}_k = \left.\begin{bmatrix} \frac{1}{\hat{\mu}_2} & 0 & 0 \end{bmatrix}\right|_{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1}} \tag{4.42}$$

4.2.2 Simulation Results

Details of the simulation environment, EKF initial conditions, and model parameters are presented in this section. Following this, individual and Monte Carlo simulations are presented and discussed in sections dedicated to the negative attitude model and positive attitude model.

Table 4.4: Parameter values used when applying the EKF in simulation. *Note: The reward value is reduced if it is not small enough for the offer to be declined.

Parameter	Range of Values $A_{out}[k] < 0$	Range of Values $A_{out}[k] > 0$	Units
N	$\{5, 10, 20\}$	$\{5, 10, 20\}$	–
K_1	$[0.01, 100.01]$	$[0.01, 100.01]$	s
μ_1	$[0.01, 100.01]$	$[0.01, 100.01]$	s/$
μ_2	$[0.2, 2.0]$	$[0.2, 2.0]$	s/mm
$A_{out}[0]$	$[-80, -40]$	$[20, 40]$	s
$R[k]$	$[0, 20]\ \forall k$	$[0, 20]^*$	$
|noise|	$[-5, 5]$	$[-5, 5]$	s
$B_d[k]$	$10 \forall k$	$200 \forall k$	s

Simulation Environment

The simulation environment is similar to the one used with the least squares estimator in Section 4.1.2. It is possible to run the EKF iteratively as the data is generated (i.e., online); however, we choose to generate the data first and then apply the EKF offline. The data is generated by simulating the model over N time steps with the parameters selected from the ranges stated in Table 4.4. In addition to the rewards and attitudes, the response to each offer is also stored. Recording the response to each offer allows the correct linearized dynamics to be chosen when applying the EKF. Uniformly distributed noise with zero mean is added to the output (this is expected to reduce occurrences of noise causing an attitude measurement to switch signs). The maximum magnitude of the noise is randomly selected from the range in Table 4.4 at the beginning of each simulation.

The EKF requires an initial state estimate, $\hat{\mathbf{x}}_{0|0}$, initial covariance estimate, $\mathbf{P}_{0|0}$, process noise covariance, $\mathbf{Q}_k$, and measurement noise covariance, $\mathbf{R}_k$. The initial state estimate is chosen to be

$$\hat{\mathbf{x}}_{0|0} = \hat{\mathbf{x}}[0] = \begin{bmatrix} \beta A_{out}[0] \\ 40 \\ R[0] \end{bmatrix} \tag{4.48}$$

where $\beta \in [0.8, 1.2]$ is randomly selected to simulate the uncertainty of the initial attitude measurement. The choice of $\hat{x}_2[0] = 40$ is arbitrary, as we have no prior knowledge of K_1. Finally, $\hat{x}_3[0]$ is chosen to be the exact input (i.e., reward) as it is fully known. The initial covariance matrix is large as to represent the uncertainty of our initial state estimate:

$$\mathbf{P}_{0|0} = \begin{bmatrix} 100 & 0 & 0 \\ 0 & 100 & 0 \\ 0 & 0 & 100 \end{bmatrix}. \tag{4.49}$$

While we can reduce the size of the $(3,3)$ element of $\mathbf{P}_{0|0}$ and increase the rate of convergence, performance was found to not differ significantly. Both the process and measurement noise covariances are assumed to be stationary (i.e., they do not change with time); as such, we drop the subscript k when referring to either covariance. The process and measurement noise covariances are chosen relative to each other; more specifically, we believe the measurement noise is larger than the process noise and select:

$$\mathbf{Q} = \begin{bmatrix} 10^{-5} & 0 & 0 \\ 0 & 5 \times 10^{-5} & 0 \\ 0 & 0 & 10^{-5} \end{bmatrix}, \tag{4.50}$$

$$\mathbf{R} = 10^{-2}. \tag{4.51}$$

A practical technique to decrease the likelihood of divergence is to artificially increase the process noise covariance [33]. Applying this technique, the $(2,2)$ element of $\mathbf{Q}$ is chosen to be larger than the other diagonal elements since we care most about the estimate of $\hat{K}_1$. The initial state estimate, initial covariance estimate, process noise covariance, and measurement noise covariance values shown in (4.48)-(4.51) are used when applying the EKF, except where otherwise stated.

Negative Attitude

We begin by discussing the results of a single simulation in which attitude remains negative and all offers are accepted. The parameters used are: $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, $B_d[k] = 10$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix}-41 & 40 & 10\end{bmatrix}$, $\hat{\mu}_1 = 7$, and $\hat{\mu}_2 = 1$. The negative attitude model EKF equations described by (4.31)-(4.42) are applied to produce the estimates in Table 4.5. The estimates produced by the EKF are excellent. When subjected to the largest amount of noise, $\hat{K}_1$ is already well within 5% of the actual value after just 5 measurements. As the number of measurements increases, the estimate worsens slightly but remains within 5%. The EKF's estimate of attitude and K_1 when $|\text{noise}| = 1$ are plotted as a function of time in Figure 4.5. Figure 4.5 shows that the estimates converge extremely quickly in this case.

When the reward is reduced such that all offers are declined, the EKF's estimates are also excellent. The estimates are listed in Table 4.6. The error of the estimate is within 5%

Table 4.5: EKF parameter estimates when attitude is negative, all offers are accepted, and with $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, $B_d[k] = 10$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix}-41 & 40 & 10\end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$.

Parameter	N	|noise|			
		0	0.01	0.1	1
$\hat{K}_1$	5	2.0014	2.0012	1.9993	1.9693
	10	2.0001	2.0011	2.0102	2.0962
	20	2.0000	2.0007	2.0075	2.0735

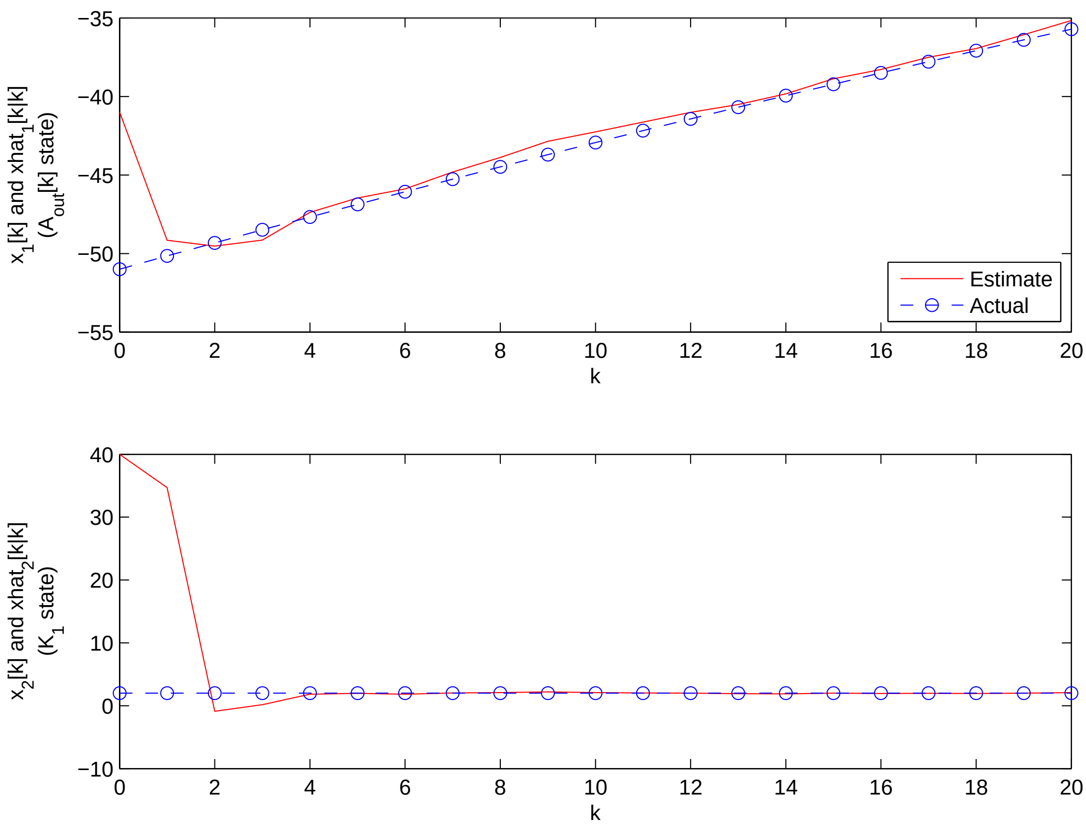

Figure 4.5: EKF estimates plotted alongside actual values generated in simulation when attitude is negative and all offers are accepted. Parameter values: $N = 20$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 2$, $|\text{noise}| = 1$, $B_d[k] = 10$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix}-41 & 40 & 10\end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$.

Table 4.6: EKF parameter estimates when attitude is negative, all offers are declined, and with $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 2$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, $B_d[k] = 10$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix}-41 & 40 & 10\end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$.

Parameter	N	|noise|			
		0	0.01	0.1	1
	5	2.0077	2.0075	2.0060	1.9756
$\hat{K}_1$	10	2.0010	1.9989	1.9804	1.7865
	20	2.0001	1.9993	1.9921	1.9174

after 5 iterations for all amounts of noise and typically remains below 5% as N increases. Figure 4.6 demonstrates, once again, that the EKF is able to rapidly converge to the actual value of K_1.

The individual simulation results presented above, while quite promising, fail to represent performance over many combinations of parameters. Monte Carlo methods are used to address this issue. At the beginning of each trial, parameters are randomly selected from the ranges listed in Table 4.4. The model is then simulated with these parameters, noise is added, the EKF is applied, and the percent error relative to the actual value of is calculated. 1,000,000 trials are performed; once all trials are simulated, the distribution of errors is examined. The EKF framework is able to handle either response to an offer and performs equally well regardless of the response; as such, the 1,000,000 trials contain both responses. It is possible that a combination of parameters may drive attitude positive during simulation. When this occurs, estimates formed using positive attitude data are discarded; however, we keep earlier estimates formed using negative attitude data from that trial (e.g., if attitude becomes positive at $k = 14$, the estimate using all 20 measurements is discarded, while the estimates using the first 5 and 10 measurements are kept).

Figure 4.7 shows the error distribution of $\hat{K}_1$ for $N = \{5, 10, 20\}$ in the first, second, and third subplots respectively. In addition to showing N, the title of each subplot contains the mean error over all trials and the fraction of all trials with less than 10% error. In this simulation $\hat{\mu}_1 = \mu_1$ and $\hat{\mu}_2 = \mu_2$ (i.e., they are known precisely). The majority of iterations produce estimates within 20% of the actual value of K_1. Further, the error distribution narrows as N is increased. Table 4.7 lists the proportion of trials that have less than 10% or 20% error as N is varied. Recall that we expect to have only 10 measurements in an experimental environment; Table 4.7 shows that we can expect 88.9% of estimates to be within 20% of the actual value when we have 10 measurements. While the EKF's estimates are great, the conditions under which they were generated are unrealistic.

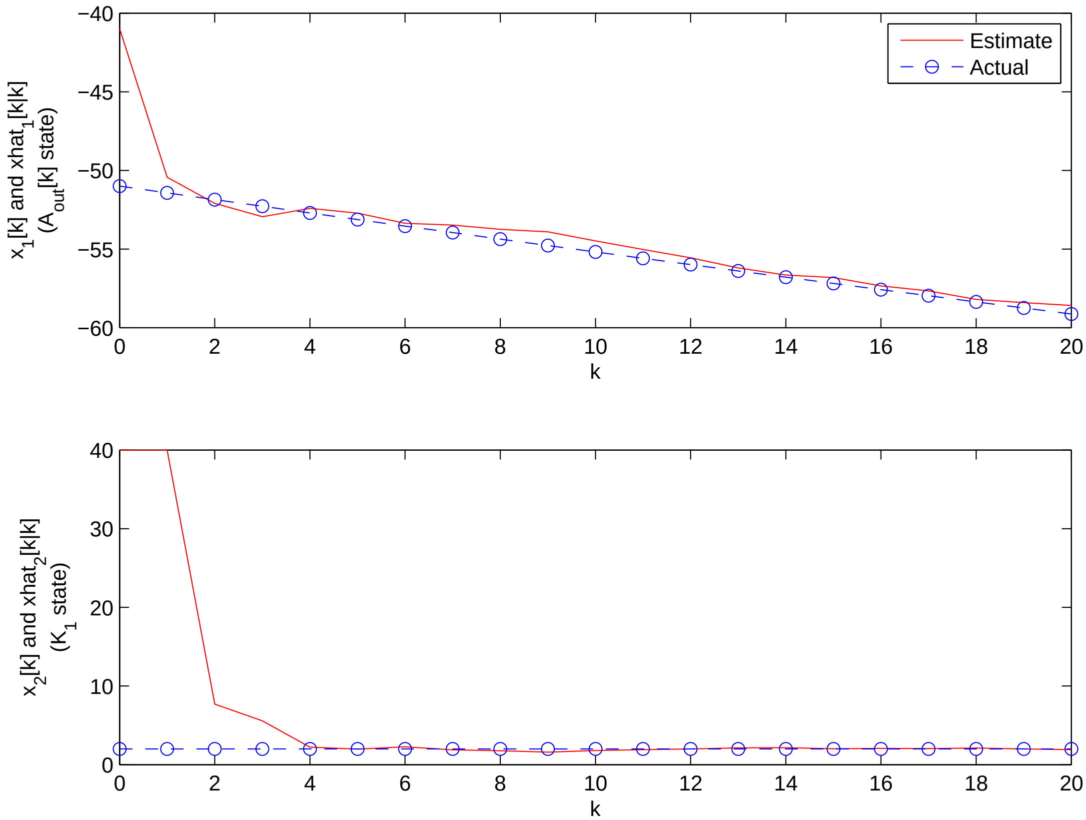

Figure 4.6: EKF estimates plotted alongside actual values generated in simulation when attitude is negative and all offers are declined. Parameter values: $N = 20$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = 1$, $B_d[k] = 10$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} -41 & 40 & 10 \end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$.

Table 4.7: Summary of EKF performance over 1,000,000 trials when attitude is negative and with randomly selected parameters when $\hat{\mu}_1 = \mu_1$ and $\hat{\mu}_2 = \mu_2$.

N	% of iterations with $\lvert\text{error}\rvert$ <10%	% of iterations with $\lvert\text{error}\rvert$ <20%
5	62.1	76.6
10	78.9	88.9
20	89.3	95.0

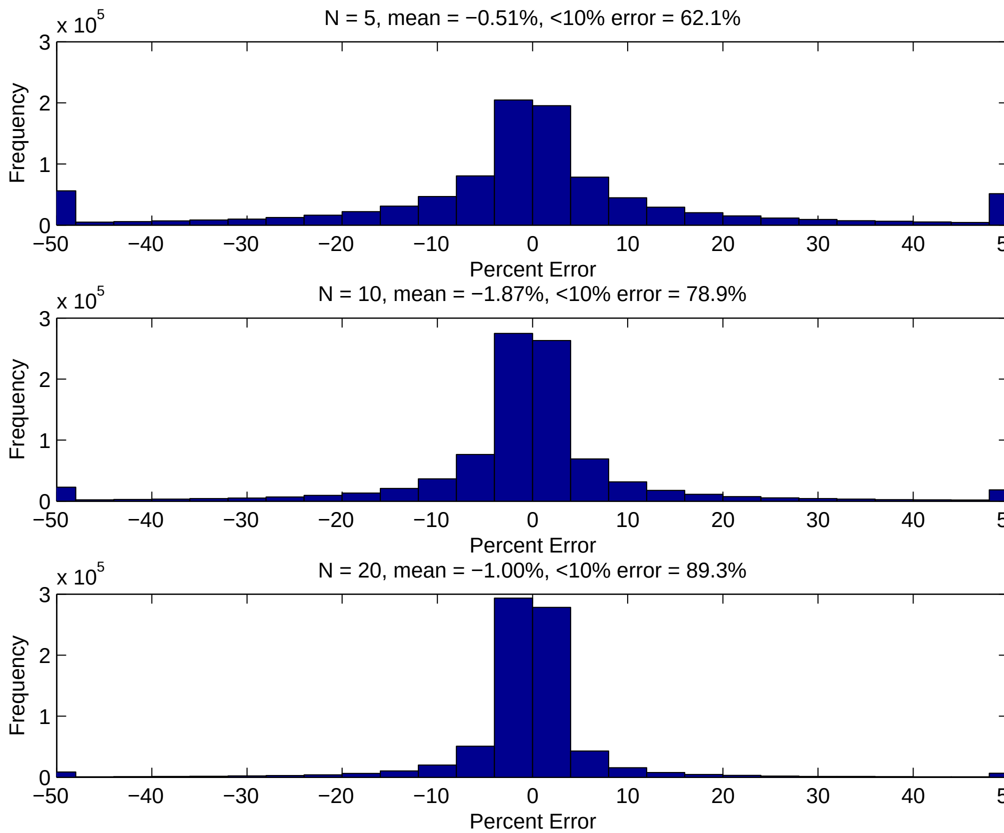

Figure 4.7: Distribution of errors of EKF estimate when attitude is negative and parameters are randomly selected from ranges listed in Table 4.4, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$. Simulation performed with 1,000,000 trials.

Table 4.8: Summary of EKF performance over 1,000,000 trials when attitude is negative and with randomly selected parameters when μ_1 and μ_2 are not known precisely.

N	% of iterations with $\lvert$error$\rvert$ <10%	% of iterations with $\lvert$error$\rvert$ <20%
5	37.9	66.9
10	43.7	77.8
20	45.6	82.9

Since knowing the exact values of μ_1 and μ_2 is highly unlikely, we add uncertainty to the estimates of μ_1 and μ_2 as follows

$$\hat{\mu}_1 = \gamma\mu_1 \tag{4.52}$$

$$\hat{\mu}_2 = \delta\mu_2 \tag{4.53}$$

where γ and δ are randomly selected from the interval $[0.8, 1.2]$ in each trial. The resulting error distribution is shown in Figure 4.8. The error distribution is considerably wider than when $\hat{\mu}_1$ and $\hat{\mu}_2$ are set to their actual values (i.e., Figure 4.7). Despite the widened error distribution, the majority of estimates are still within 20% of the actual value of K_1 and the error still decreases as N increases. More specifically, 77.8% of estimates are within 20% of the actual value after 10 measurements. Table 4.8 summarizes the performance of the estimator. The values listed in Table 4.8 are, unsurprisingly, lower than those listed in Table 4.7.

Positive Attitude

A single simulation of the positive attitude model is performed with $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = 20$, $R[k] = 0.1$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, $B_d[k] = 40$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} 16 & 40 & 0.1 \end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$. As before, the reward is constant. The reward is chosen to be small enough such that offers are declined for all 20 iterations. A suitable alternative is to dynamically reduce the size of the reward as needed to ensure offers are declined. Equations (4.33), (4.42), and (4.43)-(4.47) are used with the EKF to produce the estimates of K_1 shown in Table 4.9. The behaviour is as expected: error increases as the magnitude of noise increases, and error decreases as the number of measurements is increased. The magnitude of the noise relative to the output is larger than that used in the negative attitude model; consequently, the estimates in the final column of Table 4.9 have a larger error than the respective values in Tables 4.5 and 4.6. The EKF's estimates of

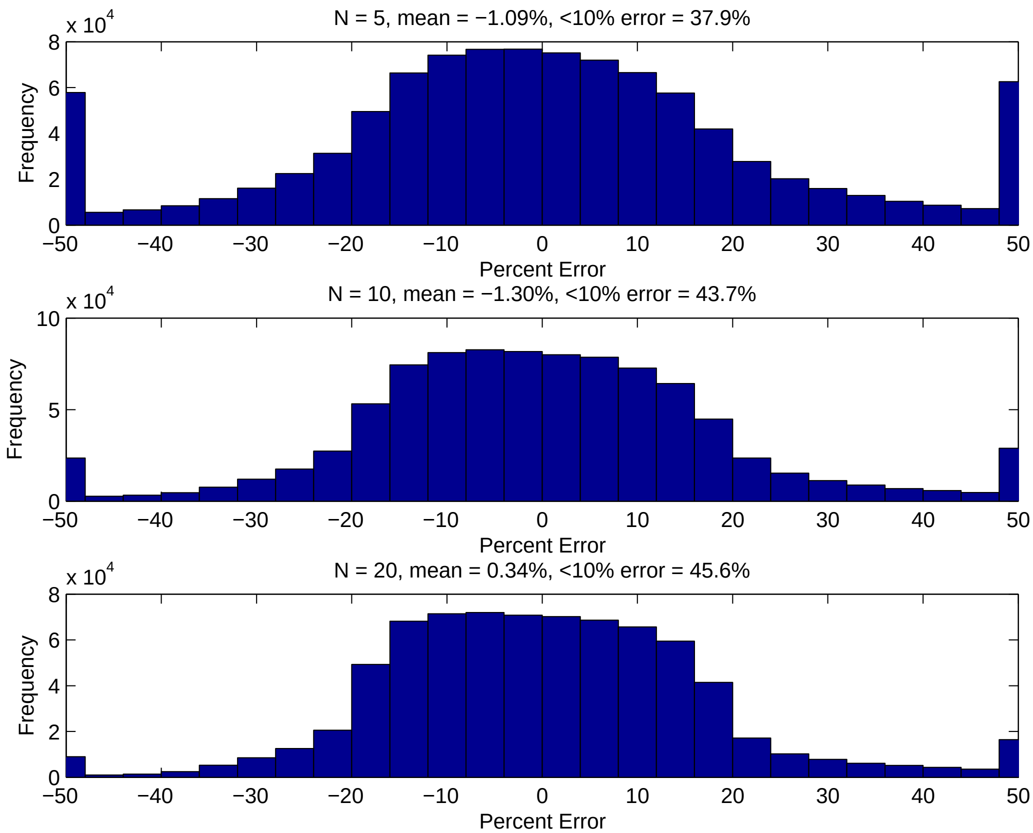

Figure 4.8: Distribution of errors of EKF estimate when attitude is negative and parameters are randomly selected from ranges listed in Table 4.4, $\hat{\mu}_1 = \gamma\mu_1$, and $\hat{\mu}_2 = \delta\mu_2$ where γ and δ are randomly selected from $[0.8, 1.2]$. Simulation performed with 1,000,000 trials.

Table 4.9: EKF parameter estimates when attitude is positive, all offers are declined, and with $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = 20$, $R[k] = 0.1$, $|\text{noise}| = \{0, 0.01, 0.1, 1\}$, $B_d[k] = 40$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} 16 & 40 & 0.1 \end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$.

Parameter	N	|noise|			
		0	0.01	0.1	1
$\hat{K}_1$	5	2.3342	2.3478	2.4687	3.2486
	10	2.0392	2.0558	2.2035	3.4870
	20	2.0046	2.0087	2.0456	2.3514

attitude and K_1 when $|\text{noise}| = 1$ are shown in Figure 4.9. Similar to the negative attitude model, the values rapidly converge.

Monte Carlo methods are used once again to simulate over a range of parameters rather than looking at an individual set of parameters. At the beginning of each trial, parameters are randomly selected from the ranges listed in the right-most column of Table 4.4. The model is simulated, measurement noise is added, the EKF is applied, and the error between $\hat{K}_1$ and K_1 is stored. After 1,000,000 trials are completed, the distribution of errors is plotted. While simulating the model, the desired behaviour is set to 100,000 to ensure that all offers are declined regardless of the parameter values selected. An alternative approach to ensure offers are declined is to reduce the size of the reward while maintaining the desired behaviour; however, this was found to reduce the change in attitude between each measurement and significantly increase the error in the estimates. While not a design constraint, this insight should be considered during the design of the experimental environment.

Figure 4.10 shows the error distribution when μ_1 and μ_2 are both known precisely. The errors are distributed across a narrow range. As shown in Table 4.10, more than 95% of estimates are within 20% of the of the true value regardless of how many measurements are used. Performance improves as the number of measurements increases.

In Figure 4.10 we assumed μ_1 and μ_2 were known. We now add up to 20% error to the estimates of μ_1 and μ_2:

$$\hat{\mu}_1 = \gamma \mu_1 \tag{4.54}$$

$$\hat{\mu}_2 = \delta \mu_2 \tag{4.55}$$

where γ and δ are randomly selected from the interval $[0.8, 1.2]$ in each trial. The resulting error distribution is shown in Figure 4.11. The error distribution is much wider than

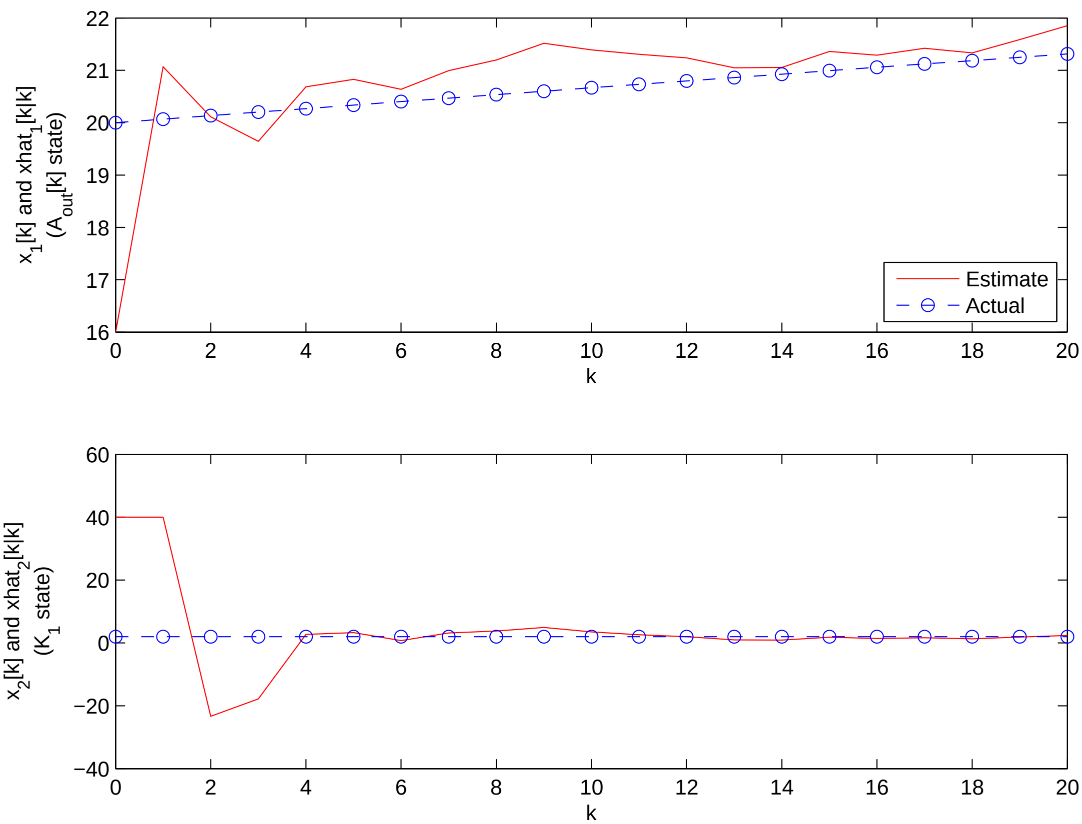

Figure 4.9: EKF estimates plotted alongside actual values generated in simulation when all offers are declined and attitude is positive. Parameter values: $N = 20$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = 20$, $R[k] = 0.1$, $|\text{noise}| = 1$, $B_d[k] = 40$, $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} 16 & 40 & 0.1 \end{bmatrix}$, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$.

Table 4.10: Summary of EKF performance over 1,000,000 trials when attitude is positive, all offers are declined, parameters are randomly selected randomly from Table 4.4, and $\hat{\mu}_1 = \mu_1$ and $\hat{\mu}_2 = \mu_2$.

N	% of iterations with $\lvert\text{error}\rvert$ <10%	% of iterations with $\lvert\text{error}\rvert$ <20%
5	93.2	96.6
10	97.3	98.6
20	98.6	99.3

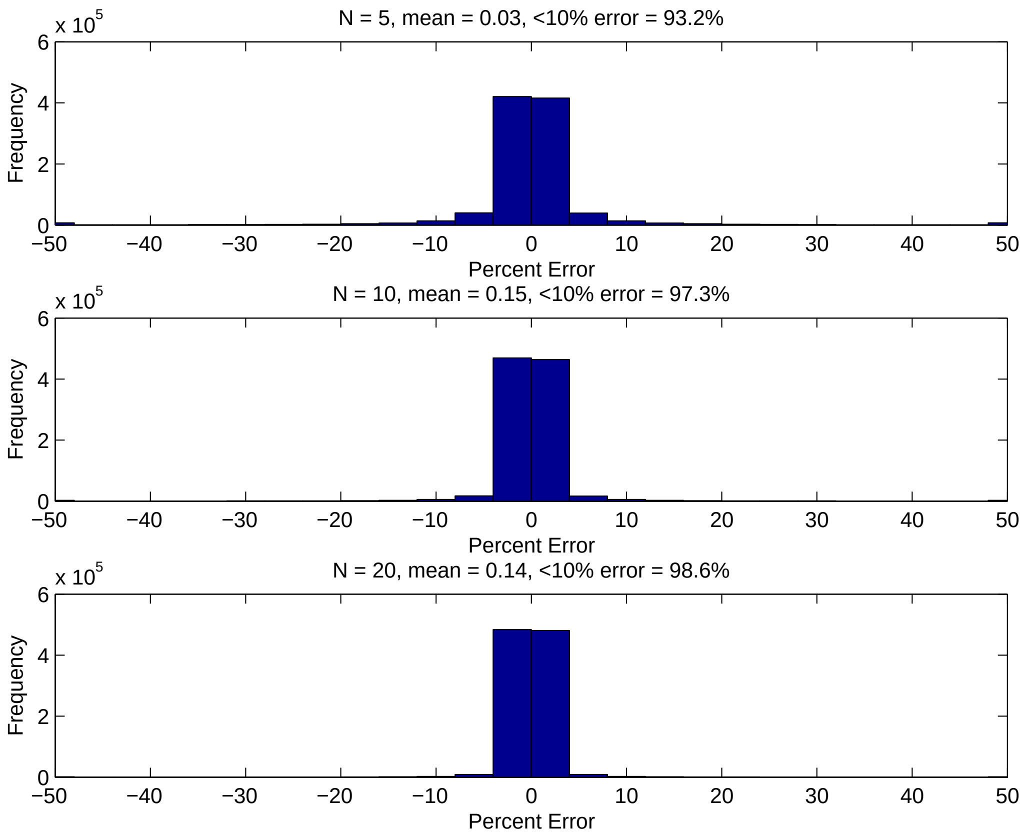

Figure 4.10: Distribution of errors of EKF estimate when attitude is positive and parameters are randomly selected from ranges listed in Table 4.4, $\hat{\mu}_1 = \mu_1$, and $\hat{\mu}_2 = \mu_2$. Simulation performed with 1,000,000 trials.

Table 4.11: Summary of EKF performance over 1,000,000 trials when attitude is positive, all offers are declined, parameters are randomly selected from ranges listed in Table 4.4, $\hat{\mu}_1 = \gamma\mu_1$, and $\hat{\mu}_2 = \delta\mu_2$ where γ and δ are randomly selected from $[0.8, 1.2]$.

N	% of iterations with |error| <10%	% of iterations with |error| <20%
5	37.8	74.9
10	36.5	72.8
20	33.8	67.7

that seen in Figure 4.10, but it remains symmetric about 0% error. With less of a peak around 0% error, it is unsurprising that the proportion of estimates within 20% error is lower. Table 4.11 summarizes the performance of the EKF when attitude is positive. As the number of measurements increases we see the opposite of what we expect, namely decreasing proportions of estimates within 10% or 20% error.

4.3 Summary of Simulations

The simulation of least squares in Section 4.1 shows that the estimates were not good. Further, while developing the specific least squares equations for our model, we assumed that the same response to all offers is given. In this simple case, the least squares estimates were unacceptable. If least squares had produced good results, there would be no guarantee that it would work well on data sets with mixed responses. It is undesirable to require that a data set have the same response to all offers for the entire set (e.g., a participant in an experimental setting would have to decline all offers). This is difficult to guarantee in an experimental setting and would certainly reduce the amount of usable data that would be collected.

This constraint is not present when applying the EKF to the negative attitude model. When the positive attitude model is considered, all offers must be declined; however, this is due to the theory behind the model rather than the choice of estimator. The EKF does, however, require that μ_1 and μ_2 be estimated through other means prior to applying the EKF. If this is done, then in both the negative and positive attitude models, the simulation results of the EKF show that more than 70% of the EKF's estimates are within 20% of the true value of K_1 after 10 measurements. We choose this approach and design our experiment accordingly.

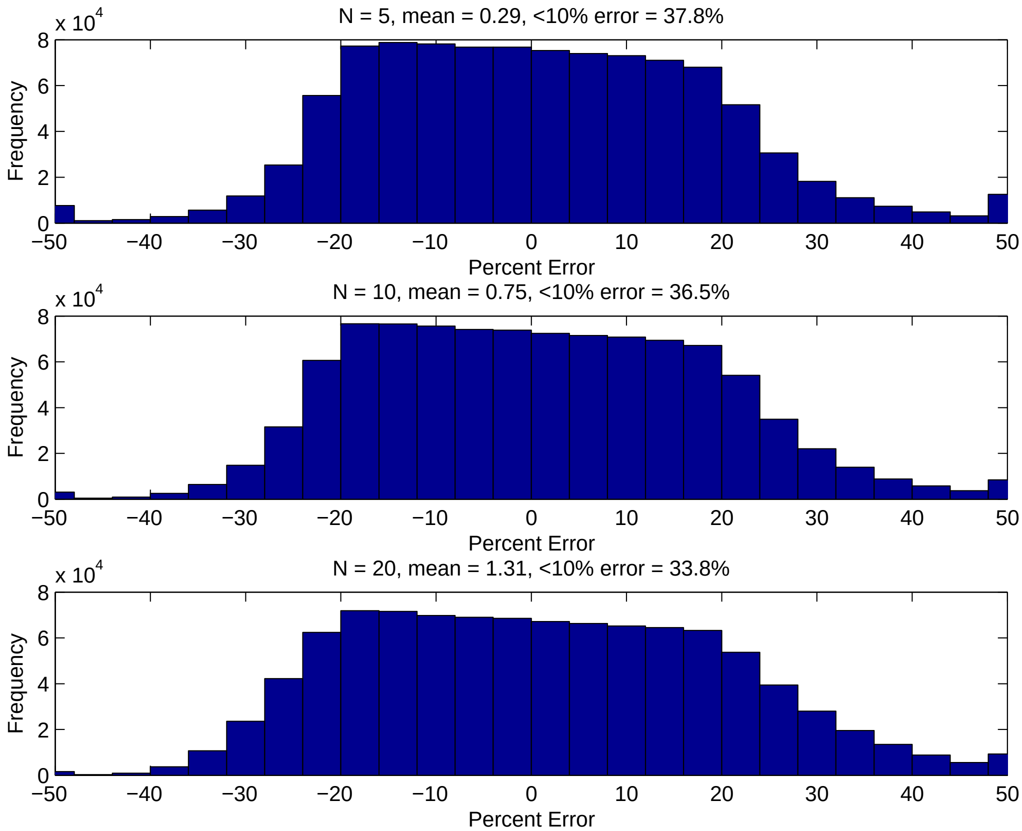

Figure 4.11: Distribution of errors of EKF estimate when attitude is positive and parameters are randomly selected from ranges listed in Table 4.4, $\hat{\mu}_1 = \gamma\mu_1$, and $\hat{\mu}_2 = \delta\mu_2$ where γ and δ are randomly selected from $[0.8, 1.2]$. Simulation performed with 1,000,000 trials.

Chapter 5

Experiment Design and Results

Validating a model requires that relevant data be collected under controlled conditions. Emphasis should be placed on minimizing external influences during data collection. A carefully designed experiment was created to validate the behaviour model presented in Section 2.1. The experiment aims to influence participants' attitude towards a sound by offering them rewards in exchange for listening to the sound.

The experiment begins with establishing ratings for eight sounds and estimating parameters needed to calculate rewards. One of the eight sounds is selected as a target sound. Participants are then offered a reward if they agree to listen to the target sound for a specified duration. This offer and reward process is iterated ten times in total and the size of the rewards and listening duration are varied at each iteration in order to excite the dynamics of the model. The flow chart in Figure 5.1 outlines the various stages of the experiment.

The validation of the model is performed by focusing on specific parts of the model one at time. As such, two variations of the experiment were run: the first focused on negative attitudes and the second focused on positive attitudes.

5.1 Experiment design

Activity

In order to validate the model an activity needed to be chosen. The criteria, in order of importance, for this activity consisted of:

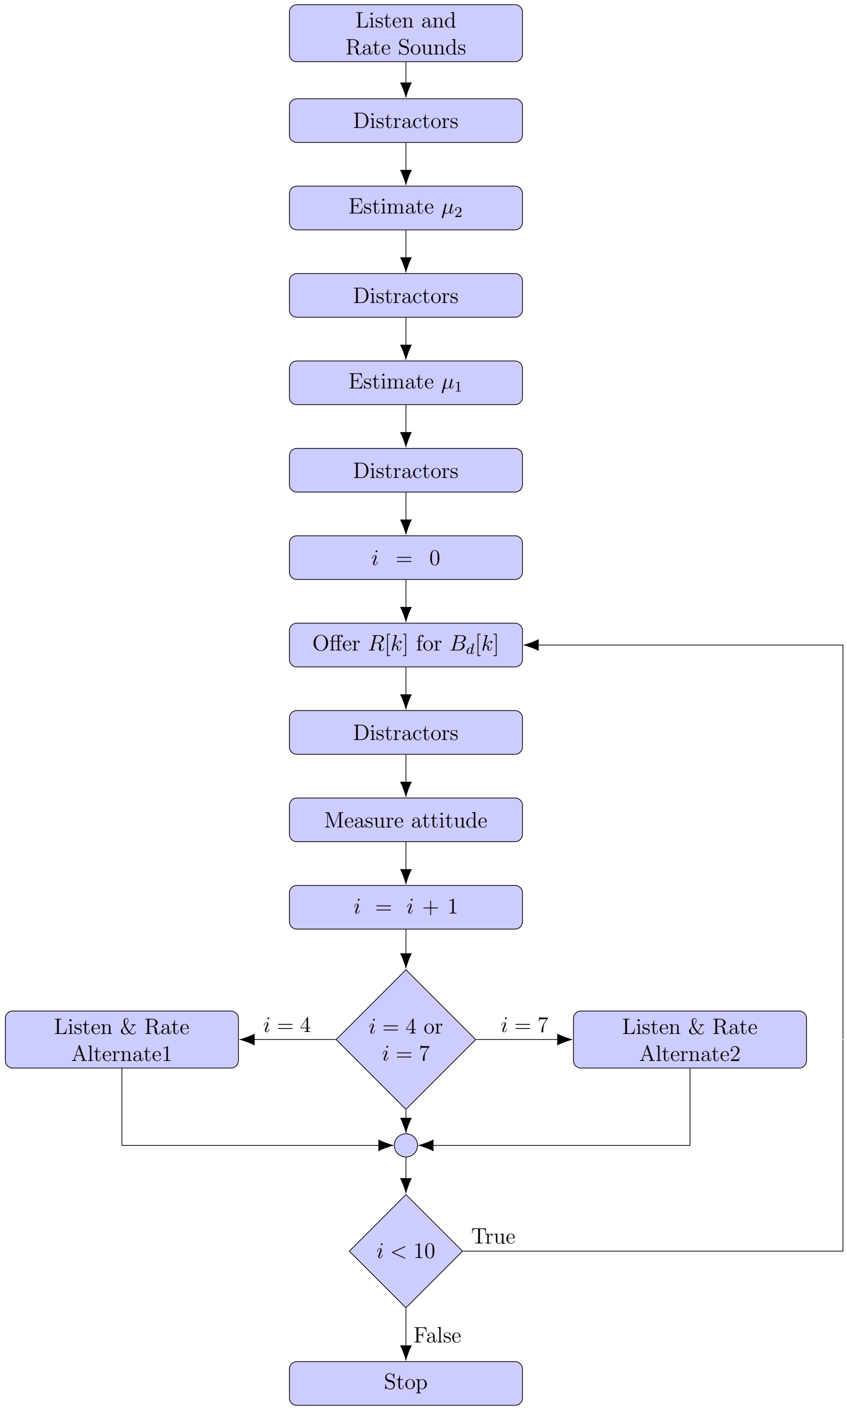

Figure 5.1: Experiment flow chart.

- performance of the activity must be quantifiable,
- attitude towards the activity must be quantifiable,
- the activity should have a small or non-existent learning curve,
- the activity should be one that any participant believes they can perform,
- the activity should not conflict with other cognitions,
- the activity should be repeatable, and
- the activity should have minimal external influences to performance (ideally the activity can be performed in a room with minimal distractions).

Listening to sounds was chosen as it satisfies all of the above criteria. Performance is easily quantified by timing, in seconds, how long a participant listens. Attitude can be measured by asking participants to rate the sound on a scale specific to the experiment. Listening to sounds is a simple activity that almost anyone can carry out. As such, it does not impose any significant restrictions on participant selection. It is expected that participants would not question their ability to perform the activity; thus, the effects of perceived behavioural control are negligible. Additionally, listening to sound clips in a study environment is unlikely to conflict with cognitions such as personal beliefs, or social norms. Playing a sound from a computer is highly repeatable. Finally, listening to sounds can take place anywhere, and requires only a computer and headphones. This allows the study to take place in an office or lab which is free from external influences.

Sounds

The first activity in the study asks participants to listen to and rate eight sounds. These sounds are Applause, Female Laughter, Sunny Day, Drill, Alien Buzzer, Longing Baby, Digital Alarm, and Ocean Wave. Refer to Appendix B for additional information about the sounds (e.g., source, author, etc.). The eight sounds provide a sufficiently large panel of sounds from which a target sound may be selected. The sounds chosen cover the range from unpleasant to pleasant.

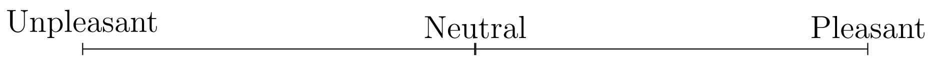

Figure 5.2: Scale upon which participants rate sounds.

Rating attitudes

The rating scale, shown in Figure 5.2, consists of a 104 mm line segment with three markers, one at either extreme and one at the center. The left most marker has "unpleasant" written above it, the center marker has "neutral," and the right most marker has "pleasant." Participants are asked to rate the pleasantness of a sound on the line scale. Ratings to the left of neutral marker are considered negative attitudes and ratings to the right of the neutral marker are considered positive attitudes. This rating is used as the measured output.

The scale has no markers other than the three described above. This is to eliminate or reduce participants' memories of past ratings. Recency effects suggest that a participant is able to best recall their most recent ratings [4]. The absence of minor gradations should diminish the participant's ability to precisely recall any rating. In doing so, the use of availability heuristics when rating sounds is reduced. An availability heuristic is a mental strategy in which participants use information that is easily recalled to influence their judgements [4]. Reducing the use of availability heuristics will decrease a participant's tendency to be consistent with their past rating behaviours (i.e., repeat the previous rating) and help obtain a true measure of their current attitude.

The distance from the neutral marker, measured in millimeters, corresponds to an attitude, measured in seconds. A challenge of the rating scale is that these two quantities are fundamentally different; as such, a conversion factor is needed to convert between the two. The conversion factor, μ_2, has units seconds/millimeter and is multiplied by the rating to convert the rating into an attitude measurement.

Estimating μ_2

After the initial rating of all eight sounds, the conversion factor μ_2 is estimated. A sound that is rated as positive or pleasant is selected, and the participant is asked, "The <positive sound> sound was one of the more pleasant sounds you rated. Listen to this sound for as long as you wish. Indicate to the researcher when you would like to begin listening and then once again when you would like to stop listening." The researcher times for how long the participant listens to the sound. The duration for which the participant listens to the positive sound is denoted T_{list} and is used with the rating of the positive sound, y_{pos}, as follows:

$$\hat{\mu}_2 = \frac{T_{list}}{y_{pos}}. \tag{5.1}$$

A positively rated sound must be available for this approach to work. Additionally, it is assumed that the rating scale is linear and μ_2 is constant across the entire range.

Estimating μ_1

Next, μ_1, the constant that relates rewards to attitudes (refer to (2.25) or (2.36)), is estimated. A sound rated as unpleasant is selected and the participant is asked the hypothetical question, "How much money would it take for you to listen to the <negative sound> sound for 60 seconds?" This question is specifically worded to ensure that no dissonance pressure arises when the participant ponders their answer. The ensures the participant's attitude remains as natural as possible. The participant's response, R_{des}, the rating to the negative sound, y_{neg}, hypothetical desired behaviour, $B_{d,hyp} = 60$ seconds, and $\hat{\mu}_2$ are used to determine $\hat{\mu}_1$:

$$\hat{\mu}_1 = \frac{B_{d,hyp} - y_{neg}\hat{\mu}_2}{R_{des}}. \tag{5.2}$$

The estimate of $\hat{\mu}_1$ is defined only if the participant responds with a positive non-zero desired reward. Additionally, at least one sound needs to have been rated as unpleasant for the hypothetical question to be valid.

In the event that a participant responds with a R_{des} that is large (typically above \$40), the researcher poses a series of questions that builds confidence in the participant's original response or more accurately estimates the value of R_{des}. The series of questions follows a binary search pattern. For example, suppose the participant responds with $R_{des} = \$50$, the researcher would lead with: "Is that to say, you would not listen to the <negative sound> sound for 60 seconds if given \$25?" If the participant would not listen to the sound for the reduced amount, the researcher stops this process and uses the original answer. If the participant indicates they would listen for the reduced amount, the researcher asks again with half of the previous offer (e.g., "Would you listen to the <negative sound> sound for 60 seconds if given \$12.50?"

Offers

The offers are carefully phrased to clearly convey the target sound, desired behaviour, and the reward. The wording was chosen to appear neutral while ensuring participants understand they have the freedom to choose. The offers in the study appear as follows: "You have the option of listening to the <target sound> sound for <behaviour desired> seconds. If you accept, you will receive $ <reward>. You may ask the researcher to play the sound if you do not recall what it sounded like." The last sentence was added to provide participants an opportunity to review the sound in the case they had forgotten it. This was a concern during the first offer as participants may not recall the names of the eight sounds they initially rated. Further, the distractors (discussed below) may hinder participants' ability to remember the name of the sound.

Rewards

Upon accepting an offer, the participant is paid while listening to the target sound. The researcher places the reward on the participant's desk. The participant keeps the rewards at the end of the study.

Distractors

One or more distractors are placed between every relevant stage of the study to allow sufficient time to pass between offers and ratings. This enables the assumption that $r_1 = r_2 = r_3 = 0$. Further, the distractors assist in hiding the true goals of the study. It is imperative that participants remain unaware of the goals of the study such that their behaviours and attitude remain as natural as possible.

Four types of distractors are used: math exercises, reading exercises, opinion exercises, and memory exercises. Each math exercise consists of four to six arithmetic operations. Reading exercises involve either reading short passages, or counting words or the number of occurrences of a letter in a passage. Opinion exercises consist of a statement and a question asking the participant to indicate their level of agreement or disagreement with the statement. Participants indicate their opinion on a scale similar to the sound rating scale. The differences in the scale are that "unpleasant" and "pleasant" are replaced with "strongly disagree" and "strongly agree" respectively. The memory exercises ask participants to recall details of earlier distractors (e.g., "In the last set of math questions, how many addition questions were there?" or "Have you read the statement on the previous page before?").

There is some semblance that distractors are related, particularly the memory exercises which require participants to recall earlier details. The opinion exercises are expected to reduce participants' memories of prior ratings. The variety of distractors should ensure that all participants are distracted such that their thoughts vary throughout the study.

Participants are also asked to listen and rate two sounds other than the target sound at two separate instances during the study. The researcher times how long the participant listens to the sounds. This serves not only as a distractor but also as an additional measurement of sound ratings.

Participant Recruitment, Selection, and Remuneration

Participant recruitment was done at the University of Waterloo. Recruitment methods included: hanging posters across campus, distributing flyers in high traffic areas, submitting a listing on the University of Waterloo's Graduate Studies website, emailing the graduate students mailing list, and emailing the electrical and computer engineering undergraduate students mailing lists.

All participants were volunteers. Participants were provided an overview of the activities they would perform during the recruitment stage as well as when they arrived to their session. This was to ensure that participants were fully aware of the tasks before they began the study. However, participants were not given complete details of the study; more specifically, the recruitment materials do not mention the offers and rewards of the study. This was done to eliminate anchoring effects. Anchoring effects capture an individual fixating on an initial piece of information, such as a the prospect of a reward or an incentive of the study [4]. Anchoring effects are detrimental to the validation effort as they influence participants' decisions regarding each offer.

The only restriction listening to sounds imposes is that participants must be able to listen to sounds through headphones. To ensure that participants were aware of this requirement, the recruitment materials state that participants will be asked to listen to sound clips. The recruitment materials can be found in Appendix D.

As an incentive to participate, participants received a slice of pizza and a beverage at the end of the study. Offering food rather than a fixed monetary incentive (e.g., a $5 gift card for participating in the study) was expected to reduce anchoring effects as it obscures the amount a participant receives. This incentive was indicated in the recruitment materials.

Recruitment proved to a be difficult task. Participant recruitment began in the middle of the Fall 2014 term and received an underwhelming response. The study continued into the Winter 2015 term. During the initial weeks of the winter term the response was better, but still much lower than expected.

Study Environment

The ideal environment is a simple, relatively empty lab or office. This type of environment minimizes external influences. The study took place in EIT 3111, a lab shared by various graduate students at the University of Waterloo. Participants performed the study seated at a desk and facing a wall. On the wall in front of the participant's seat was a painting of a flower. Being a shared space, it was not possible to fully eliminate disturbances. The seating position was chosen such that other users and equipment in the lab were typically out of sight.

Ethics

The study and related materials received full ethics clearance from the University of Waterloo's Office of Research Ethics. In accordance with ethical guidelines, participants were allowed to withdraw from the study at any time. Further, participants were fully debriefed at the end of the study. The data collected was anonymized to maintain confidentiality. This was achieved by separating is participant information from the data collected.

5.2 Experiment A - Negative Attitude

The goal of the initial experiment is to increase a participant's attitude towards a sound they initially rated as unpleasant. The assumptions and reduced model presented in Section 2.2 are realized in the study environment by incorporating the experiment design elements described in Section 5.2.1. The collected data is summarized in Section 5.2.2.

5.2.1 Design

The design elements described in this section outline specific details that are required when dealing with negative attitudes. These are in addition to the design elements described in Section 5.1.

Target Sound Selection

The researcher selects a target sound from the set of unpleasant sounds. Whenever possible, the researcher chooses a target sound whose rating falls between the -80% and -30% on

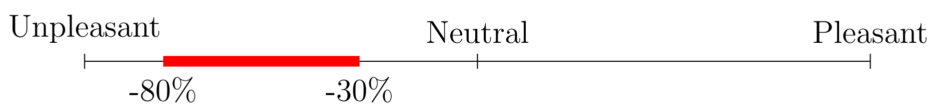

Figure 5.3: Ideal negative target sound selection range indicated by shaded portion of scale.

the scale; this is illustrated in Figure 5.3. If multiple sounds fall within this range, a target sound is chosen at random from this subset. This guideline ensures a sufficiently unpleasant target sound is chosen such that the participant's attitude remains negative for the entire study. To a lesser degree, it minimizes the likelihood that the sound is pleasant but rated unpleasant due to noise in the rating system. One final benefit to this guideline for selecting target sounds is that changes in attitude presumably stay away from saturating the low end of the scale.

Desired Behaviour and Reward Sizing

The first four offers are identical. The purpose of this is to have data that can be compared to the general trends of the model. In the case when attitude is negative and offers are accepted, the model, and theory it is based on, predicts that attitude will increase. Similarly, when offers are declined, we expect that attitudes will decrease. As the majority of applications and motivation for this work revolve around raising attitudes [24][29][8], the first four offers are sized such that, ideally, each participant accepts the offers, dissonance is nearly maximized, and near maximum attitude change is produced. Maximizing attitude change means that changes in attitude are more easily noticed and measured.

During these first four offers, the desired behaviour (measured in seconds) is held fixed at 90s. To ensure the offers at time $k = 0, 1, 2, 3$ are accepted, the reward is calculated using (2.25)-(2.27), (2.29), and the participant-specific data. To ensure the offer is accepted, (2.29) requires that $B[k+1] \geq B_d[k+1]$. Using (2.27), it follows that $BI[k+1] \geq B_d[k+1]$. Combining (2.24)-(2.26) to replace $BI[k+1]$ results in:

$$\begin{aligned} BI[k+1] &\geq B_d[k+1] \\ A_{out}[k+1] + A_{rew}[k+1] &\geq B_d[k+1] \\ A_{out}[k] + \Delta A_{out}[k] + \mu_1 R[k] &\geq B_d[k+1]. \end{aligned} \tag{5.3}$$

The change in attitude at time $k = 0$, $\Delta A_{out}[0]$, is set to 0 on the basis that the participant's attitude was not changing prior to the study. Doing so means that the first offer produces maximum dissonance, while subsequent offers may not. Recall that the

Table 5.1: Desired behaviour and reward of last six offers when attitude is negative

k	$B_d[k+1]$ (s)	$R[k]$ (\$)
4	30	2.25
5	180	0.30
6	90	0.75
7	60	2.25
8	90	1.50
9	150	1.50

measured attitude, $y[k]$, must be converted using $\hat{\mu}_2$ to attain $A_{out}[k]$. The estimate, $\hat{\mu}_1$, replaces μ_1 in (5.3). Incorporating these properties into (5.3) produces:

$$\begin{aligned} A_{out}[k] + \Delta A_{out}[k] + \mu_1 R[k] &\geq B_d[k+1] \\ \hat{\mu}_2 y[k] + \hat{\mu}_1 R[k] &\geq B_d[k+1]. \end{aligned} \tag{5.4}$$

Finally, rearranging (5.4) provides the minimum reward required for participants to accept an offer:

$$R[k] \geq \frac{B_d[k+1] - \hat{\mu}_2 y[k]}{\hat{\mu}_1}. \tag{5.5}$$

The right-hand side of (5.5) is calculated using $B_d[k+1] = 90$, the earlier calculated $\hat{\mu}_1$ and $\hat{\mu}_2$, and the participant's initial rating of the target sound, $y[0]$. The result of this calculation is rounded up to a value that simplifies payment (typically, dime, quarter, or dollar increments), and presented as the reward in the first four offers. Rounding the offer up provides a margin for error or noise produced in estimating μ_1, μ_2, or in measuring $y[0]$. Holding the four offers constant and rounding the rewards slightly decreases attitude change, contrary to the goal of maximizing attitude change. Nonetheless, these are necessary practices in the experimental environment.

The remaining six offers are identical for all participants. The goal of these offers is to excite dynamics in the model to aid in estimating K_1. With a little more effort these offers can also be compared to the model qualitatively. The values chosen are shown in Table 5.1.

5.2.2 Results

Seven participants participated in the study focusing on negative attitudes. Each participant was given a random alphanumeric identifier, namely: G2, I6, T2, V7, K4, K5, U7.

The data collected from each of these participants is summarized in this section and is presented in greater detail in Appendix C.

Model Validation

Figure 5.4 shows participant G2's data. The first of the three subplots shows measured atttitude, attitude estimated by the EKF, and the participant's response (either "Yes" or "No") to each offer. As a reminder, an offer posed at time k consists of a reward at time k, $R[k]$, and a desired behaviour at time $k+1$, $B_d[k+1]$. The attitude change that arises from the offer at time k is the change in attitude from time k to $k+1$ (i.e., $\Delta A_{out}[k] = A_{out}[k+1] - A_{out}[k]$).

Several important observations can be made while examining this subplot. First, participant G2 accepts the first four offers. Comparing the participant's attitude at time $k = 0$ to time $k = 4$, we see that attitude increases. This qualitatively follows the predictions of our model.

Next, the attitude change between time $k = 7$ and $k = 8$ does not agree with the model. The participant accepts the eighth offer at time $k = 7$; the model predicts that the participant's attitude will increase, but the data collected shows that the participant's attitude decreases. Similar observations can be made following the offers at $k = 2, 4, 5$.

Finally, as in the case at time $k = 1$, it is possible that a participant's attitude may not change at all following an offer. According to the model, attitude should remain constant only in null cases (e.g., when desired behaviour is set to zero, or when the reward is set to zero and the offer is declined).

Based on these observations, it would appear that attitude changes sometimes follow the model and other times they do not. This results in difficulty gauging the validity of the model. Measurement noise further complicates matters. Small changes in attitude may be artefacts of noise in the attitude measurement method. Examining trends over larger intervals reduces the significance of noise when noting trends; however, we are then restricted to looking at intervals in which offers are consistently accepted or declined.

In order to validate the accuracy and usefulness of the model, an objective measure must be used. In the ideal case, the model would accurately predict attitude changes all the time. For practical use, a useful model should be correct more often than not. Comparing the frequency of agreements between a participant's data and the model to the frequency of disagreements concisely captures this idea. Table 5.2 summarizes this measure for each of the participants.

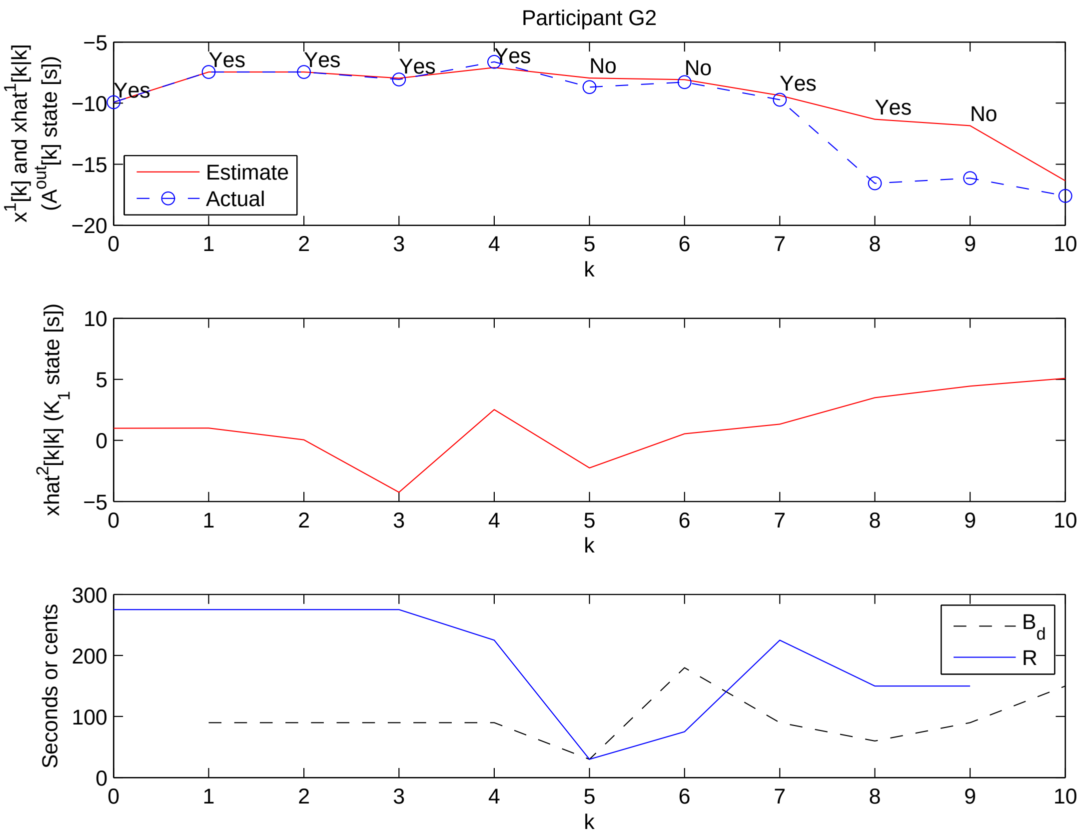

Figure 5.4: Data from Participant G2. Participant specific details: $\hat{\mu}_2 = 0.4138$ sec/mm, $\hat{\mu}_1 = 34.966$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -9.9132 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 5.0824$ sec, target sound: Digital Alarm Clock, session start time: 4:05 PM.

Table 5.2: Summary of agreement between collected data and model when attitude is negative. The final column indicates if the frequency of explicit agreements is greater than the disagreements.

Participant	# Agree	# Disagree	# No Change	# Agree > # Disagree ?
G2	5	4	1	Yes
I6	3	4	3	No
T2	5	4	1	Yes
V7	4	4	2	No
K4	3	2	5	Yes
K5	5	3	2	Yes
U7	1	0	9	Yes

The final column in Table 5.2 indicates whether the number of agreements is greater than the number of disagreements. Five of the seven participants' data showed that the model was correct more than it was not. This suggests the model has some validity; further, there is utility in the model, specifically when attitude is negative.

Identifying K_1 and Gauging the EKF's Effectiveness

The second subplot of Figure 5.4 shows the EKF's estimate of K_1. The EKF is initialized with $\hat{K}_1 = 1$. The simulations in Section 4.2.2 show that by $k = 10$, an estimate of K_1 is typically attained with less than 20% error.

The first observation in the second subplot of Figure 5.4 is that $\hat{K}_1$ is positive at larger values of k. Positive values imply that the attitude changes are consistent with the model, thus providing support for the model. As the EKF requires a number of measurements to approach the true value, it is acceptable for the initial estimates to briefly be negative. If $\hat{K}_1$ were strictly negative it would suggest that the model is not accurate.

The second observation is that $\hat{K}_1$ appears to be converging. Convergence is a strong indication that the EKF is set up well, and that the model is valid.

These two observations provide two measures, one which can be used to comment on the validity model, and another which speaks to the performance of the EKF. Table 5.3 summarizes the above two observations for each participant (full details for each participant are presented in Appendix C). It should be noted that participant U7 was not included in the summary since the measured attitude remained saturated on the measurement scale.

Table 5.3: Summary of observations of $\hat{K}_1$ when attitude is negative.

Participant	$\hat{K}_1 > 0$?	Appears to be converging?
G2	Yes	Yes
I6	No	Yes
T2	Yes	Yes
V7	Yes	Yes
K4	Yes	No
K5	Yes	No
U7	N/A	N/A

In five of the six remaining data sets, $\hat{K}_1$ was positive at the last sample; this suggests that the model is valid. Four of six estimates appear to be converging which indicates that the EKF is working effectively.

Statistical Significance

Statistical hypothesis testing can be applied to the results of the negative attitude experiment. Consider the fraction

$$x_i = \frac{\#Agree_i}{\#Agree_i + \#Disagree_i} \tag{5.6}$$

where the subscript i identifies each participant. The null hypothesis, H_0, is that the actual mean value of x_i, denoted μ, is less than or equal to 0.5. The alternate hypothesis, H_A, is the actual mean value μ is greater than 0.5. The observed mean, $\bar{x} = 0.60924$, and observed standard deviation, $\sigma = 0.184104$, over the population of 7 participants result in a t-score test statistic of $t = 1.570$. The probability of producing a t-score this extreme is 0.0837 (i.e., $p(t > 1.570) = 0.0837$). This p-value means the probability of the experimentally observed fraction mean, $\bar{x}$, given that we believe the actual mean, μ, is less than or equal to 0.5 is 0.0837. The result is significant at the $p = 0.0837$ level.

5.3 Experiment B - Positive Attitude

The second experiment aims to increase a participant's attitude towards a pleasant sound. This is done by exciting the effects of cognitive dissonance, similar to the first experiment, while not exciting overjustification effects. Section 5.3.1 describes how the assumptions in Section 2.3 are realized in the experimental environment. Section 5.3.2 summarizes the results of this experiment.

5.3.1 Design

In this section we describe the design elements specific to working with positive attitudes. These design elements are in addition to those outlined in Section 5.1.

Target Sound Selection

The researcher selects a sound initially rated as pleasant to be the target sound. The ideal target sound is rated far enough away from the "neutral" marking as to eliminate sounds that will likely cause attitude reversal to occur. Further, the ideal target sound is rated far enough away from the "pleasant" marker such that there is room to measure increases in the attitude.

Sounds rated as pleasant imply that the participant has some intrinsic desire to listen to the sound, even in the absence of a reward. If the desired behaviour in an offer is smaller the participant's behaviour intention, the participant would surely accept the offer and produce the desired behaviour or more. In doing so, they would invoke overjustification effects and not provide useful data for our purposes. In order to avoid this, the participant's attitude towards the target sound must be less than the desired behaviour contained in an offer.

A suitable guideline is to choose, when possible, a pleasant sound rated between 15% and 50% as the target sound. This region is shown in Figure 5.5. Preference should be given to sounds with lower ratings.

Desired Behaviour and Reward Selection

Again, the first four offers are held constant for a given participant. This allows the data collected to easily be compared qualitatively against the trends predicted by the model.

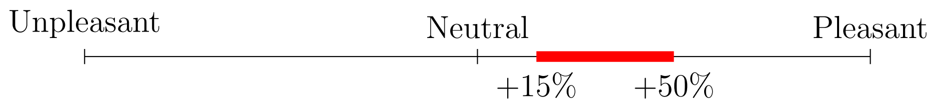

Figure 5.5: Ideal positive target sound selection range indicated by shaded portion of scale.

The offers must be sized such that participants decline each reward. This follows from (2.49) which shows only the effects of cognitive dissonance may increase attitude and (2.43) which shows cognitive dissonance is present only when an offer is declined.

As per (2.40), when a reward is declined, $B_{rel}[k] = -1$; therefore, $B[k] < B_d[k]$. Since the target sound is pleasant, the participant has a positive attitude towards listening to the sound (i.e., $A_{out}[k] > 0$). Matching these facts to the second case in (2.38) means the following conditions must hold:

$$B[k] = A_{out}[k], \tag{5.7}$$

$$BI[k] < B_d[k], \tag{5.8}$$

$$A_{out}[k] < B_d[k]. \tag{5.9}$$

Note that violating the (5.9) would contradict $B[k] < B_d[k]$ meaning the offer was not declined. Further, the target sound and desired behaviour should be selected such that these conditions are not violated by default. Using the definitions in (2.35)-(2.37) in conjunction with the above conditions produces:

$$\begin{aligned} BI[k+1] &< B_d[k+1] \\ A_{out}[k+1] + A_{rew}[k+1] &< B_d[k+1] \\ A_{out}[k] + \Delta A_{out}[k] + \mu_1 R[k] &< B_d[k+1]. \end{aligned} \tag{5.10}$$

The change in attitude at time $k = 0$, $\Delta A_{out}[0]$, is set to 0 as the participant's attitude was not changing prior to the study. The measured attitude, $y[k]$, must be converted using $\hat{\mu}_2$ to attain $A_{out}[k]$. The estimate, $\hat{\mu}_1$, replaces μ_1 in (5.10). Incorporating these properties into (5.10) produces:

$$\begin{aligned} A_{out}[k] + \Delta A_{out}[k] + \mu_1 R[k] &< B_d[k+1] \\ \hat{\mu}_2 y[k] + \hat{\mu}_1 R[k] &< B_d[k+1]. \end{aligned} \tag{5.11}$$

Finally, rearranging (5.11) provides an upper bound on the reward:

$$R[k] < \frac{B_d[k+1] - \hat{\mu}_2 y[k]}{\hat{\mu}_1}. \tag{5.12}$$

Satisfying the inequality in (5.12) causes the reward to be declined and attitude to increase. As attitude increases, the right-hand side of (5.12) decreases. This means the upper limit on the reward decreases after each successive declined offer. Since the first four offers are held constant, the reward must be sized such that (5.12) is upheld at each time. One final note is that any reward chosen will produce the largest change in attitude from $k = 0$ to $k = 1$.

To help ensure participants decline each offer, the reward is set to the right-hand side of (5.12) discounted by a factor of $\frac{1}{2}$ and rounded down. A drawback is that dissonance, and thus attitude change, is reduced from its maximum. An option that increases dissonance while still encouraging offers to be declined is to double the desired behaviour. In doing so, the offer requires even more of a participant in order to receive a reward. This option maintains a significant, meaningful reward; thus, maintaining dissonance levels and ensuring attitude change is large enough to be measured.

$R[k]$ is calculated using $B_d[k+1] \in \{90, 180\}$, $\hat{\mu}_1$, $\hat{\mu}_2$, and the participant's initial rating of the target sound, $y[0]$. The value of the desired behaviour is 90s for half of the participants and 180s for the other half.

The final six offers are composed of varying desired behaviours and small rewards that should promote participants to decline offers. The offers are varied such that the dynamics are excited and the performance of the EKF can be examined. Two different sets of desired behaviour are used, the first set is the same as in earlier experiments, and the second is double that of the first. These values are shown in Table 5.4. The rewards are sized for each participant such that, when possible, they are below the right hand side of (5.12). The disclaimer "when possible" is included since it is possible that an ideal target sound is not available and the participant's attitude towards the chosen target sound is above the desired behaviour. In these instances the reward is set to the smallest denomination available.

5.3.2 Results

Eight participants completed the study focusing on cognitive dissonance when attitude is positive. The participants were each given a unique random alphanumeric identifier, namely: D6, E4 Q8, C9, P9, X4, S1, M4. In this section, we summarize the data collected from the participants. A more detailed analysis of each participant's data is presented in Appendix C.

Table 5.4: Desired behaviour of last six offers when attitude is positive

k	Set 1 $B_d[k+1]$ (s)	Set 2 $B_d[k+1]$ (s)
4	30	60
5	180	360
6	90	180
7	60	120
8	90	180
9	150	300

Model Validation

Each participant's data is plotted in the same format as that of Figure 5.6. The first subplot displays measured attitude, estimated attitude, and the response (either "Yes" or "No") to the offer posed at each time. The second subplot shows $\hat{K}_1$, the EKF's estimate of the cognitive dissonance proportionality constant. The third and final subplot shows the reward, $R[k]$, and the associated desired behaviour, $B_d[k+1]$.

Focusing on the first subplot of Figure 5.6, we observe that participant M4 declines many of the offers. Further, we see that following the declined offer at $k = 0$, the participant's attitude decreases. This decrease in attitude contradicts what we expect given the simplified model. The same observation can be made between $k = 7$ and $k = 9$.

The first subplot also shows that participant M4's attitude remains unchanged twice. This observation does not fit into the model. It is possible that the reward was not large enough to create a measurable attitude change.

The final observation that can be made from the first subplot of Figure 5.6 is that attitude may increase following a declined offer; this is seen three times from $k = 2$ to $k = 5$. This observation agrees directly with the model.

Using the same approach in Section 5.2.2, we objectively measure the model's accuracy by comparing the frequency of agreements between participant data and the model to the frequency of disagreements. This measure is motivated by the notion that a useful model is accurate more often than not. In addition, this allows us to capture and consider the various observations that are made when examining the first subplot.

In order to excite the effects of cognitive dissonance and validate the model, participants must decline offers. Despite our efforts to create offers that would be declined, five

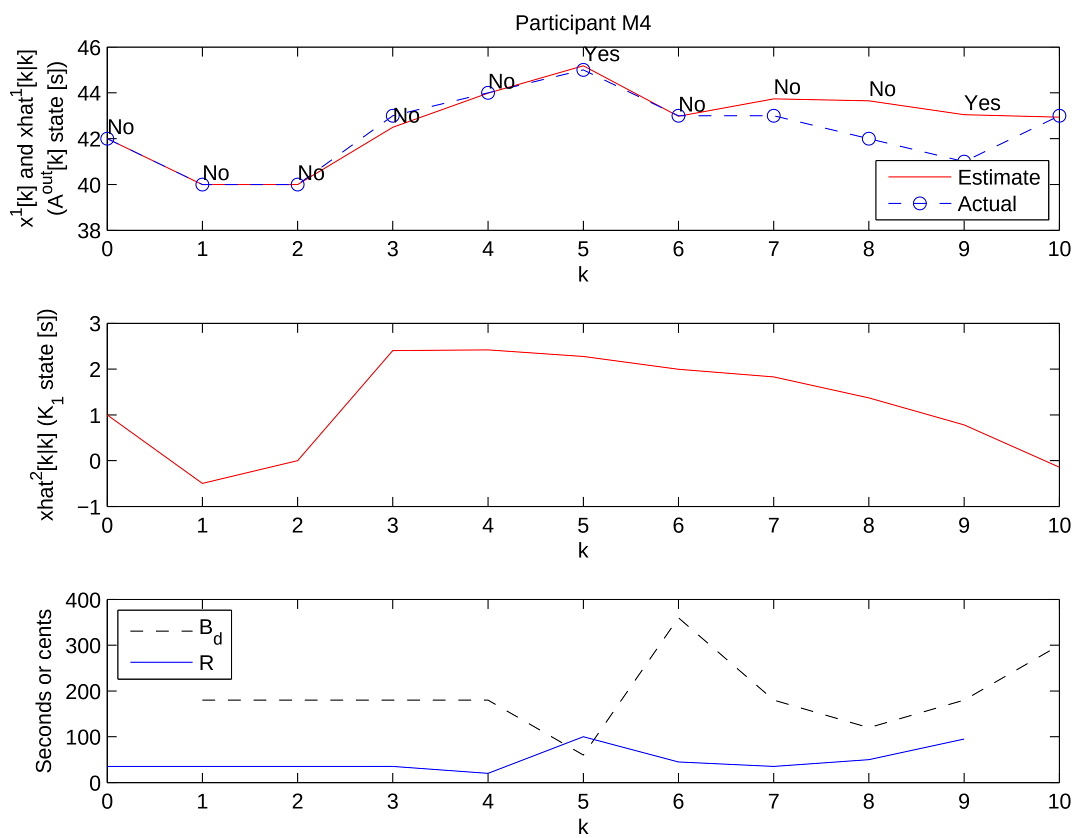

Figure 5.6: Data from Participant M4. Participant specific details: $\hat{\mu}_2 = 1.0000$ sec/mm, $\hat{\mu}_1 = 190.000$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 42.0000 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = -0.1455$ sec, target sound: Sunny Day, session start time: 12:00 PM.

Table 5.5: Summary of participants in positive attitude study.

Participant	$A_{out}[k] > 0 \, \forall \, k$?	# accepted	# declined
D6	No	9	1
E4	Yes	9	1
Q8	No	10	0
C9	No	10	0
P9	No	3	7
X4	Yes	1	9
S1	No	10	0
M4	No	2	8

Table 5.6: Summary of agreement between collected data and model when attitude is positive. The final column indicates if the frequency of explicit agreements is greater than the disagreements.

Participant	# Agree	# Disagree	# No Change	# Agree > # Disagree ?
P9	1	2	4	No
M4	3	3	2	No

participants accepted all or nearly all of the offers. Since accepting an offer does not excite cognitive dissonance, the data collected from these five participants is not very useful for validation. Additionally, two participants' attitudes do not remain positive for the entire study; consequently, their data cannot be used to comment on the validity of the model. Table 5.5 summarizes the response of all the participants. The participants whose attitude became negative in the study are shaded in red, while the participants that accepted all or many of the offers are shaded in yellow. Participant E4 falls into both categories.

As participants P9 and M4 are the only participants that declined the majority of their offers and had attitudes that remained positive, we use the data collected during their sessions with our measure. Table 5.6 lists: the frequency of agreements and disagreements between the model and the participant's data, frequency of no changes in attitude, and the result of the measure. The final column of Table 5.6 shows us that the model and collected data do not agree more often than they disagree for either participant. Thus the results of this experiment do not allow us to conclude that our model of cognitive dissonance when attitude is positive is valid.

The second subplot of Figure 5.6 displays $\hat{K}_1$, produced by the EKF, as it changes

Table 5.7: Summary of observations of $\hat{K}_1$ when attitude is positive.

Participant	$\hat{K}_1 > 0$?	Appears to be converging?
P9	No	No
M4	No	No

with each new measurement. A positive final value that appears to have settled or, at a minimum, appears to be converging suggests the model is valid. This is not the case in Figure 5.6; rather, $\hat{K}_1$ does not converge and is negative at $k = 10$.

Table 5.7 summarizes observations made from the second subplot for participants P9 and M4. Once again the data does not allow us to conclude that this model is valid.

Given the lack of support for the model and the limited set of usable data, no conclusions are made regarding the selection of EKF tuning parameters.

Chapter 6

Conclusions and Future Work

The goal of this thesis was to validate the core of the model presented in Ni, Kulic, and Davison's paper [24] while simultaneously identifying parameters of the model. More specifically, the aspects related to the theory of planned behaviour and cognitive dissonance theory are focused on. The validation task is split into smaller pieces by dealing with negative and positive attitudes separately. Extended Kalman Filtering (EKF) along with carefully worded questions was chosen as the most promising way to identify μ_1, μ_2, and K_1. An experiment was designed to collect the relevant data from human participants.

6.1 Model Validation

In Section 5.2.2, the model is shown to be more often correct rather than incorrect. This is encouraging and suggests that the negative attitude model has utility. The negative attitude model, and more generally, the full model, contain assumptions in order to manage complexity. It follows that deviations from the model (i.e., instances in which data collected disagrees with the model) may arise from violations of these underlying assumptions. Stronger conclusions may be reached by recruiting additional participants and applying the same analysis as that used in Section 5.2.2. Further validation of the negative attitude model is possible by applying the control scheme, such as the one outlined in [24], in an experimental environment.

No conclusions about the validity of the positive attitude model may be drawn at this time as much of the data collected was unusable. There are several explanations for the poor quality of data: First, the experiment design was not successful at making participants decline all the offers. Second, in a few instances, offers did not satisfy all the logical requirements and resulted in unusable data. Third, contrary to the model, the experiment does not provide an opportunity to participants with positive attitudes to perform the behaviour (i.e., listen to the sound) after declining an offer. Given these issues, revisions to the experiment are required in order to validate cognitive dissonance in the positive attitude model. The experiment must be redesigned with the following in mind:

- The strategy used for sizing offers must be modified; increasing the desired behaviour while maintaining a significant reward is expected to make participants decline offers while maintaining measurable amounts of attitude change.
- Offers should satisfy all logical constraints outlined in Section 5.3.1. It is recommended that offers be dynamically prepared rather than using predetermined values. This will eliminate erroneous offers, thus ensuring the data collected is usable.
- Participants should be provided with a opportunity to perform the behaviour when they decline an offer. This opportunity must be fluidly incorporated into the flow of the experiment, and also be measured to ensure consistency with the model.

Following the validation of the positive attitude model, the transitions from negative to positive attitude and positive to negative attitude should be verified.

6.2 Parameter Estimation

Standard least squares was shown to not be a suitable approach for simultaneously estimating K_1 and μ_1. The reason for least squares' poor performance should be further investigated. Additional analysis should be performed to ensure the input sufficiently excites the system and that $\Phi^T\Phi$ is well conditioned. Additionally, the performance of least squares should be re-evaluated in simulation when K_1 is the only unknown parameter being estimated.

While the EKF was deemed a suitable method for identifying K_1, further optimizations to the EKF may still be available. The system may be redefined such that there is one less state (i.e., redefine the input $u[k]$). Further, the uniform noise model assumed in simulation may be modified to better reflect the noise present in the experimental data. Additionally, the behaviour of the EKF should be further studied when attitude is positive and all offers are declined. The proportion of estimates with 10% or 20% error decreases as the number of measurement samples used increases (refer to Table 4.11). This is contrary to the expected behaviour of the EKF.

The approaches used to estimate μ_1 and μ_2 are less than ideal. Only a single measurement is used in each approach to identify the respective parameter. Future work should incorporate a more robust method to estimate each parameter. A method using multiple measurements is preferred. The data collected when participants listen to and rate alternate sounds (refer to Figure 5.1) may be used to improve $\hat{\mu}_2$.

It is recommended that parameters identified from each participant be stored in a library or database. From this library we may gain relevant statistical knowledge to apply alternate identification approaches (such as MAP). Alternatively, the EKF may be more accurately initialized using the mean of $\hat{K}_1$'s in the library. Further, knowledge of the parameters may enable default control techniques to be developed (i.e., a controller designed using generalized knowledge of the family of plants).

6.3 Experiment Design

All future experiments should incorporate statistical analysis tools (e.g., hypothesis testing) to analyze experimental data. Additionally, future experiments should include a control group whenever possible. A suitable control group for the experiment in Chapter 5 would be one in which participants go through the experiment, but are not offered any rewards (the wording at each offer should be modified as not to mention a reward). The degree to which attitudes change in the control group may be compared to attitude change in participants not part of the control group.

The distractors used in the experiment were effective at masking the true purpose of the study. During the debriefing period participants were asked what they believed the study was about. None of the participants accurately identified the purpose of the study. One participant believed the sounds were an important part of the experiment due to the frequent offers and ratings.

After carrying out the experiment, it was noted that the distractors may have introduced additional cognitions. Cognitions such as boredom (due to a variety of mundane activities), frustration and anger (due to repeated activities), or emotional responses (arising from reading passages) may have influenced participants' attitudes towards the target sounds. These additional effects are not incorporated in the model. Future experiments should consider incorporating distractors that minimize unintended effects.

The sound clips used may not have been challenging enough to listen to in the negative attitude experiment. Said differently, the sounds clips may not have been very unpleasant at all. Additionally, the sound clips may not have been been engaging enough. Participants' often looked bored, uninterested, or looked around the room as they listened to the sounds following an offer. Future experiments should consider using more extreme sounds or perhaps an even more extreme activity. A more engaging alternative is to use speeches or lectures; subsequent offers would continue the speech or lecture where the participant last left off.

APPENDICES

Appendix A

The Importance of Observability in Extended Kalman Filtering

To facilitate simultanenously identifying K_1, μ_1, and μ_2, we develop the relevant details required to apply the EKF. The poor estimation performance when using this setup is presented. We note that the linearized system is frequently unobservable and address this issue by defining alternate outputs. Ultimately, our attempts fail to produce a system that is observable.

Recall definitions (4.27), (4.28), and (4.29):

$$u[k] = R[k-1], \tag{4.27}$$

$$y[k] = \frac{A_{out}[k]}{\mu_2}, \tag{4.28}$$

$$\mathbf{x}[k] = \begin{bmatrix} x_1[k] \\ x_2[k] \\ x_3[k] \\ x_4[k] \\ x_5[k] \end{bmatrix} = \begin{bmatrix} A_{out}[k] \\ K_1[k] \\ \mu_1[k] \\ \mu_2[k] \\ u[k] \end{bmatrix}. \tag{4.29}$$

When attitudes are negative, the state transition function, f, can be defined as follows:

$$f\left(\mathbf{x}[k], u[k]\right) = \begin{bmatrix} x_1[k] + a[k] \\ x_2[k] \\ x_3[k] \\ x_4[k] \\ u[k] \end{bmatrix} \tag{A.1}$$

where $a[k]$ is determined by the response to an offer and is defined:

$$a[k] = \begin{cases} +x_2[k]\dfrac{|x_1[k]|}{|x_1[k]| + |x_3[k]x_5[k]|} & \text{if offer accepted at time } k \\ -x_2[k]\dfrac{|x_3[k]x_5[k]|}{|x_1[k]| + |x_3[k]x_5[k]|} & \text{if offer declined at time } k. \end{cases} \tag{A.2}$$

The response to each offer must be known to apply the EKF. This is trivially recorded both in simulation and in experimental environments. The parameters K_1, μ_1, and μ_2 are constants; the dynamics describing these parameters in (A.1) state that the next value is simply the previous value. The output function, h, is defined as follows:

$$h\left(\mathbf{x}[k]\right) = \frac{x_1[k]}{x_4[k]}. \tag{A.3}$$

The Jacobian of the state transition function depends on the response to the offer. As such we use the recorded response to each offer to select the appropriate linearized dynamics. We use superscripts A and D to denote the Jacobian of the accepted and declined offers respectively. For the accepted offers, we have

$$\mathbf{F}_{k-1}^{A} = \left.\begin{bmatrix} \frac{\partial f_1^A}{\partial x_1[k]} & \frac{\partial f_1^A}{\partial x_2[k]} & \frac{\partial f_1^A}{\partial x_3[k]} & \frac{\partial f_1^A}{\partial x_4[k]} & \frac{\partial f_1^A}{\partial x_5[k]} \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}\right|_{\substack{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1} \\ u[k]=u[k-1]}} \tag{A.4}$$

where:

$$\frac{\partial f_1^A}{\partial x_1[k]} = 1 + \frac{x_2[k]}{|x_1[k]| + |x_3[k]x_5[k]|} \cdot \frac{x_1[k]}{|x_1[k]|} - \frac{x_1[k]x_2[k]}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)^2} \tag{A.5}$$

$$\frac{\partial f_1^A}{\partial x_2[k]} = \frac{|x_1[k]|}{|x_1[k]| + |x_3[k]x_5[k]|} \tag{A.6}$$

$$\frac{\partial f_1^A}{\partial x_3[k]} = \frac{-\,|x_1[k]|x_2[k]}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)^2} \cdot \frac{x_3[k]x_5[k]^2}{|x_3[k]x_5[k]|} \tag{A.7}$$

$$\frac{\partial f_1^A}{\partial x_4[k]} = 0 \tag{A.8}$$

$$\frac{\partial f_1^A}{\partial x_5[k]} = \frac{-\,|x_1[k]|x_2[k]}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)^2} \cdot \frac{x_3[k]^2x_5[k]}{|x_3[k]x_5[k]|}, \tag{A.9}$$

and for the declined offers, we have

$$\mathbf{F}_{k-1}^{D} = \begin{bmatrix} \frac{\partial f_1^D}{\partial x_1[k]} & \frac{\partial f_1^D}{\partial x_2[k]} & \frac{\partial f_1^D}{\partial x_3[k]} & \frac{\partial f_1^D}{\partial x_4[k]} & \frac{\partial f_1^D}{\partial x_5[k]} \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \Bigg|_{\substack{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1} \\ u[k]=u[k-1]}} \tag{A.10}$$

where:

$$\frac{\partial f_1^D}{\partial x_1[k]} = 1 + \frac{x_2[k]|x_3[k]x_5[k]|}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)^2} \cdot \frac{x_1[k]}{|x_1[k]|} \tag{A.11}$$

$$\frac{\partial f_1^D}{\partial x_2[k]} = \frac{-\,|x_3[k]x_5[k]|}{|x_1[k]| + |x_3[k]x_5[k]|} \tag{A.12}$$

$$\frac{\partial f_1^D}{\partial x_3[k]} = \frac{-\,x_2[k]x_3[k]x_5[k]^2}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)\left(|x_3[k]x_5[k]|\right)} + \frac{x_2[k]x_3[k]x_5[k]^2}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)^2} \tag{A.13}$$

$$\frac{\partial f_1^D}{\partial x_4[k]} = 0 \tag{A.14}$$

$$\frac{\partial f_1^D}{\partial x_5[k]} = \frac{-\,x_2[k]x_3[k]^2x_5[k]}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)\left(|x_3[k]x_5[k]|\right)} + \frac{x_2[k]x_3[k]^2x_5[k]}{\left(|x_1[k]| + |x_3[k]x_5[k]|\right)^2}. \tag{A.15}$$

The Jacobian of the output function is:

$$\mathbf{H}_k = \begin{bmatrix} \frac{1}{x_4[k]} & 0 & 0 & \frac{-x_1[k]}{x_4[k]^2} & 0 \end{bmatrix} \Big|_{\mathbf{x}[k]=\hat{\mathbf{x}}_{k-1|k-1}}. \tag{A.16}$$

The results of a single simulation are shown in Figure A.1. In this simulation attitude remains negative and all all offers are accepted. The parameters used are: $N = \{5, 10, 20\}$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = 0$, $B_d[k] = 10$, and $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix}-51 & 2 & 7 & 1 & 10\end{bmatrix}$. Note that the initial state estimate are the exact state values of the system, the most ideal value possible. This is not say that all of the EKF requirements are set ideally; the error covariance estimate, process noise covariance, and measurement noise covariance values used in all the EKF simulations are:

$$\mathbf{P}_{0|0} = \begin{bmatrix} 100 & 0 & 0 & 0 & 0 \\ 0 & 100 & 0 & 0 & 0 \\ 0 & 0 & 100 & 0 & 0 \\ 0 & 0 & 0 & 100 & 0 \\ 0 & 0 & 0 & 0 & 100 \end{bmatrix},$$

$$\mathbf{Q} = \begin{bmatrix} 10^{-5} & 0 & 0 & 0 & 0 \\ 0 & 5\times10^{-5} & 0 & 0 & 0 \\ 0 & 0 & 10^{-5} & 0 & 0 \\ 0 & 0 & 0 & 10^{-5} & 0 \\ 0 & 0 & 0 & 0 & 10^{-5} \end{bmatrix},$$

$$\mathbf{R} = 10^{-2}.$$

Figure A.1 shows that none of the estimates, including the attitude estimate, approach their true values after 20 measurements. There is a severe lack of convergence to the correct values despite $\hat{y}_k$ being driven to $y[k]$. Simulating with alternate parameter choices returns similar results.

Figure A.2 shows the error distribution of each of the parameters over 1,000,000 trials. Parameters were selected from the ranges listed in Table 4.4; additionally, the initial estimates of $A_{out}[k]$, μ_1, and μ_2 were chosen with up to 20% uncertainty. The error distributions of $\hat{K}_1$, $\hat{\mu}_1$, and $\hat{\mu}_2$ are shown in the first, second, and third subplot of Figure A.2 respectively. The distributions are abysmal, estimates are unlikely to be within the 20% of the actual parameter value. Simultaneous parameter estimation is clearly not a good idea.

To investigate why parameters fail to converge when simultaneously estimating multiple parameters we look at the observability of the system. The observability matrix, $\mathcal{O}$, of this system is:

$$\mathcal{O} = \begin{bmatrix} \mathbf{H}_k \\ \mathbf{H}_k\mathbf{F}_{k-1} \\ \mathbf{H}_k\mathbf{F}_{k-1}^2 \\ \mathbf{H}_k\mathbf{F}_{k-1}^3 \\ \mathbf{H}_k\mathbf{F}_{k-1}^4 \\ \mathbf{H}_k\mathbf{F}_{k-1}^5 \end{bmatrix}. \tag{A.17}$$

The system is observable when $\mathcal{O}$ is full rank (i.e., $rank(\mathcal{O}) = 5$ in this setup). When the observability in Figure A.1 is checked at each iteration, we find that $rank(\mathcal{O}) = 3$. Simulating over other choice of parameters returns the same result.

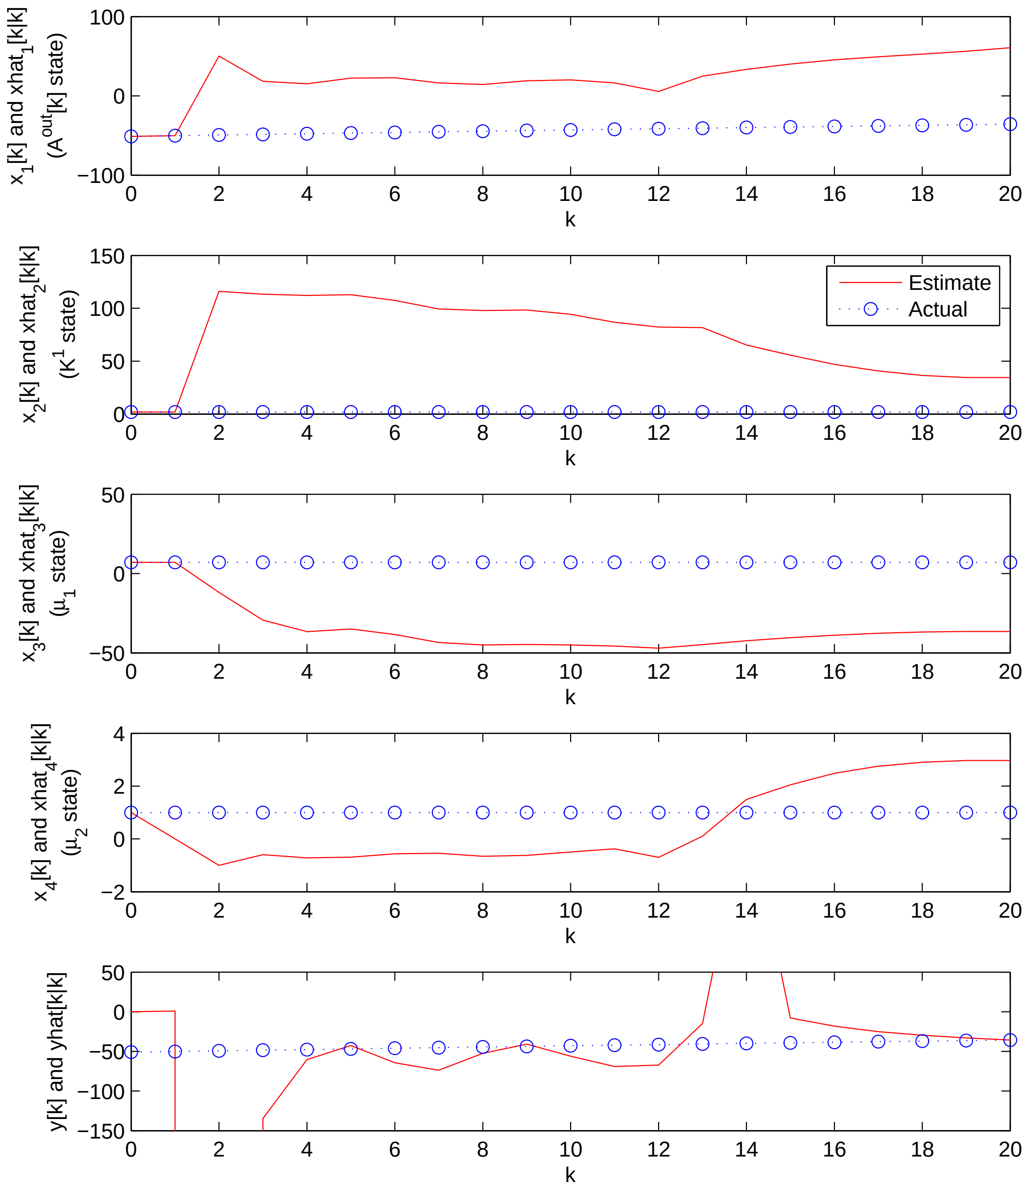

Figure A.1: EKF estimates plotted alongside actual values generated in simulation when all offers are declined. Parameter values: $N = 20$, $K_1 = 2$, $\mu_1 = 7$, $\mu_2 = 1$, $A_{out}[0] = -51$, $R[k] = 10$, $|\text{noise}| = 0$, $B_d[k] = 10$, and $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} -51 & 2 & 7 & 1 & 10 \end{bmatrix}^T$.

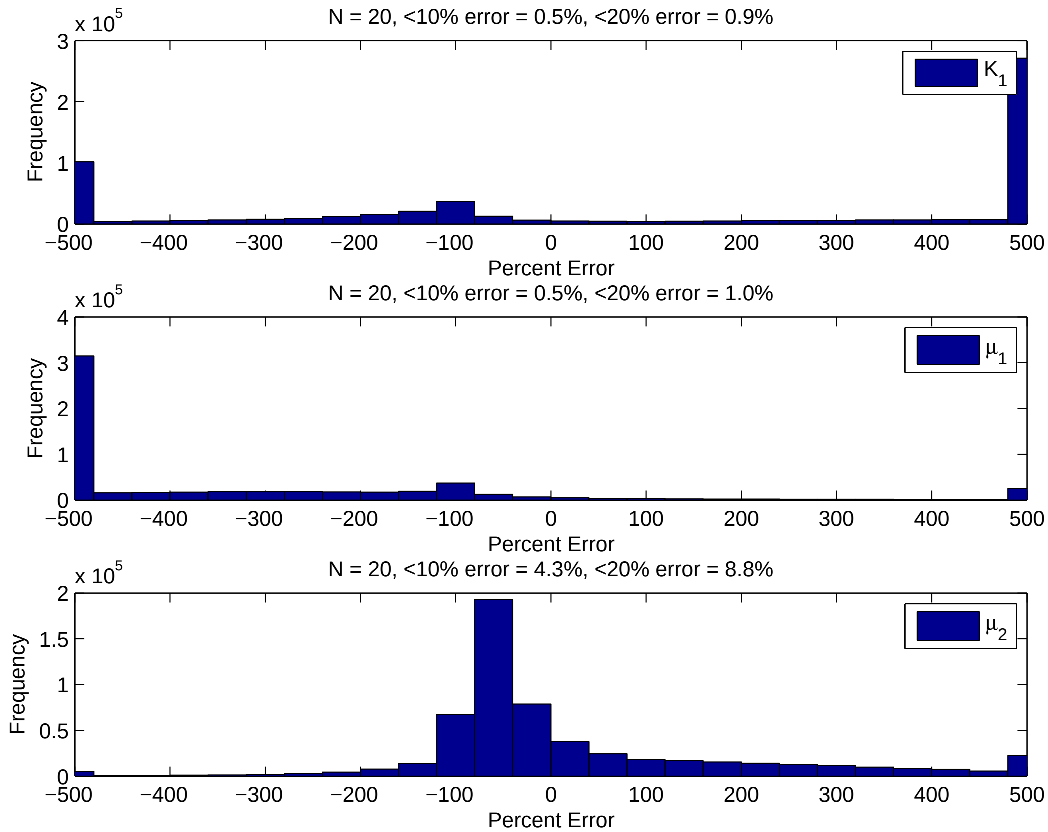

Figure A.2: Distribution of errors of EKF estimates when attitude is negative, and all parameters are simultaneously estimated. Parameters selected from Table 4.4 and $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} \zeta A_{out}[k] & 40 & \gamma\mu_1 & \delta\mu_2 & R[k] \end{bmatrix}^T$ where ζ, γ, and δ are randomly selected from $[0.8, 1.2]$.

Observability may be improved by either changing the dynamics of the system or by modifying the outputs (i.e., measurements). The dynamics of the system can only be modified by reducing the number of parameters being estimated. The output may be modified to include additional measurements of the system. Practically, no additional measurements exist; the output of the attitude measurement tool, $\frac{x_1[k]}{x_4[k]}$, is all that is available. A technique to include additional outputs is to define new outputs based on existing states and outputs. As we do not require that the EKF be run online, we may look at non-causal outputs. We redefine the output to include an alternate output based on a future attitude and resize the measurement noise covariance matrix:

$$h\left(\mathbf{x}[k]\right) = \begin{bmatrix} \dfrac{x_1[k]}{x_4[k]} \\ \dfrac{x_1[k+2]}{x_4[k+2]} \end{bmatrix}, \tag{A.18}$$

$$\mathbf{R} = \begin{bmatrix} 10^{-2} & 0 \\ 0 & 10^{-2} \end{bmatrix}. \tag{A.19}$$

The redefined output in (A.18) may only increase $rank(\mathcal{O})$ by 1 which is not enough. In order to further investigate, we assume μ_2 is known and reduce the number of parameters being estimated (and, consequently, the number of states and dimensions of the error covariance matrix). In doing so, we find that the system is now observable; however, the EKF's estimates still fail to converge to the correct values. The same is true with the alternate output definition in (A.20):

$$h\left(\mathbf{x}\right) = \begin{bmatrix} \dfrac{x_1[k]}{x_4[k]} \\ \left(\dfrac{x_1[k+1]}{x_4[k+1]}\right)^3 \end{bmatrix}. \tag{A.20}$$

Investigating further, the full rank observability matrices are close to being singular. As such, it is possible the rank of the observability matrix is not truly indicative of the observability of the system.

Appendix B

Sound Clips

The following sounds were used in the experiment. Some sounds were renamed for use in the experiment; these names are shown parentheses.

Sound: Applause
Author: Mike Koenig
License: Attribution 3.0
URL: http://soundbible.com/989-10-Second-Applause.html

Sound: Laughter (Female Laughter)
Author: Ezwa
License: Public Domain
URL: http://www.pdsounds.org/sounds/laughter_0

Title: Sunny Day
Author: stephan
License: Public Domain
URL: http://soundbible.com/1661-Sunny-Day.html

Sound: Drill
Author: Mike Koenig
License: Attribution 3.0
URL: http://soundbible.com/1074-Drill.html

Sound: Alien Buzzer

Author: Kevan
License: Public Domain
URL: http://soundbible.com/1811-Annoying-Alien-Buzzer.html

Sound: Long Wanting Cry (Longing Baby)
Author: Natalie
License: Public Domain
URL: http://www.pdsounds.org/sounds/baby_long_wanting_cry

Sound: Alarm Clock (Digital Alarm)
Author: UncleKornicob
License: Public Domain
URL: http://soundbible.com/1787-Annoying-Alarm-Clock.html

Title: Crisp Ocean Waves (Ocean Wave)
Author: Mike Koenig
License: Attribution 3.0
URL: http://soundbible.com/1936-Crisp-Ocean-Waves.html

Appendix C

Observations

In this appendix the experimental data is presented. Before the data is presented, an overview of data collection, analysis, and experiment details is provided.

C.1 Interpreting the Data

Each participant's data is presented in a figure containing three subplots. The first subplot shows the measured attitude, estimated attitude, and whether the participant accepted the offer posed at interval k. The measured attitude, in units of seconds, is shown with a blue dashed line and blue circles at each marker. The EKF's estimated attitude, also in units of seconds, is shown using a red solid line.

The second subplot displays the EKF's estimate of K_1 using a solid red line.

The third and final subplot displays details of each offer. The reward, in units of cents, is plotted in a blue solid line. The behaviour desired, in units of seconds, is plotted using a black dashed line. The offer at time k consists of the reward at time k and the desired behaviour at time $k+1$. For example, the third offer (i.e., the offer at $k = 2$) would consist of $R[2]$ and $B_d[3]$.

In the experiment, attitude is measured by asking participants to place a tick on a horizontal line segment. The line contains three markings, "unpleasant" and "pleasant" at the left and right extremes, and "neutral" at the center. The tick's position relative to the "neutral" marking is measured, in millimetres, and scaled to seconds using an estimate of the conversion constant, $\hat{\mu}_2$.

The conversion constant, μ_2, is estimated prior to offers being made in the experiment. The estimate, $\hat{\mu}_2$, is calculated using the participant's rating, in millimetres, of a randomly selected positive sound (i.e., a sound for which the participant placed a tick between "neutral" and "pleasant") and the length of time the participant listened to the sound. The researcher surreptitiously times the listening time.

The parameter μ_1 is also estimated prior to offers being made in the experiment. The estimate, $\hat{\mu}_1$, is calculated using $\hat{\mu}_2$, the participant's rating, in millimetres, of a negative sound (i.e., a sound for which the participant placed a tick between "unpleasant" and "neutral"), and their response to the hypothetical question: "How much money would it take for you to listen to <negative sound> sound for 60 seconds?"

The desired behaviour, $B_d[k]$, at each time was fixed for all participants in the negative attitude experiment. In the positive attitude experiment, the desired behaviour was selected from one of two sets. The first set was the same as that used in the negative attitude experiment. The second set doubled the values in the first set.

The first four rewards are specific to each participant. In the negative attitude experiment, these four rewards were sized such that participants would accept the offer but also such that dissonance pressure would be near maximum. While in the positive attitude experiment, the first four rewards were sized such participants would decline the offer. The reward was sized using the participant specific parameters $\hat{\mu}_1$ and $\hat{\mu}_2$, their rating of the target sound, and the desired behaviour value. In some sessions focusing on negative attitudes, the researcher ramped up the reward when the participant failed to accept the offer.

The final six rewards were fixed for all participants in the negative attitude study. In the positive attitude study, the final six offers were adjusted for each participant to encourage them to decline offers.

The tuning parameters used to obtain EKF estimates in the first and second subplot of each figure are the same for all participants. These parameters are:

$$\begin{aligned}
\mathbf{P}_{[0|0]} &= \begin{bmatrix} 100 & 0 & 0 \\ 0 & 100 & 0 \\ 0 & 0 & 100 \end{bmatrix}, \\
\mathbf{Q} &= \begin{bmatrix} 10^{-5} & 0 & 0 \\ 0 & 5 \times 10^{-5} & 0 \\ 0 & 0 & 10^{-5} \end{bmatrix}, \\
\mathbf{R} &= 10^{-2}.
\end{aligned}$$

Specific details for each participant are presented in the caption below each figure. These details include the values of $\hat{\mu}_2$, $\hat{\mu}_1$, the initial state estimate $\hat{x}_{[0|0]}$, the final value estimate of K_1 (i.e., $\hat{K}_1$ at $k = 10$), the unpleasant target sound, and the start time of the session.

C.2 Negative Attitude Participant Data

The data from each participant's session is presented in the following pages. The data is presented in chronological order.

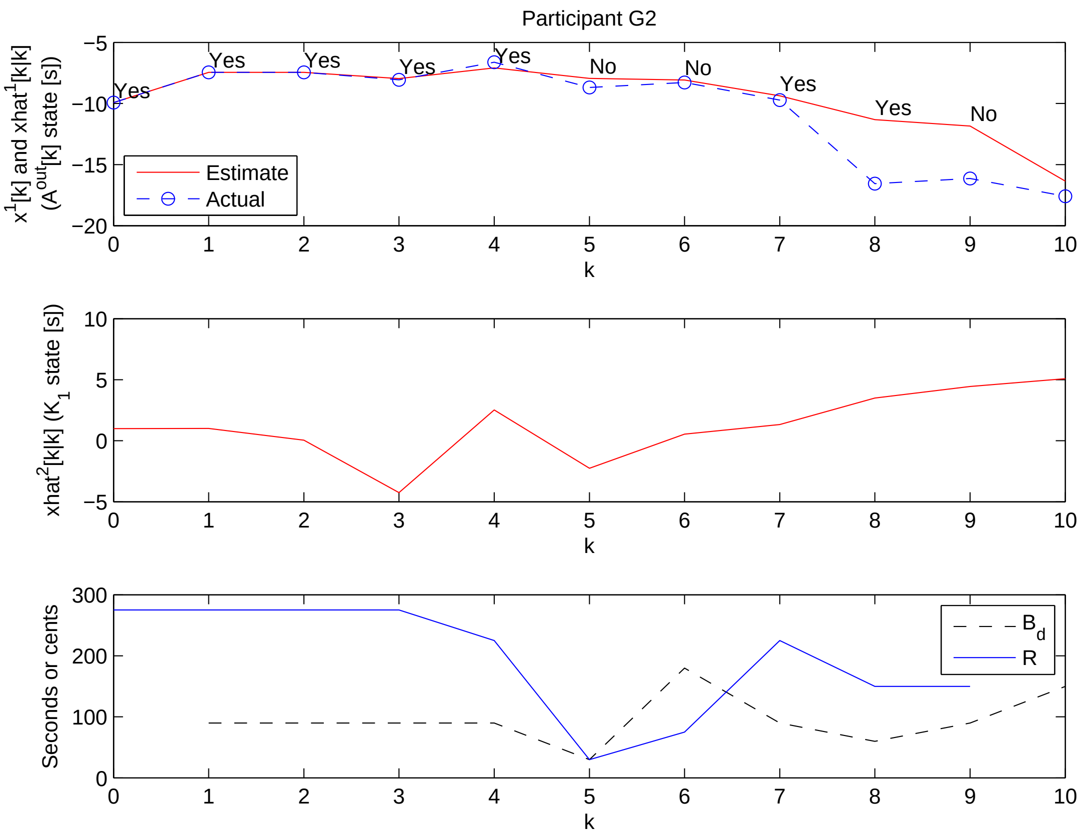

Figure C.1: Data from Participant G2. Participant specific details: $\hat{\mu}_2 = 0.4138$ sec/mm, $\hat{\mu}_1 = 34.966$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -9.9132 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 5.0824$ sec, target sound: Digital Alarm Clock, session start time: 4:05 PM.

Reviewing participant G2's attitude and behaviour in the first subplot of Figure C.1, we see that the first four offers are accepted and that attitude increases slightly during this time. This qualitatively follows the model's prediction. At $k = 8$ we see the attitude estimate deviates from the measured attitude. After the eighth offer is accepted there is a decrease in attitude; this change is opposite of what the model predicts. As the EKF considers both the model and the measurement, its attitude estimate is in between that of the prediction and measurement. This can be viewed as the EKF filtering noisy measurements.

A single small change in attitude can be attributed, or even dismissed, due to the effects of measurement noise. However, the changes over multiple samples are noteworthy; the increase in attitude between $k = 0$ and $k = 5$ suggest the model is indeed able to predict how attitude changes in response to offering rewards for specific behaviours.

The second subplot in Figure C.1 shows that $\hat{K}_1$ is likely converging to a positive value. This is important as negative values would indicate that the model is not accurate. Further, wild swings in $\hat{K}_1$ would mean that the EKF is not performing well. Overall the EKF is performing well.

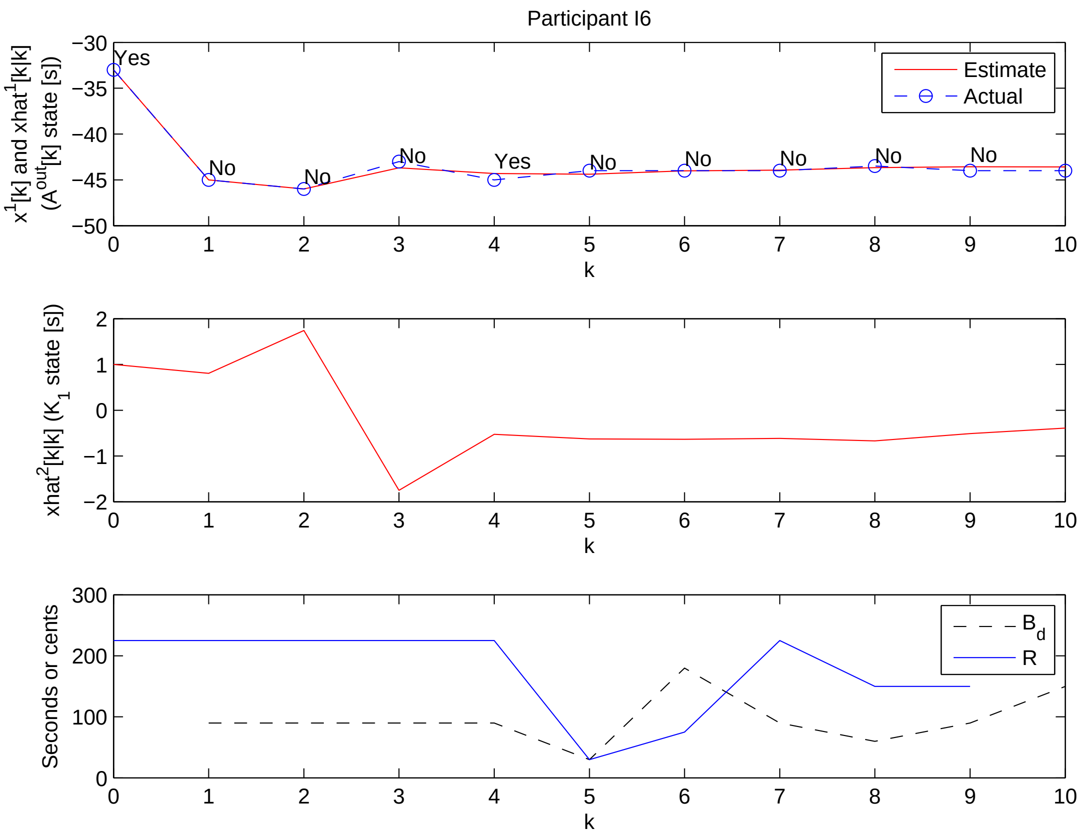

Figure C.2: Data from Participant I6. Participant specific details: $\hat{\mu}_2 = 1.000$* sec/mm, $\hat{\mu}_1 = 27.000$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -33 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = -0.3902$ sec, target sound: Drill, session start time: 5:30 PM. *See discussion below.

The first subplot of Figure C.2 shows a decrease in attitude after the first offer is accepted. This behaviour and change in attitude is contrary to the model. The large decrease in attitude may be caused by learning effects; that is, the participant may have realized one or more of the following: (a) after listening to the sound again, they disliked the sound more than they originally indicated; (b) after listening to all the sounds, they had rated the target sound too high; or (c) how long 90 seconds truly is. These learning effects may have influenced the participant to decline the majority of later offers. (To prevent (a), and to a lesser degree, (c), later participants were asked to listen to each sound for 30 seconds during the initial rating. No methods to prevent (b) were implemented in the study.) The learning effects could be viewed as noise of larger magnitude acting on the initial attitude measurement. There is minimal attitude change after $k = 2$. The participant's attitude did not change enough to qualitatively support or refute the model. Perhaps more time is needed for attitudes to change, or other factors are at play.

The second subplot in Figure C.2 shows $\hat{K}_1$ taking a negative value. This suggests the model is not correct. Recall that K_1 is the proportional constant that relates dissonance pressure to change in attitude. The measurements show little to no change in attitude despite there being some dissonance pressure; thus, the EKF should produce $\hat{K}_1$ values close to 0. Looking closely, the estimate of $\hat{K}_1$ appears to be approaching 0. This suggests that the EKF is working as intended.

The "*" in the caption of Figure C.2 indicates this is an assumed value for $\hat{\mu}_2$, not the true value determined from the study. Participant I6 rated all sounds as negative (i.e., between "unpleasant" and "neutral"). As such, no sounds to which the participant had a positive attitude towards were available to determine $\hat{\mu}_2$. In order to make use of the data collected in the study, it was assumed that $\hat{\mu}_2 = 1.000$. This assumption changes the scale of attitude measurements, the size of the first four offers, and impacts the estimates of $\hat{\mu}_1$ and $\hat{K}_1$.

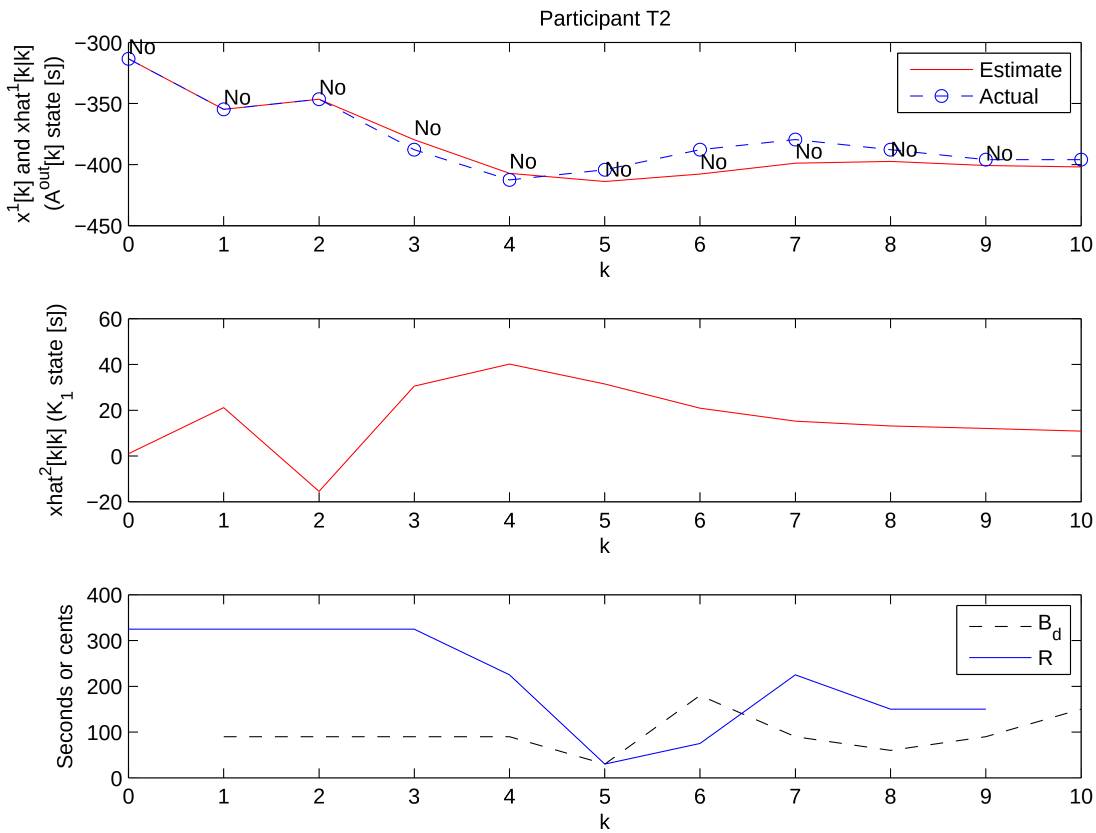

Figure C.3: Data from Participant T2. Participant specific details: $\hat{\mu}_2 = 8.250$ sec/mm, $\hat{\mu}_1 = 124.500$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -313.5 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 10.9025$ sec, target sound: Digital Alarm Clock, session start time: 4:30 PM.

Reviewing participant T2's behaviour and attitude in the first subplot of Figure C.3, we see that the first four offers are declined and that attitude decreases during the initial four offers. The first four offers may not have been large enough to entice participant T2 to accept the offers as we had intended. This suggests that the estimates of $\hat{\mu}_2$ and $\hat{\mu}_1$ are not ideal or other factors were at play. Nonetheless, the participant's attitude still follows the trend in our model; that is, when offers are declined, attitudes will decrease. The attitude at $k = 2$ increased marginally from $k = 1$; this may be attributed to the noisy attitude measurement rating system.

Beyond the first four offers, all there is a small increase in attitude despite declining all offers. These small increases are possibly within the range of measurement noise. Looking at the change in attitude from beginning to end, there is a definite decrease in attitude. This trend is consistent with our model.

The $\hat{K}_1$ value shown in the second subplot of Figure C.3 is arguably converging to a positive value. Convergence to a positive value supports the model's validity and shows that the EKF is working well.

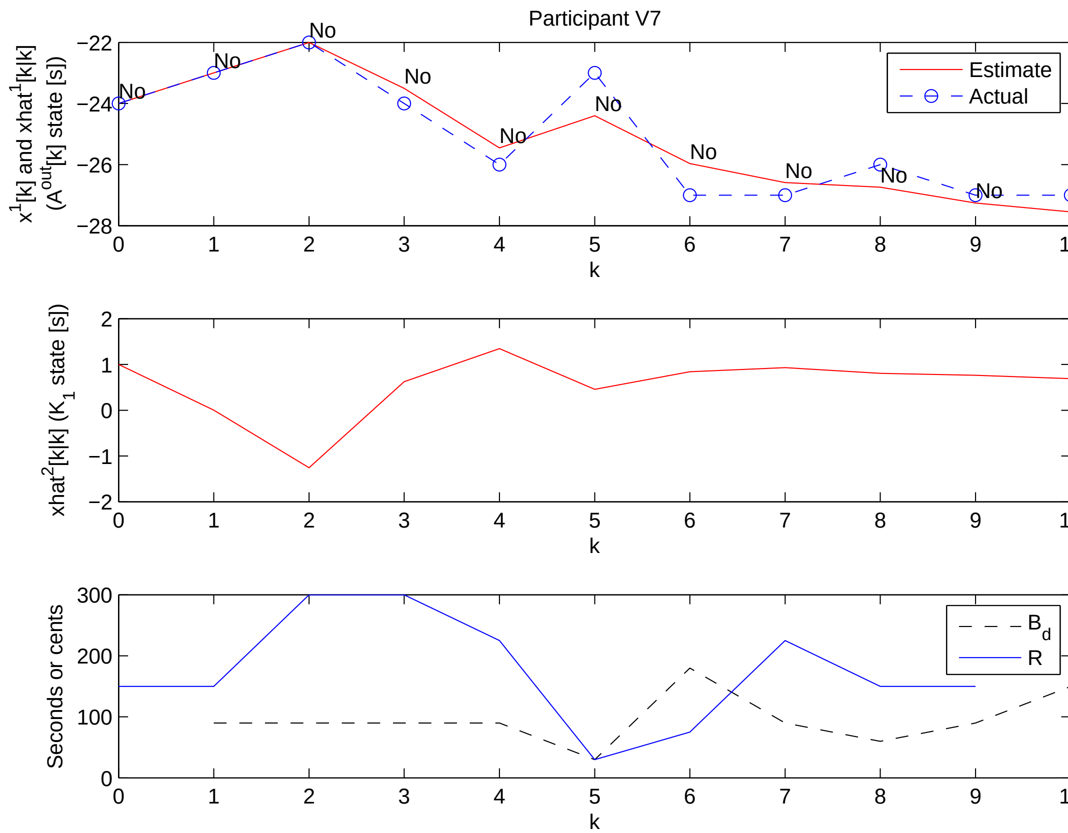

Figure C.4: Data from Participant V7. Participant specific details: $\hat{\mu}_2 = 1.000$* sec/mm, $\hat{\mu}_1 = 60.000$** sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -24 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 0.6911$ sec, target sound: Female Laughter, session start time: 5:30 PM. *See discussion below. **See discussion below.

Reviewing participant V7's attitude in the first subplot of Figure C.4 reveals that all offers were declined. The participant's measured attitude appears to be noisy; however, looking at the overall change in attitude from $k = 0$ to $k = 10$ we see a decrease in attitude. The data provided by participant V7 supports the model as it shows that attitude decreases when offers are declined.

Reviewing the second subplot of Figure C.4, we can see that $\hat{K}_1$ is converging to a positive value. This suggests that the model is correct. Furthermore it suggests the EKF is effective at estimating the value of K_1 given a small sequence of data.

The "*" in the caption of Figure C.4 indicates this is an assumed value for $\hat{\mu}_2$, not the true value determined from the study. Participant V7 indicated they do not wish to listen to the "Applause" sound, a sound they had previously rated positively. It was necessary to assume a value of $\hat{\mu}_2$ in order to continue the session; thus, it was assumed $\hat{\mu}_2 = 1.000$. This assumption affects the scale of attitude measurements, the size of the first four offers, and impacts the estimates of $\hat{\mu}_1$ and $\hat{K}_1$.

Similarly, the "**" in the caption of C.4 indicates this is an assumed value for $\hat{\mu}_1$ and not the true value determined from the data collected in the study. Participant V7 responded it would take $0 for them to listen to "Female Laughter" for 60 seconds. After this response, participant V7 explained they would not accept money to listen to an unpleasant sound; they believed it would lead to increasingly extreme offers (i.e., more unpleasant activities for larger rewards). The value of $\hat{\mu}_1$ is proportional to the inverse of the reward, and thus a divide by zero error is present. To continue the study and analyze the participant's responses, it was assumed that participant V7 had responded $1 rather than $0.

Both the preceding assumptions greatly affect the size of the first four offers. Insufficient rewards would result in the participant declining offers rather than accepting them (as was the goal). This is not believed to be the true reason for the declined offers in this data set; rather, the participant's comments suggest other factors were at play. These factors are far outside the scope of the model and the study.

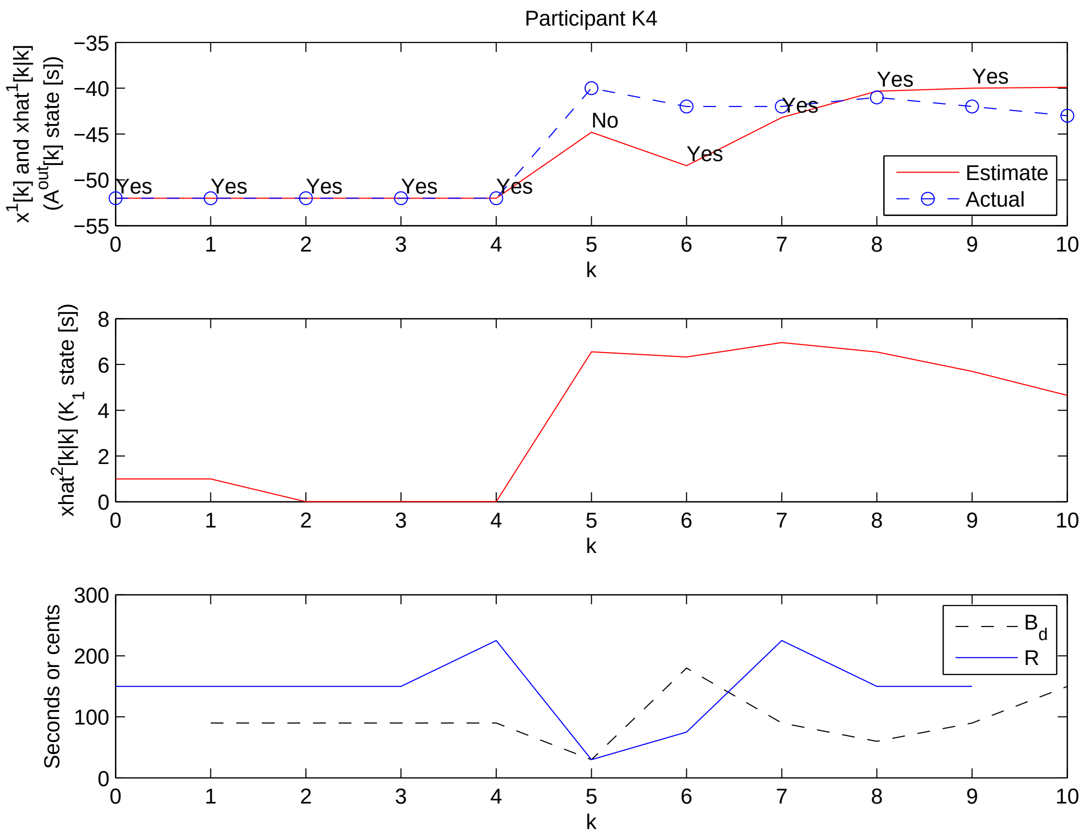

Figure C.5: Data from Participant K4. Participant specific details: $\hat{\mu}_2 = 1.000$* sec/mm, $\hat{\mu}_1 = 60.000$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -52 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 4.6455$ sec, target sound: Applause, session start time: 11:15 AM. *See discussion below.

The first subplot of Figure C.5 shows that participant K4's attitude remains unchanged until after the fifth offer. Participant K4's attitude was saturated on the low end of the scale; thus, we are unable to measure changes in attitude during the first few offers. We can only speculate that the participant's attitude was much more negative, steadily increased, and finally appeared on the scale after accepting the fifth offer. The change in attitude following offers five through eight all agree with the model, while offers nine and ten do not. Subsequent measurements may once again be within the range of noise, but the long term trend is a definite increase in attitude. This agrees with the observation that the majority of offers were accepted.

Throughout the study distractors are employed to prevent participants from figuring out the true purpose of the study. One distractor asks participants to listen to and rate an alternate sound. Between offers 4 and 5, participant K4 was asked to listen to the "Air Horn" sound and then rate it. After listening to this sound, the participant's attitude towards the target sound ("Applause") increased. It is possible that their attitude towards the target sound was influenced by the less pleasant "Air Horn" sound.

Participant K4 rated all sounds as unpleasant and placed ticks at the far left edge of the scale; thus, it was not possible to avoid the saturation issues by picking an alternate unpleasant sound.

The second subplot of Figure C.5 shows that $\hat{K}_1$ is steadily approaching a positive value. The estimated value does not appear to have settled as well as it does for other participants. This is caused by the repeated saturated attitude measurements which fail to convey the change in attitude speculated in the first five attitude measurements.

The "*" in the caption of Figure C.5 indicates this is an assumed value for $\hat{\mu}_2$, not the true value determined from the study. As mentioned earlier, Participant K4 did not rate any sounds as pleasant (i.e., all sounds were rated between "unpleasant" and "neutral"). As such, it was not possible to time how long the participant listened to a pleasant sound. Further, the participant indicated they would not listen to an unpleasant sound. In order to make use of the data they provided in the study, it was assumed that $\hat{\mu}_2 = 1.000$. This assumption changes the scale of attitude measurements, the size of the first four offers, and impacts the estimates of $\hat{\mu}_1$ and $\hat{K}_1$.

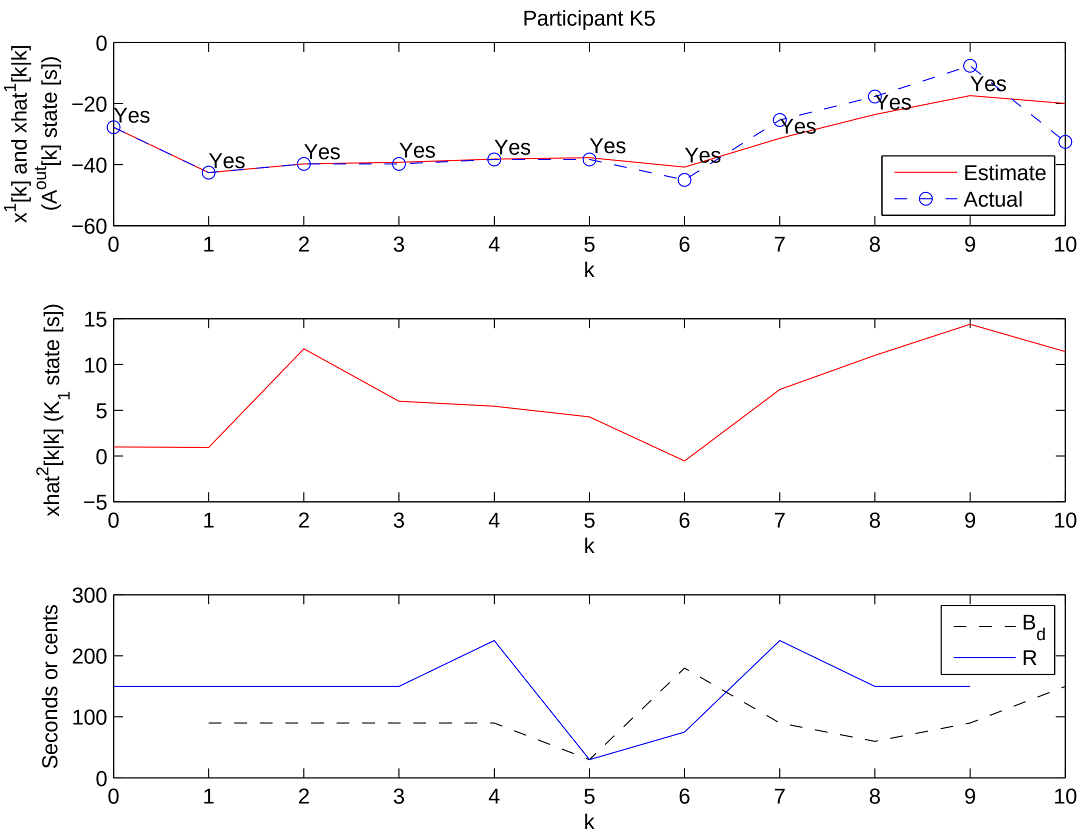

Figure C.6: Data from Participant K5. Participant specific details: $\hat{\mu}_2 = 0.9583$ sec/mm, $\hat{\mu}_1 = 87.792$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -27.7907 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 11.3986$ sec, target sound: Longing Baby, session start time: 12:30 PM.

Participant K5 accepted all the offers during their session. Reviewing the participant's behaviour and attitude in the first subplot of Figure C.6, we see that attitude increases following seven of the accepted offers. The decrease in attitude from $k = 0$ to $k = 1$ may be caused by the learning effects discussed for participant G2. More specifically, the participant may have adjusted their rating of the target sound after having listened to all the sounds. This further supports the notion that initial attitude measurements are noisier than subsequent measurements.

Subsequent measurements show small changes in attitude that are presumably within the range of measurement noise. Examining the larger trend, attitude increased from the first few measurement of the session to the last few measurements of the session. Given that all offers were accepted, the trend in attitude is consistent with the model.

The second subplot of Figure C.6 shows that $\hat{K}_1$ is positive. The estimate appears to be changing more rapidly than other participants. The estimate of K_1 would greatly benefit from additional samples.

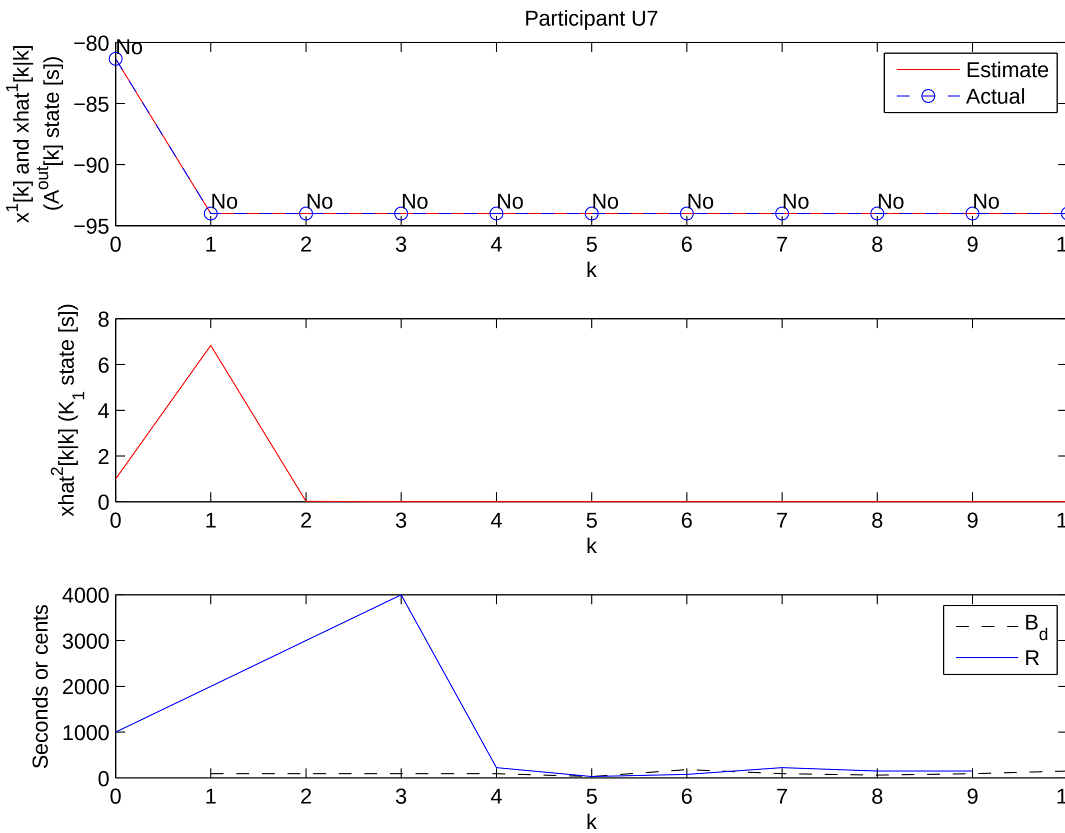

Figure C.7: Data from Participant U7. Participant specific details: $\hat{\mu}_2 = 1.8077$ sec/mm, $\hat{\mu}_1 = 5.654$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} -81.3465 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 0.0004$ sec, target sound: Longing Baby, session start time: 2:30 PM.

Participant U7 declined all the offers during their session. Plotting the participant's behaviour and attitude, shown in the first subplot of Figure C.7, reveals a decrease in attitude after declining the first offer. This single change in attitude, which is large and beyond the range of expected measurement noise, supports the model. The decrease in attitude saturates the measurement scale. The participant's attitude measurements remain saturated for all future attitude measurements. Beyond the first offer, we can speculate that the participant's attitude continued to decrease. This provides limited support for the model.

The second subplot shows minimal change in $\hat{K}_1$. The lack of observable change in attitude, due to saturation of the measurement, results in the EKF's estimate of K_1 being invalid. This data set cannot be used to verify or comment on the EKF's effectiveness.

In an attempt to have the participant accept offers and also move their attitude out of saturation, the researcher increased the rewards during the first four offers. This can be seen in the third subplot in Figure C.7; the first offer was \$10 and the fourth offer was the highest at \$40. These efforts were unsuccessful as the participant continued to decline the rewards and their attitude measurement remained saturated. It is interesting to note that these offers were the highest out of all participants.

C.3 Positive Attitude Participant Data

The data from sessions in which the target sound was initially positive is presented in the following pages. The data is presented in chronological order. The desired behaviour values used for the first four participants are shown under the heading "Set 1" in C.1, while the last four participants were offered the values under the heading "Set 2".

Table C.1: Desired behaviour sets.

k	Set 1 $B_d[k+1]$ (s)	Set 2 $B_d[k+1]$ (s)
0,1,2,3	90	180
4	30	60
5	180	360
6	90	180
7	60	120
8	90	180
9	150	300

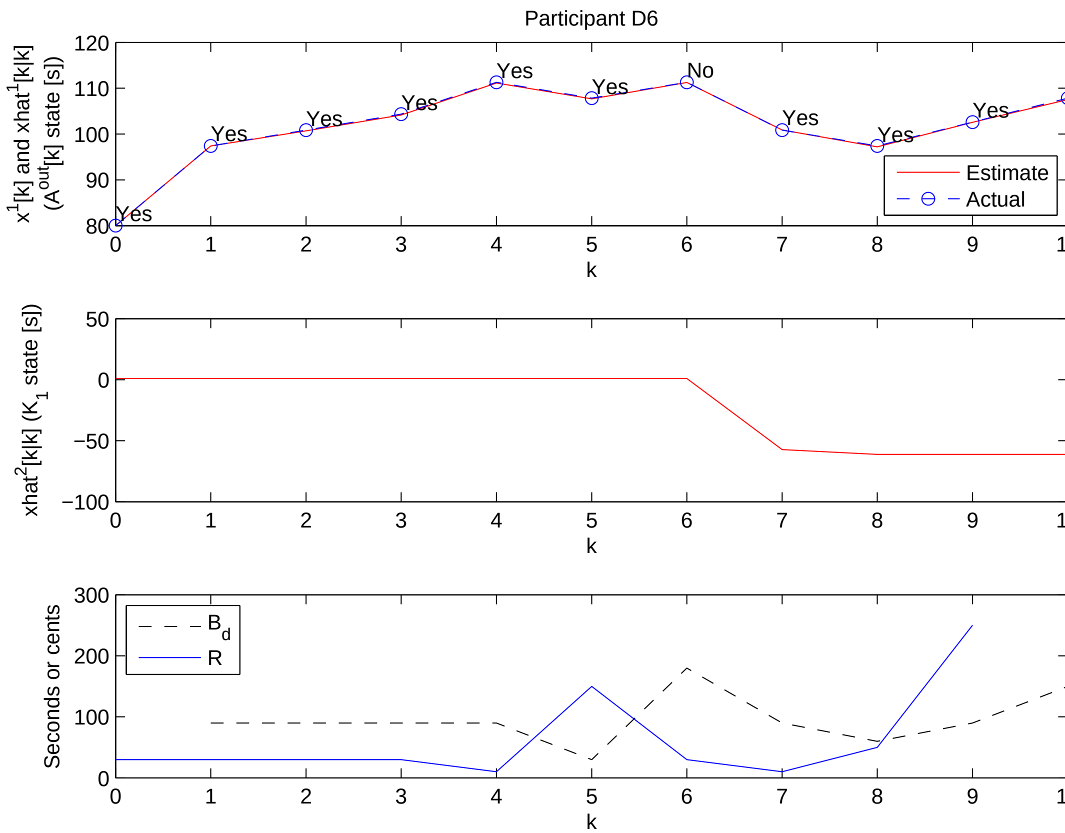

Figure C.8: Data from Participant D6. Participant specific details: $\hat{\mu}_2 = 3.4783$ sec/mm, $\hat{\mu}_1 = 16.277$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 80.0009 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 1.0000$ sec, target sound: Sunny Day, session start time: 11:20 AM.

Participant D6 accepted all but one offer in their session. As a result it is extremely difficult to make any strong conclusions regarding cognitive dissonance when attitude is positive.

The estimate $\hat{K}_1$ does not change until an offer is accepted. This is because K_1 is not part of the overjustification dynamics that arise when when offers are accepted. Further, as the EKF has only one usable sample, $\hat{K}_1$ is not a reliable estimate. Additionally, the single usable sample did not move in the direction expected. No conclusions can be made regarding the effectiveness of the EKF when applied to participant D6's data.

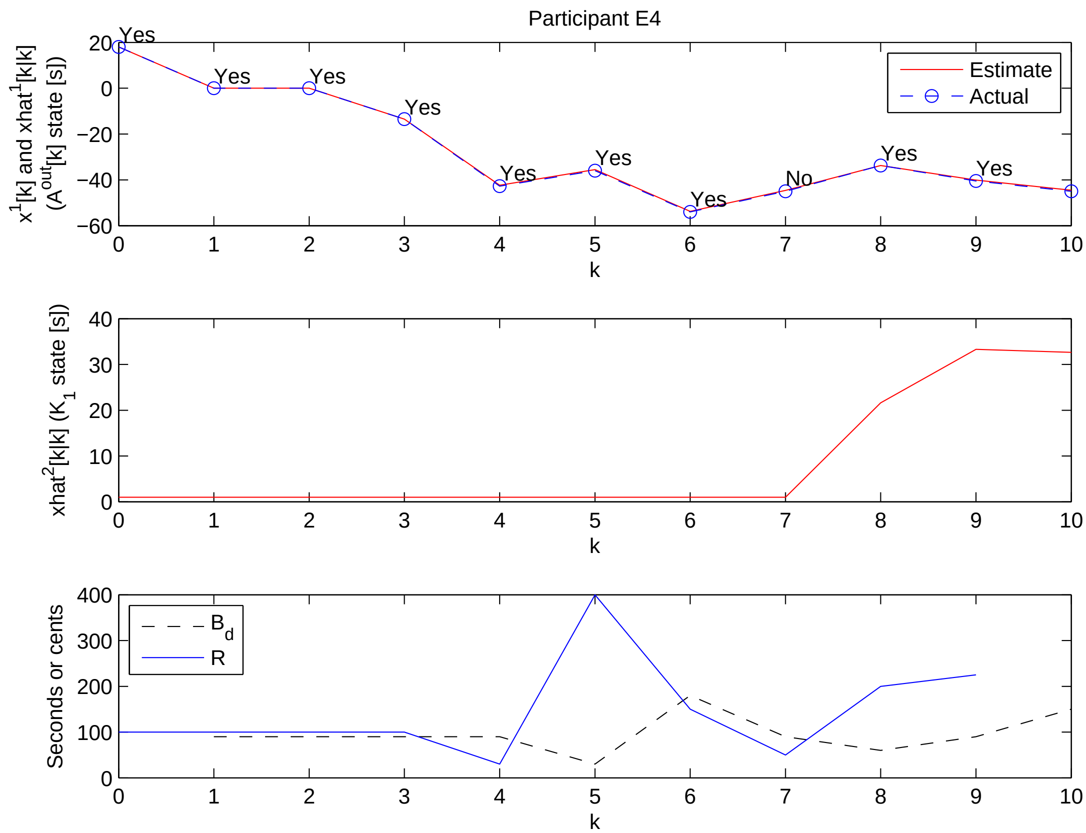

Figure C.9: Data from Participant E4. Participant specific details: $\hat{\mu}_2 = 2.2500$ sec/mm, $\hat{\mu}_1 = 30.000$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 18.0000 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 32.6571$ sec, target sound: Applause, session start time: 11:00 AM.

Participant E4's attitude is driven negative through the offers. The attitude decreases to 0 after accepting the initial offer at time $k = 0$. This is not the behaviour sought in this experiment. It it is possible overjustification effects are at play during this first offer. The later offers violate the assumption that attitude is positive.

The estimate $\hat{K}_1$ is unreliable as there is only one sample from which the EKF may divulge information regarding $\hat{K}_1$. More importantly, the simplified model is not even applicable when the offer is declined as attitude is negative rather than positive. No conclusions are made based on participant E4's data.

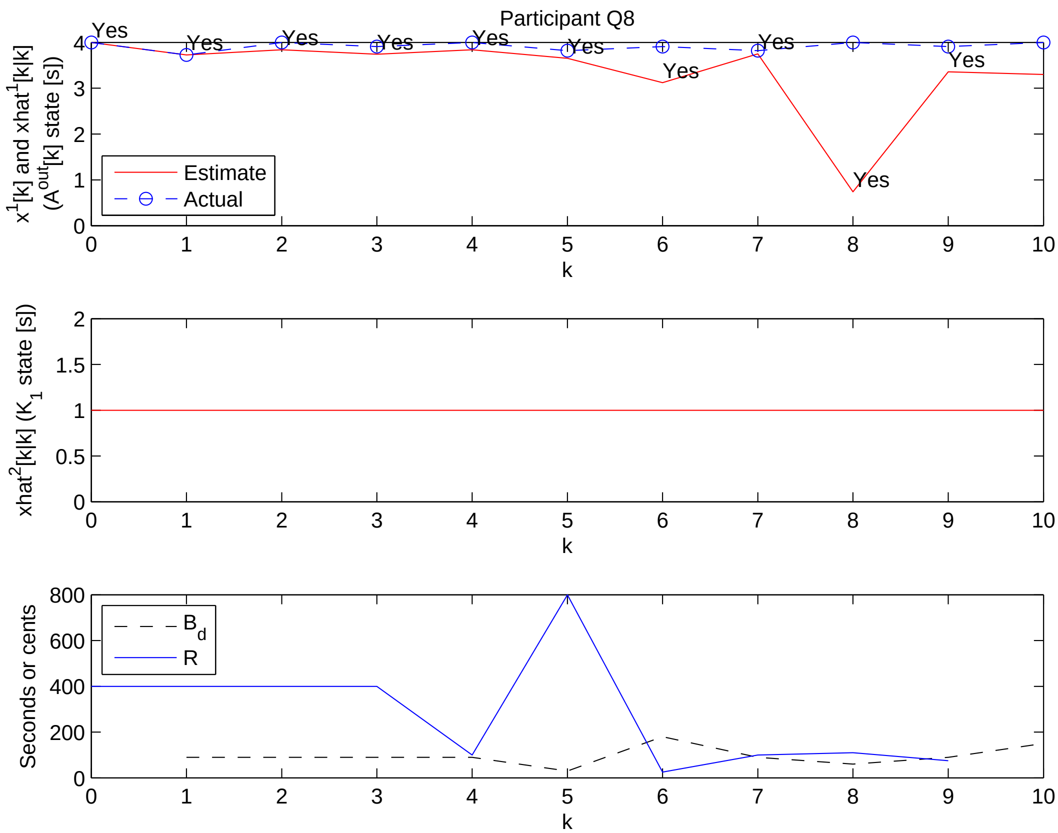

Figure C.10: Data from Participant Q8. Participant specific details: $\hat{\mu}_2 = 0.0909$ sec/mm, $\hat{\mu}_1 = 6.291$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 3.9996 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 1.0000$ sec, target sound: Ocean Wave, session start time: 2:40 PM.

Participant Q8 accepted all the rewards; this is the exact opposite of the behaviour the experiment aims to produce. As a result, the data collected is unsuitable for validating the effects of cognitive dissonance when attitude is positive.

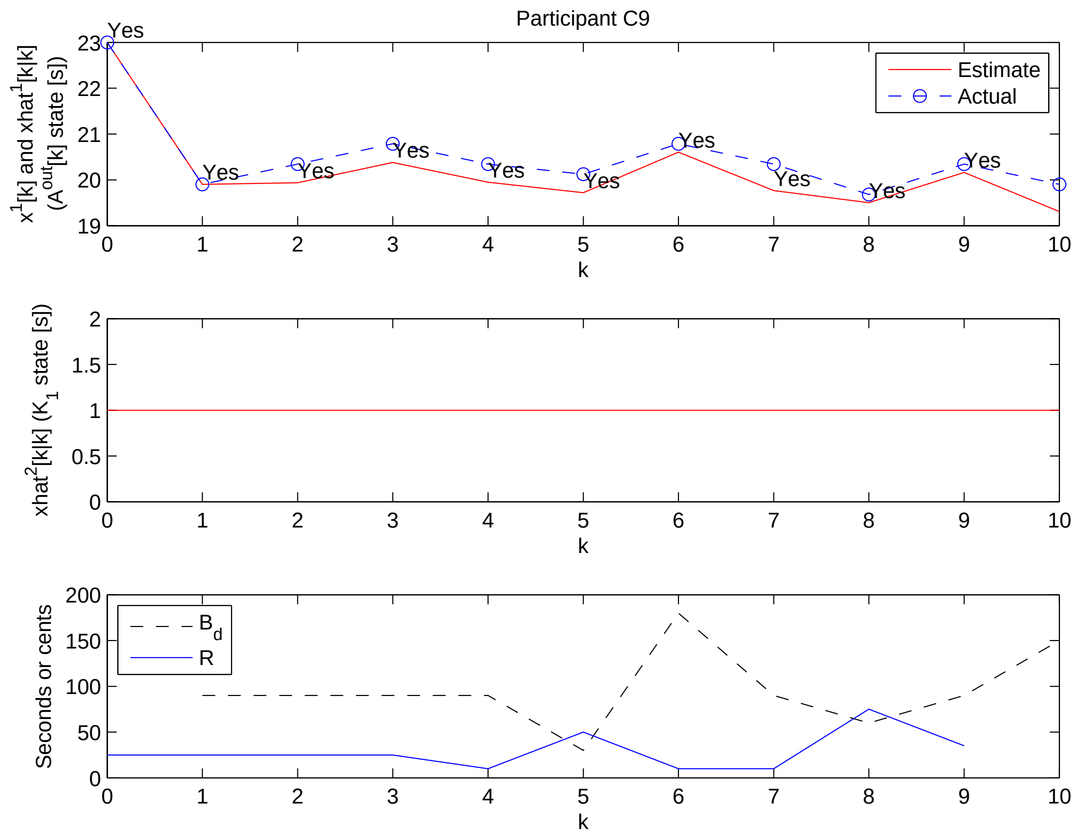

Figure C.11: Data from Participant C9. Participant specific details: $\hat{\mu}_2 = 0.4423$ sec/mm, $\hat{\mu}_1 = 27.667$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 22.9996 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 1.0000$ sec, target sound: Ocean Wave, session start time: 4:00 PM.

Participant C9 accepted all the rewards, and thus failed to excite the effects of cognitive dissonance. As a result, the data collected is unsuitable for validating cognitive dissonance's effects when attitude is positive. Further, the EKF cannot estimate parameters related to cognitive dissonance when rewards are accepted.

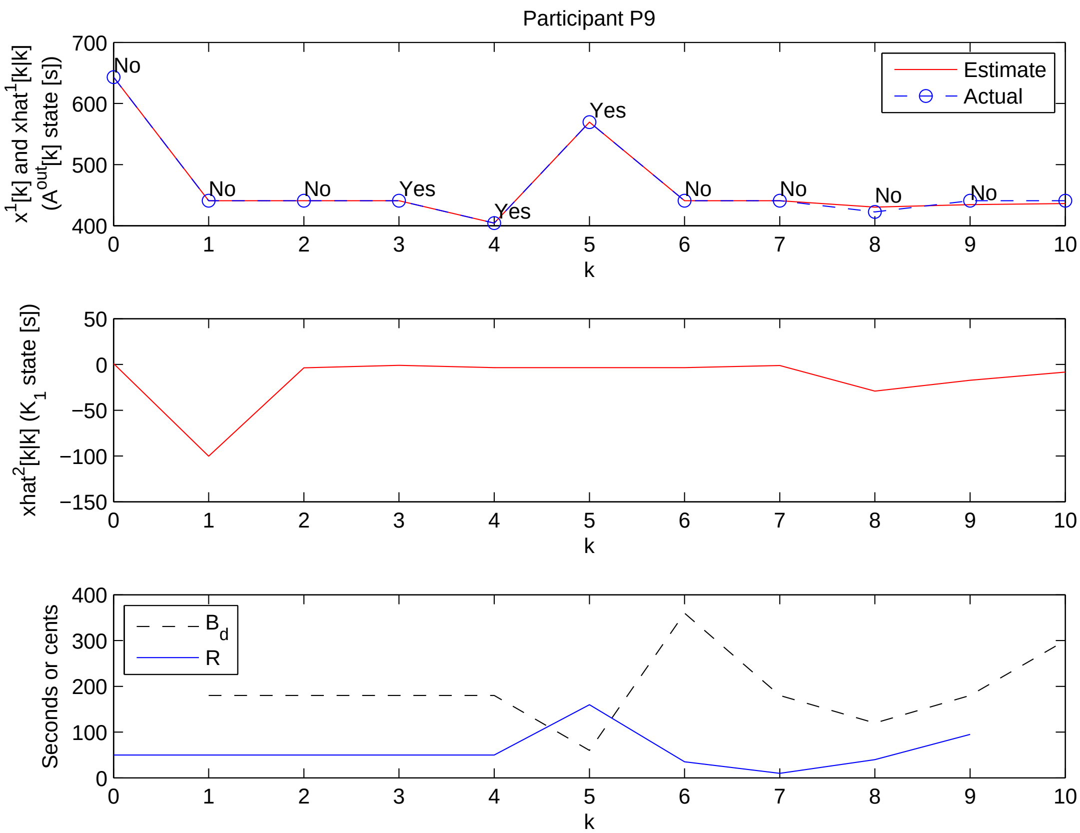

Figure C.12: Data from Participant P9. Participant specific details: $\hat{\mu}_2 = 18.3750$ sec/mm, $\hat{\mu}_1 = 112.688$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 643.1250 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = -8.3028$ sec, target sound: Sunny Day, session start time: 1:00 PM.

Participant P9 declined three of the first four offers. In the interval $K \in [0, 3]$, the participant's attitude decreased or remained constant. This opposes the the trend predicted by the model when the attitude is positive and an offer is declined. The participant declines the final four offers and we see their attitude increases only after the ninth offer. This increase in attitude is preceded by a minor decrease in attitude; it is possible this was simply a noisy measurement.

Participant P9's attitude remained above the desired behaviour during the entire experiment. This implies that the participant has an intrinsic desire to listen to the sound that is higher than the desired behaviour of the offer. It is expected that the participant will accept the offer and listen for more than the desired behaviour. This is direct opposition of our model. Further, the condition $A_{out}[k] < B_d[k]$, that is required to calculate a positive reward from an offer, is not satisfied.

The researcher made two errors during participant P9's session. First, he failed to realize that the condition $A_{out}[k] < B_d[k]$ was not satisfied. Second, he incorrectly calculated the rewards. If the reward had been calculated correctly, the result would have been negative indicating to the researcher that that $A_{out} \not< B_d[k]$.

With all of the above observation considered, it appears that the model does not reflect the participant's attitudes and behaviours well.

The estimate $\hat{K}_1$ is negative. This follows from the fact that attitude typically decreased following declined offers. Not all measurements were usable by the EKF as three of the ten offers were accepted. No comments on the convergence of $\hat{K}_1$ are made. With $\hat{K}_1$ negative, it weakly opposes the model when attitude is positive.

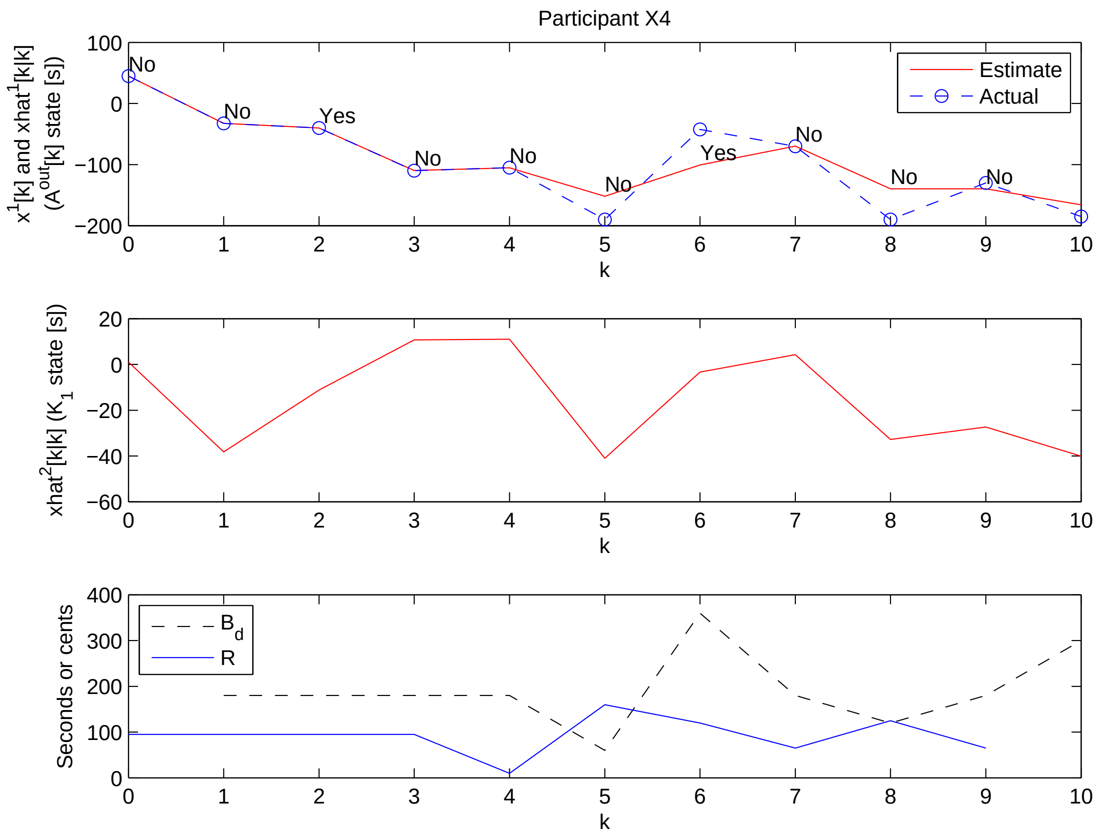

Figure C.13: Data from Participant X4. Participant specific details: $\hat{\mu}_2 = 5.0000$ sec/mm, $\hat{\mu}_1 = 70.000$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 45.0000 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = -40.1030$ sec, target sound: Ocean Wave, session start time: 2:30 PM.

Participant X4 declined the first two offers; however, their attitude did not increase. After declining the first offer, the participant's attitude decreased to be negative and remained there for the remainder of the session. With attitude negative for the majority of the session, the assumptions made in the model and the reward sizing are violated. The model no longer applies and no conclusions can be made.

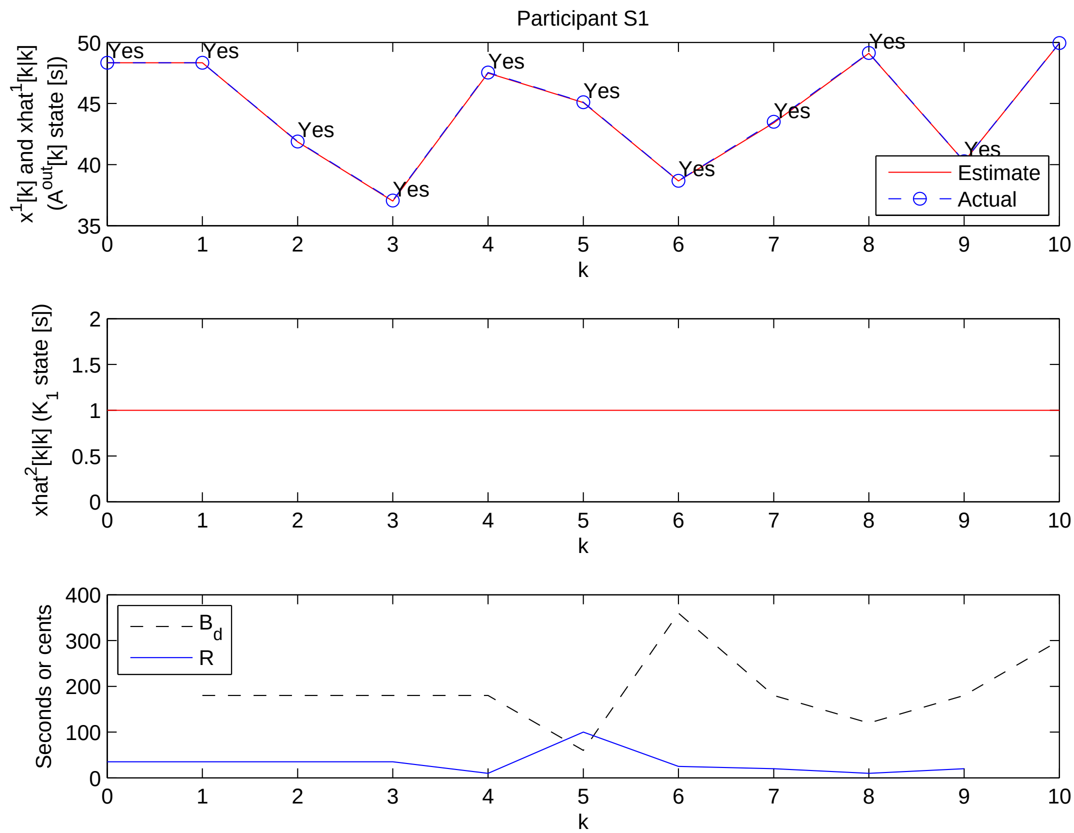

Figure C.14: Data from Participant S1. Participant specific details: $\hat{\mu}_2 = 1.6111$ sec/mm, $\hat{\mu}_1 = 197.333$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 48.3330 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = 1.0000$ sec, target sound: Ocean Wave, session start time: 11:00 AM.

Participant S1 accepted all of the offers. In doing so, cognitive dissonance does not apply and we cannot assess the participant's attitude against the model.

The EKF is unable to estimate parameters related to cognitive dissonance when rewards are accepted as cognitive dissonance does apply in these scenarios. As no usable measurements are available, $\hat{K}_1$ remains constant.

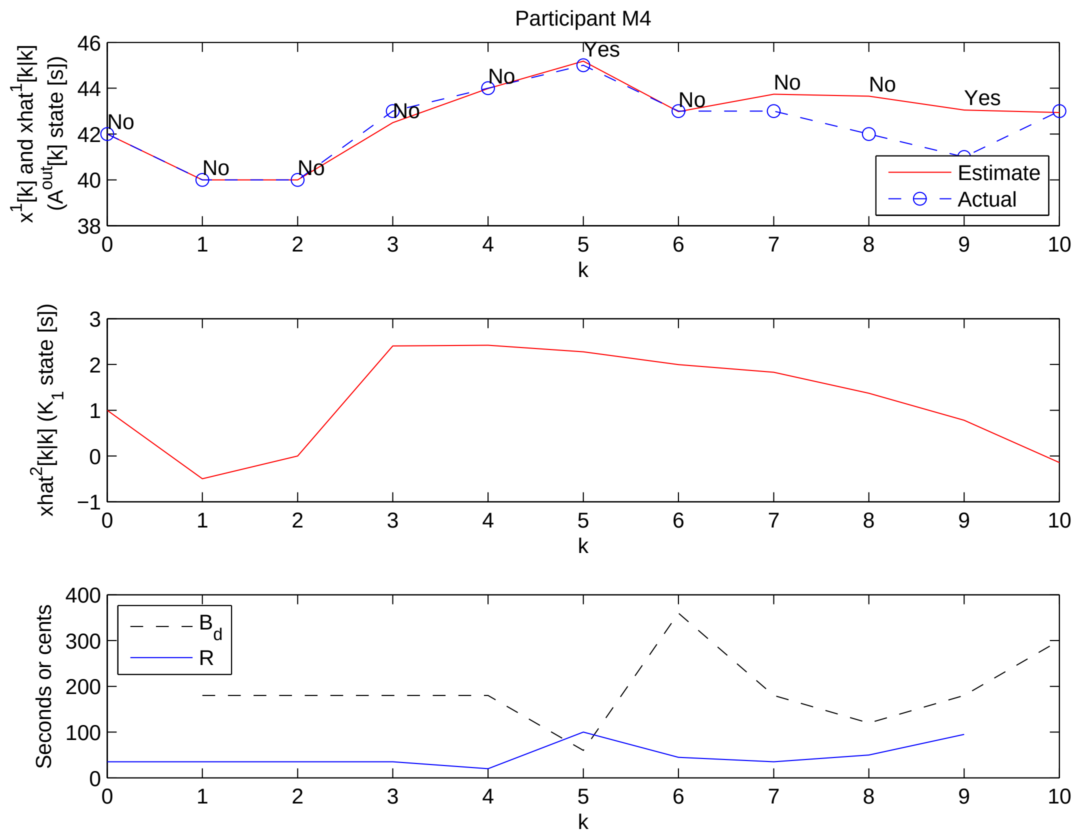

Figure C.15: Data from Participant M4. Participant specific details: $\hat{\mu}_2 = 1.0000$ sec/mm, $\hat{\mu}_1 = 190.000$ sec/\$, $\hat{x}_{[0|0]} = \begin{bmatrix} 42.0000 & 1 & 75 \end{bmatrix}^T$, final $\hat{K}_1 = -0.1455$ sec, target sound: Sunny Day, session start time: 12:00 PM.

Participant M4 declined eight offers. Of these eight offers, three were followed by an increase in attitude, three were followed by a decrease in attitude, and two were followed with no change in attitude. There is equal evidence for and against the model; thus, it is difficult to confidently conclude anything.

The estimate of K_1 ends on a negative value. Additionally, it does not appear to have converged. This would suggest the model is unable to predict attitudes following behaviour. It should be noted that not all measurements were usable as some offers were accepted. Having fewer measurements is known to produce less accurate estimates.

Appendix D

Study Materials

This appendix contains the documents used in the study. Minor changes to the formatting of these documents were made while incorporating these documents into this thesis.

D.1 Study

This section contains the complete study. The target sound, offers, and other blank spaces were filled in by the researcher as the study was run. Participants were given one sheet of paper at a time, the horizontal rule below indicates a new page in the study.

University of Waterloo Thoughts and Behaviour Experiment

Welcome, participant!

This study will last exactly 60 minutes; the person conducting the study (the researcher) will indicate the study is complete after this time has passed. During the study, please refrain from using electronic devices (e.g., phones, computers, etc.). The researcher is unable to answer any questions once the study begins to ensure they are following the same protocol for each participant. Do you have any questions right now? You can ask questions at the end of the study.

Take your time in all the tasks.

Remember that participation is voluntary and that you may withdraw from the study at any time.

Please write here your participant code:

On the next page you will rate how pleasant or unpleasant various sounds are after listening to each one for 30 seconds.

When you are ready to begin, put on the provided headphones and ask the researcher to play the first sound. After listening to the sound, rate the sound by placing a tick mark on the scale. The researcher will **play the sound for 30s**; you may listen to the sound again by asking the researcher to play the sound. Listen and rate each sound.

Note: Please let the researcher know if the volume is too loud or too quiet.

Applause

Unpleasant Neutral Pleasant

Female Laughter

Unpleasant Neutral Pleasant

Sunny Day

Unpleasant Neutral Pleasant

Drill

Unpleasant Neutral Pleasant

Alien Buzzer

Unpleasant Neutral Pleasant

Longing Baby

Unpleasant Neutral Pleasant

Digital Alarm

Unpleasant Neutral Pleasant

Ocean Wave

Unpleasant Neutral Pleasant

Fill in the blanks below:

$26 - 62 =$ __________

$60 + 16 =$ __________

$47 \times 4 =$ __________

$55 \div 11 =$ __________

The letter "n" appears __________ times in the statement below.

The beige hue on the waters of the loch impressed all, including the French queen, before she heard that symphony again, just as young Arthur wanted.

How much money would it take for you to listen to the _______________ sound for 60 seconds?

$____.______

Read the following paragraph:

He stood over the body in the fading light, adjusting the hair and putting the finishing touches to the simple toilet, doing all mechanically, with soulless care. And still through his consciousness ran an undersense of conviction that all was rightthat he should have her again as before, and everything explained. He had had no experience in grief; his capacity had not been enlarged by use. His heart could not contain it all, nor his imagination rightly conceive it. He did not know he was so hard struck; that knowledge would come later, and never go. Grief is an artist of powers as various as the instruments upon which he plays his dirges for the dead, evoking from some the sharpest, shrillest notes, from others the low, grave chords that throb recurrent like the slow beating of a distant drum. Some natures it startles; some it stupefies. To one it comes like the stroke of an arrow, stinging all the sensibilities to a keener life; to another as the blow of a bludgeon, which in crushing benumbs.

In the previous set of math questions, how many questions were there?

__________ questions

The __________________ sound was one of the more pleasant sounds you rated. Listen to this sound for as long as you wish. Indicate to the researcher when you would like to begin listening and then once again when you would like to stop listening.

Fill in the blanks below:

$4 \times 27 =$ __________

$62 \div 2 =$ __________

$47 - 97 =$ __________

$12 + 91 =$ __________

$74 + 54 =$ __________

$67 \times 3 =$ __________

You have the option of listening to the _______________ sound for __________ seconds. If you accept, you will receive $ ____.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

In the last set of math questions, how many division questions were there?

__________ questions

How many words are in the statement below?

Exploring the zoo, we saw every kangaroo jump and quite a few carried babies

__________ words

Rate the _______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant Neutral Pleasant

The letter "m" appears __________ times in the statement below.

The beige hue on the waters of the loch impressed all, including the French queen, before she heard that symphony again, just as young Arthur wanted.

Do you agree or disagree with the following statement:

A litre of milk should cost more than a litre of gasoline.

Indicate your choice by placing a tick on the scale below

Strongly Disagree Neutral Strongly Agree

You have the option of listening to the _______________ sound for __________ seconds. If you accept, you will receive $ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Fill in the blanks below:

$42 + 11 =$ __________

$15 \times 7 =$ __________

$48 \div 4 =$ __________

$11 + 101 =$ __________

Rate the _______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant Neutral Pleasant

Read the following paragraph:

Neptune is the eighth and farthest planet from the Sun in the Solar System. It is the fourth-largest planet by diameter and the third-largest by mass. Among the gaseous planets in the Solar System, Neptune is the most dense. Neptune is 17 times the mass of Earth and is slightly more massive than its near-twin, Uranus, which is 15 times the mass of Earth but not as dense. Neptune orbits the Sun at an average distance of 30.1 astronomical units. Named after the Roman god of the sea, its astronomical symbol is , a stylised version of the god Neptune's trident.

The letter "i" appears __________ times in the statement below.

Pack my box with five dozen liquor jugs.

You have the option of listening to the ______________ sound for __________ seconds. If you accept, you will receive \$ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Fill in the blanks below:

$62 - 46 =$ __________

$6 + 17 =$ __________

$7 \times 11 =$ __________

$114 \div 3 =$ __________

Rate the ______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant Neutral Pleasant

Do you agree or disagree with the following statement:

The Ontario government should focus on renewable energy production. (Renewable energy production includes solar energy, wind power, hydropower, and geothermal energy.)

Indicate your choice by placing a tick on the scale below:

Strongly Disagree — Neutral — Strongly Agree

You have the option of listening to the ______________ sound for __________ seconds. If you accept, you will receive $ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Fill in the blanks below:

$15 - 39 =$ __________

$42 \div 14 =$ __________

$26 + 62 =$ __________

$54 \times 2 =$ __________

$12 \times 9 =$ __________

Rate the ______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant — Neutral — Pleasant

Read the following paragraph.

The point was we took this all very seriously. They had taken away our flowers, our summer days, and our bonuses, we were on a wage freeze and a hiring freeze, and people were flying out the door like so many dismantled dummies. We had one thing still going for us: the prospect of a promotion. A new title: true, it came with no money, the power was almost always illusory, the bestowal a cheap shrewd device concocted by management to keep us from mutiny, but when word circulated that one of us had jumped up an acronym, that person was just a little quieter that day, took a longer lunch than usual, came back with shopping bags, spent the afternoon speaking softly into the telephone, and left whenever they wanted that night, while the rest of us sent e-mails flying back and forth on the lofty topics of Injustice and uncertainty.

Listen to the ______________ sound.

Rate this sound by placing a tick on the scale:

Unpleasant Neutral Pleasant

You have the option of listening to the ______________ sound for _________ seconds. If you accept, you will receive $ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

The letter "m" appears _________ times in the statement below:

Now listen! This stock could help you make huge amounts of money in weeks!

Fill in the blanks below:

$22 - 1 =$ _________

$58 + 92 =$ _________

$8 \times 15 =$ _________

$17 + 12 =$ _________

$76 + 13 =$ _________

Rate the ______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant Neutral Pleasant

In the last set of math questions, how many addition questions were there?

_________ questions

You have the option of listening to the _______________ sound for __________ seconds. If you accept, you will receive $ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

How many words are in the statement below?

Pack my box with five dozen liquor jugs.

__________ words

Fill in the blanks below.

$2 + 109 =$ __________

$89 + 29 =$ __________

$10 \times 15 =$ __________

$16 - 12 =$ __________

$48 \div 8 =$ __________

$8 \times 6 =$ __________

Rate the _______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant — Neutral — Pleasant

The letter "n" appears __________ times in the statement below:

The beige hue on the waters of the loch impressed all, including the French queen, before she heard that symphony again, just as young Arthur wanted.

Have you read the statement on the previous page before? Answer "yes" or "no."

____

You have the option of listening to the _______________ sound for __________ seconds. If you accept, you will receive $ ___._____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Read the following paragraph:

It was but a hurried parting in a common street, yet it was a sacred remembrance to these two common people. Utilitarian economists, skeletons of schoolmasters, Commissioners of Fact, genteel and used-up infidels, gabblers of many little dogs-eared creeds, the poor you will have always with you. Cultivate in them, while there is yet time, the utmost graces of the fancies and affections, to adorn their lives so much in need of ornament; or, in the day of your triumph, when romance is utterly driven out of their souls, and they and a bare existence stand face to face, Reality will take a wolfish turn, and make an end of you.

Rate the _______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant | Neutral | Pleasant

In the last set of math questions, how many questions were there?

__________ questions

Listen to the _______________ sound.

Rate this sound by placing a tick on the scale:

Unpleasant | Neutral | Pleasant

Read the following statement:

Exploring the zoo, we saw every kangaroo jump and quite a few carried babies.

You have the option of listening to the ______________ sound for __________ seconds. If you accept, you will receive $ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Fill in the blanks below:

$60 + 16 =$ __________

$26 \times 5 =$ __________

$47 - 4 =$ __________

$55 + 11 =$ __________

Rate the ______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant | Neutral | Pleasant

Do you agree or disagree with the following statement?

A litre of gasoline should cost more than a litre of milk.

Indicate your choice by placing a tick on the scale below.

Strongly Disagree | Neutral | Strongly Agree

You have the option of listening to the ______________ sound for __________ seconds. If you accept, you will receive $ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Fill in the blanks below:

$20 - 59 =$ __________

$17 - 5 =$ __________

$47 + 39 =$ __________

$55 \times 3 =$ __________

$64 \div 8 =$ __________

$72 - 12 =$ __________

Rate the _______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary.

Unpleasant — Neutral — Pleasant

Read the following paragraphs:

Ever since the child had learned to walk he had been his mother's and father's despair and delight, for there never was such a boy for wandering, for climbing up things, for getting into and out of things. That night, he had been woken by the sound of something on the floor beneath him falling with a crash. Awake, he soon became bored, and had begun looking for a way out of his crib. It had high sides, like the walls of his playpen downstairs, but he was convinced that he could scale it. All he needed was a step...

He pulled his large, golden teddy bear into the corner of the crib, then, holding the railing in his tiny hands, he put his foot onto the bear's lap, the other foot up on the bear's head, and he pulled himself up into a standing position, and then he half-climbed, half-toppled over the railing and out of the crib.

You have the option of listening to the _______________ sound for __________ seconds. If you accept, you will receive \$ ___.____. You may ask the researcher to play the sound if you do not recall what it sounded like.

Fill in the blanks below:

$14 \times 6 =$ __________

$62 - 16 =$ __________

$48 + 55 =$ __________

$66 - 72 =$ __________

The letter "s" appears __________ times in the paragraph below:

Back in June we delivered oxygen equipment of the same size.

Rate the _______________ sound by placing a tick on the scale. You may ask the researcher to listen to it if necessary:

Unpleasant Neutral Pleasant

The letter "m" appears __________ times in the statement below:

The beige hue on the waters of the loch impressed all, including the French queen, before she heard that symphony again, just as young Arthur wanted.

D.2 Recruitment Poster and Flyer

Department of Electrical and Computer Engineering
University of Waterloo

PARTICIPANTS NEEDED FOR RESEARCH IN SYSTEM MODELLING

We are looking for volunteers to take part in a study on thoughts and behaviour.

As a participant in this study, you will be asked to complete simple activities such as reading short passages, listening to sound clips, solving simple math problems, and rating these activities.

Your participation would involve 1 session lasting 60 minutes. Sessions will be held throughout January & February 2015.

Pizza and a drink will be provided at the end of the session.

For more information about this study, or to volunteer for this study, please contact Nikesh Parsotam at nparsotam@uwaterloo.ca

This study has been reviewed by, and received ethics clearance through a University of Waterloo Research Ethics Committee.

Alternatively, scan to volunteer for this study!

D.3 Recruitment Email

Hello,

My name is Nikesh Parsotam and I am a MASc student working under the supervision of Professor Dan Davison in the Department of Electrical and Computer Engineering at the University of Waterloo. The reason that I am contacting you is that we are conducting a study that examines the dynamics of peoples' thoughts and behaviour. We are currently seeking participants in this study.

Participation in this study involves coming into our laboratory in the EIT building and completing a variety of simple activities. Activities include reading short passages, solving basic math problems, listening to sound clips, and rating these activities. This study will take 1 hour of your time. In appreciation of your participation we will provide pizza and a drink.

Sessions are available daily between 9 am and 5 pm until Monday, January 19, 2015.

If you are interested in participating, please contact me at nparsotam@uwaterloo.ca and list 2 dates and times, for when you would like to participate. I will then send a confirmation email, and provide you with further information concerning the location of the study. If you have to cancel your appointment, please email me at nparsotam@uwaterloo.ca.

I would like to assure you that the study has been reviewed and received ethics clearance through a University of Waterloo Research Ethics Committee. The final decision about participation is yours.

Sincerely,
Nikesh Parsotam

D.4 Information and Consent Form

Date:	January 2015 - February 2015
Title of Project:	Parameter Estimation of a Human Behaviour Model
Faculty Supervisor:	Dan Davison, Electrical and Computer Engineering, x35338, ddavison@uwaterloo.ca
Student Investigator:	Nikesh Parsotam, Electrical and Computer Engineering, x37257, nparsotam@uwaterloo.ca

Study Overview
I am a Master's student in the Department of Electrical and Computer Engineering at the University of Waterloo conducting research under the supervision of Dr. Dan Davison. You are invited to participate in a study examining the dynamics of a person's thoughts and behaviour.

What You Will Be Asked to Do
As a participant in this study, you will be asked to complete a number of simple tasks. These tasks include reading short passages, solving basic math problems, listening to sound clips, and rating these activities. The reading exercises involve reading up to a half page. The math questions involve addition, subtraction, multiplication and division. You may be asked to listen to sounds for up to ten minutes. The study will be administered via pen and paper, except for listening to sound clips which will be done using a computer and headphones. The headphones are sanitized between sessions using an alcohol based sanitizer.

During the study, the researcher is not able to answer questions as to ensure that the same protocol is followed for each participant. We will invite participants to ask any questions at the end of the study. Additionally, feel free to ask the research questions you may have before beginning the study.

Participation and Remuneration
Participation in this study is voluntary, and the study, once started, will take 60 minutes of your time. As a gesture of appreciation for your participation, we will accompany you to the EIT Café and buy for you a slice of pizza and a drink (i.e. soda, juice, coffee, tea, water) at the end of the study.

If you do not wish to complete any of the activities or tasks, please indicate to the researcher that you would like to withdraw from this study. Alternatively, you may decide to withdraw at any time by informing the researcher.

Benefits of the Study
The benefits of this study include validating a proposed model and establishing confidence

in a parameter estimation method. There are no personal benefits to participation.

Risks to Participation in the Study
We want you to be aware of the possible risks associated with participation in this research. You may find the volume level too loud, or too low; simply ask the researcher to adjust the volume as required.

Confidentiality
All information you provide is considered completely confidential; indeed, your name and email address will not be included or in any other way associated, with the data collected in the study. You will not be identified individually in any way in any written reports of this research. Electronic data, with identifying information removed, will be kept indefinitely following publication of the research. Printed data will be securely stored in a locked room in EIT 3114 for 1 year then destroyed by confidential shredding; only researchers associated with this study have access to the data.

Questions and Research Ethics Clearance
If, after reviewing this letter, you have any questions about this study, or would like additional information to assist you in reaching a decision about participation, please feel free to ask the student investigator or the faculty supervisor listed at the top of this letter.

I would like to assure you that this study has been reviewed and received ethics clearance through a University of Waterloo Research Ethics Committee. However, the final decision about participation is yours. If you have any comments or concerns resulting from your participation in this study, please contact Dr. Maureen Nummelin, the Director, Office of Research Ethics, at 1-519-888-4567, Ext. 36005 or maureen.nummelin@uwaterloo.ca.

Thank you for your interest in our research and for your assistance with this project.

Consent of Participant

I have read the information presented in the information letter about a study being conducted by Nikesh Parsotam under the supervision of Dr. Dan Davison of the Department of Electrical and Computer Engineering at the University of Waterloo. I have had the opportunity to ask any questions related to this study, to receive satisfactory answers to my questions, and any additional details I wanted. I am aware that I may withdraw from the study without loss of participation credit at any time by advising the researchers of this decision.

This project has been reviewed by, and received ethics clearance through a University of Waterloo Research Ethics Committee. I was informed that if I have any comments or concerns resulting from my participation in this study, I may contact the Director, Office of Research Ethics, at 1-519-888-4567, x36005.

With full knowledge of all foregoing, I agree, of my own free will, to participate in this study. By signing this consent form, you are not waiving your legal rights or releasing the investigator(s) or involved institution(s) from their legal and professional responsibilities.

____
Print Name

____
Signature of Participant

____
Dated at Waterloo, Ontario

____
Witnessed

D.5 Debrief and Consent Form

Study Title:	Parameter Estimation of a Human Behaviour Model
Faculty Supervisor:	Dan Davison, Electrical and Computer Engineering, x35338, ddavison@uwaterloo.ca
Student Investigator:	Nikesh Parsotam, Electrical and Computer Engineering, x37257, nparsotam@uwaterloo.ca

We greatly appreciate your participation in our study, and thank you for spending the time helping us with our research. When you began the study, you were told that the purpose of this study was to examine the dynamics of peoples' thoughts and behavior. However, the study was more focused than we explained at the beginning. Specifically, in this study we were investigating how a reward would affect your attitude and behavior towards listening to an unpleasant or pleasant sound. A classical theory called cognitive dissonance theory predicts that people experience a pressure called cognitive dissonance pressure when attitude and behaviour towards a task are inconsistent and, moreover, that the person will change their attitude or behavior to try to reduce this pressure. Overjustification theory predicts that rewards given to perform an activity one enjoys or views as pleasant will result in a reduction in intrinsic enjoyment towards that activity. Each time you accepted or declined a reward offer, you experienced mild cognitive dissonance pressure or the effects of overjustification and we monitored how your attitude towards the sound changed over the duration of the session. We have devised a mathematical model that captures these effects and we are interested in validating the model and identifying key parameters of the model that are unique to each participant. The offers were based on your behaviour when asked to listen to a pleasant sound, your response to how much money would be required to have you listen to an unpleasant sound, and your behaviour throughout the study. Our goal was to have participants accept the offers without offering too large of a reward.

In this study, a series of activities was given for you to carry out. Many of these activities, such as the math, reading exercises and some opinion questions, were distractors included only to allow for sufficient time to elapse between reward offers and sound ratings. These distractors also masked the true goals of this study. Additionally, during recruitment as well as the information and consent period, we did not explain that you would be offered monetary rewards to listen to a sound during the study. Finally, we told you the study would last exactly 60 minutes. This was to reduce the urge for participants to hurry through the experiment. The researcher was told not to answer questions to minimize influences to a participant's behaviours and responses. We could not give you complete information about the study before your involvement because it may have influenced your choices during the study in a way that would make investigations of the research question invalid. The reason that we used partial disclosure in this study was that we needed your behaviour and attitudes to be unaffected by the study objectives. We apologize for omitting details and for providing you with fictional information about the purpose of and tasks in our study. We hope that you understand the need for partial disclosure now that the purpose of the study has been more fully explained to you. We would also like to assure you that most research does not involve the use of partial disclosure.

If any of the questions or exercises in this study caused you to feel uncomfortable, please feel free to contact Nikesh Parsotam, anytime, at nparsotam@uwaterloo.ca. You can also contact my faculty supervisor, Dan Davison, at 519-888-4567 x35338 or by email at ddavison@uwaterloo.ca. Also please feel free to contact Dr. Maureen Nummelin, the Director, Office of Research Ethics, at 1-519-888-4567, x36005 or maureen.nummelin@uwaterloo.ca if you have concerns or comments resulting from your participation.

The information you provided will be kept confidential by not associating your name with your responses. The data will be stored with all identifying or potentially identifying information removed. Electronic data will be stored indefinitely on password protected computers. Printed data will be kept in a locked room in EIT 3114 for 1 year (7 years for the receipt-of-reward form) then destroyed by confidential shredding. No one other than the researchers will have access to the data.

Once all the data are collected and analyzed for this project, I plan on sharing the experimental data (stripped of all personal identification information) with the research community through seminars, conferences, presentations and journal articles. If you are interested in receiving more information regarding the results of this study, or if you have any questions or concerns, please contact Nikesh Parsotam at the email address listed above. The study is expected to be complete by October 10, 2015.

Because the study involves some aspects that you were not told about before starting, it is very important that you not discuss your experiences with any other participants who potentially could be in this study until after the end of the term. If people come into the study knowing about our specific objectives, as you can imagine, it could influence their results, and the data we collect would be not be useable. Also, since you will be given a copy of this feedback letter to take home with you, please do not make this available to other students. Moreover, because some elements of the study are different from what was originally explained, we have another consent form for you to read and sign if you are willing to allow us to use the information that you have provided. This form is a record that the purpose of the study has been explained to you, and that you are willing to allow your information to be included in the study.

We really appreciate your participation, and hope that this has been an interesting experience for you.

POST-DEBRIEFING CONSENT FORM FOR STUDIES INVOLVING DECEPTION (IN-LAB)

Study Title:	Parameter Estimation of a Human Behaviour Model
Faculty Supervisor:	Dan Davison, Electrical and Computer Engineering, x35338, ddavison@uwaterloo.ca
Student Investigator:	Nikesh Parsotam, Electrical and Computer Engineering, x37257, nparsotam@uwaterloo.ca

During the debriefing session, I learned that it was necessary for the researchers to disguise the real purpose of this study. I realize that this was necessary since having full information about the actual purpose of the study might have influenced the way in which I responded to the tasks and this would have invalidated the results. Thus, to ensure that this did not happen, some of the details about the purpose of the study initially were not provided (or were provided in a manner that slightly misrepresented the real purpose of the study). However, I have now received a complete verbal and written explanation as to the actual purpose of the study and have had an opportunity to ask any questions about this and to receive acceptable answers to my questions.

I have been asked to give permission for the researchers to use my anonymized data (or information I provided) in their study, and agree to this request.

I am aware this study has been reviewed and received ethics clearance through a University of Waterloo Research Ethics Committee and I may contact the Director, Office of Research Ethics at 519-888-4567 x36005, if I have any concerns or comments resulting from my involvement in this study.

Participant's Name: ____

Participant's Signature: ____________________________________

Date: ____

Witness' Name: ____

Witness' Signature: _______________________________________

D.6 Participant Remuneration Form

In accordance with the University of Waterloo's procedures for remuneration, participants are asked to sign the "University of Waterloo Research Participants Acknowledgement of Receipt of Remuneration and Self-Declare Income" form upon receiving remuneration. The remuneration procedures and form may be found by following the link below.

Remuneration procedures: https://uwaterloo.ca/finance/guidance-procedures/procedures-info/remuneration-research-participants
Remuneration form: https://uwaterloo.ca/finance/sites/ca.finance/files/uploads/files/ParticipantRemunerationandSelfDeclarationFormFinal.pdf
Form issued: December 20, 2011

D.7 Ethics Form 101

ORE OFFICE USE ONLY

ORE #______________

APPLICATION FOR ETHICS REVIEW OF RESEARCH INVOLVING HUMAN PARTICIPANTS

Please remember to PRINT AND SIGN the form and forward with all attachments to the Office of Research Ethics, Needles Hall, Room 1024.

A. GENERAL INFORMATION

1. Title of Project: Parameter Estimation of a Human Behaviour Model

2. a) Principal and Co-Investigator(s)

NEW As of May 1, 2013, all UW faculty and staff listed as investigation must complete the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans Tutorial, 2nd Ed. (TCPS2) prior to submitting an ethics application. The tutorial takes at least three hours; it has start and stop features.

Name	Department	Ext:	e-mail:

2. b) Collaborator(s)

NEW As of May 1, 2013, all UW faculty and staff listed as investigation must complete the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans Tutorial, 2nd Ed. (TCPS2) prior to submitting an ethics application. The tutorial takes at least three hours; it has start and stop features.

Name	Department	Ext:	e-mail:

3. Faculty Supervisor(s)

NEW As of May 1, 2013, all UW faculty and staff listed as investigation must complete the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans Tutorial, 2nd Ed. (TCPS2) prior to submitting an ethics application. The tutorial takes at least three hours; it has start and stop features.

Name	Department	Ext:	e-mail:
Dan Davison	Electrical & Computer Engineering	35338	ddavison@uwaterloo.ca

4. Student Investigator(s)

Name	Department	Ext:	e-mail:	Local Phone #:
Nikesh Parsotam	Electrical & Computer Engineering	37257	nparsotam@uwaterloo.ca	4164646752

5. Level of Project: MASc **Specify Course:**

Research Project/Course Status: New Project\Course

6. Funding Status (If Industry funded and a clinical trial involving a drug or natural product or is medical device testing, then Appendix B is to be completed):

Is this project currently funded? Yes

- If Yes, provide Name of Sponsor and include the title of the grant/contract: NSERC : Decentralized Regulation, Tracking, and Disturbance Rejection for Multi-agent Systems: A Targeting Approach
- If No, is funding being sought OR if Yes, is additional funding being sought? No
- Period of Funding: 2012-2017

7. Does this research involve another institution or site? No
If Yes, what other institutions or sites are involved:

8. Has this proposal, or a version of it, been submitted to any other Research Ethics Board/Institutional Review Board? No

9. For Undergraduate and Graduate Research:

Has this proposal received approval of a Department Committee? Not Dept. Req.

10. a) Indicate the anticipated commencement date for this project: 10/10/2014

b) Indicate the anticipated completion date for this project: 10/10/2015

11. Conflict of interest: Appendix B is attached to the application if there are any potential, perceived, or actual financial or non-financial conflicts of interest by members of the research team in undertaking the proposed research.

B. SUMMARY OF PROPOSED RESEARCH

1. Purpose and Rationale for Proposed Research

a. Describe the purpose (objectives) and rationale of the proposed project and include any hypothesis(es)/research questions to be investigated. For a non-clinical study summarize the proposed research using the headings: Purpose, Aim or Hypothesis, and Justification for the Study. For a clinical trial/medical device testing summarize the research proposal using the following headings: Purpose, Hypothesis, Justification, and Objectives.
Where available, provide a copy of a research proposal. For a clinical trial/medical device testing a research proposal is required:
Purpose
The purpose of this study is to validate a proposed model of human attitude and identify key parameters in the model. The model links attitudes and behaviour. The dynamics explored in this study assume a simplified model, largely based on the theory of cognitive dissonance.

Hypothesis
The dynamics of a human's attitude towards a behaviour and the resulting behaviour are related to the rewards offered for producing the desired behaviour.

Justification
In the study, we will investigate whether the model correctly predicts the behaviour of participants being rewarded for producing a desired behaviour via a sequence of monetary rewards. Cognitive dissonance theory suggests that a pressure arises when attitudes and behaviours are inconsistent (e.g. dissonance pressure arises when a participant declines a

desirable reward that is offered to carry out an undesirable behaviour). In this study we create inconsistencies and observe the participants' behaviour and attitude. In this study, listening to an unpleasant sound will be the desired behaviour and a monetary reward is used. Confirmation of the hypothesis could enable a variety of applications, one of which would be promoting healthy choices.

Objectives
There are two desired outcome of this study. The first is to verify that the model's trends agree with actual human behaviour observed in this study. The second, provided the first outcome is achieved, is to use parameter estimation techniques to estimate three model parameters for each participant. Once these parameters are estimated, the model can be used to predict the participant's attitude and behaviour.

b. In lay language, provide a one paragraph (approximately 100 words) summary of the project including purpose, the anticipated potential benefits, and basic procedures used.
The purpose of this study is to validate a model of human attitude. Validation of the model could lead to development of applications that improve healthy life choices. Participants will be offered a reward to listen to a sound for a specific duration; shortly after their attitude will be measured through a paper questionnaire. This process will be iterated with varying rewards and durations.

C. DETAILS OF STUDY

1. Methodology/Procedures

a. Indicate all of the procedures that will be used. Append to form 101 a copy of all materials to be used in this study.

Survey(s) or questionnaire(s) (in person) Some are standardized.

b. Provide a detailed, sequential description of the procedures to be used in this study. For studies involving multiple procedures or sessions, provide a flow chart. Where applicable, this section also should give the research design (e.g., cross-over design, repeated measures design).
Through the use of printed instructions, participants are asked to do the following:

1. Listen to 8 sounds and rate the pleasantness or unpleasantness of each sound on a linear scale

2. Complete tasks that do not pertain to the objective of the study (e.g. math questions, counting exercises, reading exercises, provide their opinion)

3. State how much of a reward would be required to listen to an unpleasant sound (determined from 1) for 60 seconds

4. Answer questions that do not pertain to the objective of the study (e.g. math questions, counting exercises, reading exercises, provide their opinion)

5. Listen to a pleasant sound (determined from 1) for as long as they like

6. Answer questions that do not pertain to the objective of the study (e.g. math questions, counting exercises, reading exercises, provide their opinion)

7. Accept or decline a reward to listen to an unpleasant sound (determined from 1)

8. Answer questions that do not pertain to the objective of the study (e.g. math questions,

counting exercises, reading exercises, provide their opinion)

9. Rerate the unpleasant sound from 7

10. Answer questions that do not pertain to the objective of the study (e.g. math questions, counting exercises, reading exercises, provide their opinion)

(Steps 7 to 10 repeated 9 more times)

c. Will this study involve the administration/use of any drug, medical device, biologic, or natural health product? No

2. Participants Involved in the Study

a. Indicate who will be recruited as potential participants in this study.
UW Participants:
Undergraduate students
Graduate students
Faculty and/or Staff

b. Describe the potential participants in this study including group affiliation, gender, age range and any other special characteristics. Describe distinct or common characteristics of the potential participants or a group (e.g., a group with a particular health condition) that are relevant to recruitment and/or procedures. Provide justification for exclusion based on culture, language, gender, race, ethnicity, age or disability. For example, if a gender or sub-group (i.e., pregnant and/or breastfeeding women) is to be excluded, provide a justification for the exclusion.

Participants should be able to listen to sounds through standard headphones & read, write and speak in English. The targeted behaviour is listening to sounds and thus persons with difficulty hearing sounds are excluded. The study will be conducted in English; the questionnaire is written in English and the person conducting the survey is able to communicate in English.

To mitigate perceptions of influence or authority, we will avoid contacting participants that are currently being taught by the investigator (e.g. ECE students in 2A this term).

c. How many participants are expected to be involved in this study? For a clinical trial, medical device testing, or study with procedures that pose greater than minimal risk, sample size determination information is to be provided.

100

3. Recruitment Process and Study Location

a. From what source(s) will the potential participants be recruited?
UW undergraduate or graduate classes
Other UW sources: Handout flyers to pedestrians on campus
mailing lists of students at UW (e.g. graduate studies mailing list, faculty member mailing lists)

b. Describe how and by whom the potential participants will be recruited. Provide a copy of any materials to be used for recruitment (e.g. posters(s), flyers, cards, advertisement(s), letter(s), telephone, email, and other verbal scripts).
Potential participants will be recruited via email, posters, flyers. The student investigator will send emails to some or all of the mailing lists below. The student investigator will place posters on the various board across campus only if the owner of the bulletin boards approve. The student investigator will distribute flyers following the sample script in Appendix B. Mailing lists: ECE undergraduate students, ECE graduate students, and ECE faculty & staff members

c. Where will the study take place? On campus: EIT 3111

4. Remuneration for Participants

Will participants receive remuneration (financial, in-kind, or otherwise) for participation? Yes

If Yes, provide details:

Monetary rewards will be offered to have participants listen to a sound they rated as unpleasant for a period of 10s to 300s. The participant will be informed they may either accept or decline each reward when it is offered. The range of each individual reward ranges from $0.10 to $10 and there will be a total of 10 offers. All rewards are given shortly after they accept. Participants will be offered a slice of pizza and a drink (i.e. soda, juice, coffee, tea, water) from the EIT Café upon completion of the study or should they choose to withdraw at any point.

5. Feedback to Participants

Describe the plans for provision of study feedback and attach a copy of the feedback letter to be used. Wherever possible, written feedback should be provided to study participants including a statement of appreciation, details about the purpose and predictions of the study, restatement of the provisions for confidentiality and security of data, an indication of when a study report will be available and how to obtain a copy, contact information for the researchers, and the ethics review and clearance statement.

Written feedback will be provided to participants after the debriefing letter is provided. Participants are given the option to provide their contact information should they wish to receive the results related to the study.

The feedback letter is part of the debrief letter (attached in Appendix E)

D. POTENTIAL BENEFITS FROM THE STUDY

1. Identify and describe any known or anticipated direct benefits to the participants from their involvement in the project.

There is no direct benefit to the participants outside of the remuneration.

2.Identify and describe any known or anticipated benefits to the scientific community/society from the conduct of this study.

The scientific community will benefit from the outcome of this study as it will be used to validate a proposed behavioural model. Further, confidence in the parameter estimation model will be gained.

E. POTENTIAL RISKS TO PARTICIPANTS FROM THE STUDY

1. For each procedure used in this study, describe any known or anticipated risks/stressors to the participants. Consider physiological, psychological, emotional, social, economic risks/stressors. A study–specific current health status form must be included when physiological assessments are used and the associated risk(s) to participants is minimal or greater.

Minimal risks anticipated.

Participants may find the volume level to be too loud.

Participants' attitude towards the sounds used in the study may change in either direction.

2. Describe the procedures or safeguards in place to protect the physical and psychological health of the participants in light of the risks/stressors identified in E1.

Instructions are given instructions on how to adjust the volume. Participants have the option to withdraw from the study at any time.

F. INFORMED CONSENT PROCESS

1. What process will be used to inform the potential participants about the study details and to obtain their consent for participation?
Information letter with written consent form

2. If written consent cannot be obtained from the potential participants, provide a justification for this.

3. Does this study involve persons who cannot give their own consent (e.g. minors)? No

G. ANONYMITY OF PARTICIPANTS AND CONFIDENTIALITY OF DATA

1. Provide a detailed explanation of the procedures to be used to ensure anonymity of participants and confidentiality of data both during the research and in the release of the findings.
Participants' information (i.e. names) will be kept in hard copy form and only for record-keeping of consent and payments. Participants' data (i.e. attitudes) contain no identifiers to allow one to deduce the identity of a participant (i.e. it is inherently anonymized). To keep track of data for analysis purposes, an alpha numerical ID will be assigned to each data set. This alpha numerical ID will not be linked, at any time, to participant information. The participant's consent form is not attached in any way to the data gathered in the study. Participants who complete the consent form will continue on to the study. Shortly after completion of the participant's session, the consent form will be transferred to Dan Davison's office (EIT 3114) to be stored in a locked filing cabinet for a period of 1 year, after which they will be confidentially shredded. Participants will need to sign a receipt (of rewards accepted, see appendix E) upon completion or withdrawal from the study. Shortly after completion of the participant's session, receipts will be transferred to Dan Davison's office (EIT 3114) to be stored in a locked filing cabinet for a period of 7 years, after which they will be confidentially shredded. The hard copy of the anonymized study data will be stored in Dan Davison's office (EIT 3114) in a locked filing cabinet and will be available to the researchers only. The hard copy will be confidentially shredded after 1 year. An electronic copy will be created and will be stored on the researcher's password protected computer.

2. Describe the procedures for securing written records, video/audio tapes, questionnaires and recordings. Identify (i) whether the data collected will be linked with any other dataset and identify the linking dataset and (ii) whether the data will be sent outside of the institution where it is collected or if data will be received from other sites. For the latter, are the data de-identified, anonymized, or anonymous?
Study responses on paper will be stored in a locked filing cabinet in a researchers locked office. Study responses transcribed to electronic data will be stored on password protected computers. (i) The data collected will not be linked to any other data set. (ii) The data collected will not be sent outside of the institution where it was collected.

3. Indicate how long the data will be securely stored and the method to be used for final disposition of the data.
Paper Records
    Confidential shredding after 7 year(s).
Paper records, other than receipts, will be confidentially shredded 1 year after study. Electronic Data, anonymized data set to be retained indefinitely on the researcher's password protected computer.
Location: Hard disk of researcher(s) password protected computers

4. Are there conditions under which anonymity of participants or confidentiality of data cannot be guaranteed?
No

H. PARTIAL DISCLOSURE AND DECEPTION

1. Will this study involve the use of partial disclosure or deception? Partial disclosure involves withholding or omitting information about the specific purpose or objectives of the research study or other aspects of the research. Deception occurs when an investigator gives false information or intentionally misleads participants about one or more aspects of the research study. Yes

If Yes,
(i) explicitly state if it is partial disclosure and/or deception,
(ii) if applicable, describe the partial disclosure, that is what information is being withheld or omitted concerning the purpose/objectives or procedures,
(iii) if applicable, describe all of the deception(s) to be used in this study, AND
(iv) provide a justification for each use of partial disclosure and deception.

(i) Partial disclosure.
(ii)
a. No mention is made in the recruitment or information and consent documents that participants will be offered monetary rewards to listen to a sound at various points during the study. This is such that the initial rating of sounds is as unbiased as possible.
b. The math exercises, reading exercises and some opinion questions serve as distractions. This is to allow for sufficient time to elapse in accordance with cognitive dissonance theory.

(iv)
a. The reason is that knowledge of the monetary reward may influence the participant's decision making process and consequently affect their attitudes towards the behaviour. Recording attitude and analyzing changes in attitude is a key part of the study; not disclosing this allows us to have a more accurate measure of attitudes.
b. Sufficient time is required between offers and resulting behaviours for each participant's cognitions to be formed independently.

If Yes, **outline the process to be used to debrief participants**.
Provide a hard copy of the written debriefing sheet for participants, the researcher's verbal debriefing script (if applicable), and for deception, the materials used to obtain consent following debriefing

Full Disclosure is provided at the end of the study. A copy of the disclosure is found in Appendix E.

Researchers must ensure that all supporting materials/documentation for their applications are submitted with the signed, hard copies of the ORE form 101/101A. Note, materials shown below in bold are normally required as part of the ORE application package. The inclusion of other materials depends on the specific type of projects.

Protocol Involves a Drug, Medical Device, Biologic, or Natural Health Product
If the study procedures include administering or using a drug, medical device, biologic, or natural health product that has been or has not been approved for marketing in Canada then the researcher is to complete Appendix A. Appendix A is to be attached to each of the one copy of the application that are submitted to the ORE. Information concerning studies involving a drug, biologic, natural health product, or medical devices can be found on the ORE website.

Please **check** below all appendices that are attached as part of your application package:

- Recruitment Materials: A copy of any poster(s), flyer(s), advertisement(s), letter(s), telephone or other verbal script(s) used to recruit/gain access to participants.

- Information Letter and Consent Form(s)*. Used in studies involving interaction with participants (e.g. interviews, testing, etc.)
- Data Collection Materials: A copy of all survey(s), questionnaire(s), interview questions, interview themes/sample questions for open-ended interviews, focus group questions, or any standardized tests.
- Feedback letter *
- Debriefing Letter: Required for all studies involving deception.

* Refer to sample letters.

NOTE: The submission of incomplete application packages will increase the duration of the ethics review process.

To avoid common errors/omissions, and to minimize the potential for required revisions, applicants should ensure that their application and attachments are consistent with the *Checklist For Ethics Review of Human Research Application*

Please note the submission of incomplete packages may result in delays in receiving full ethics clearance. We suggest reviewing your application with the Checklist For Ethics Review of Human Research Applications to minimize any required revisions and avoid common errors/omissions.

INVESTIGATORS' AGREEMENT

I have read the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans, 2nd Edition (TCPS2) and agree to comply with the principles and articles outlined in the TCPS2. In the case of student research, as Faculty Supervisor, my signature indicates that I have read and approved this application and the thesis proposal, deem the project to be valid and worthwhile, and agree to provide the necessary supervision of the student.

NEW As of May 1, 2013, all UW faculty and staff listed as investigators must complete the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans Tutorial, 2nd Ed. (TCPS2) prior to submitting an ethics application. Each investigator is to indicate they have completed the TCPS2 tutorial. If there are more than two investigators, please attach a page with the names of each additional investigator along with their TCPS2 tutorial completion information.

______________________________	____________________
Print and Signature of Principal Investigator/Supervisor	**Date**

Completed TCPS2 tutorial:
___YES ___NO ___ In progress

______________________________	____________________
Print and Signature of Principal Investigator/Supervisor	**Date**

Completed TCPS2 tutorial:
___YES ___NO ___ In progress

Each student investigator is to indicate if they have completed the Tri-Council Policy Statement, 2nd Edition Tutorial (http://pre.ethics.gc.ca/eng/education/tutorial-didacticiel/). If there are more than two student investigators, please attach a page with the names of each additional student investigator along with their TCPS2 tutorial completion information.

Signature of Student Investigator **Date**

Completed TCPS2 tutorial:
___YES ___NO ___ In progress

Signature of Student Investigator **Date**

Completed TCPS2 tutorial:
___YES ___NO ___ In progress

FOR OFFICE OF RESEARCH ETHICS USE ONLY:

Maureen Nummelin, PhD **Date**
Chief Ethics Officer
OR
Julie Joza, MPH
Senior Manager, Research Ethics
OR
Sacha Geer, PhD
Manager, Research Ethics

ORE 101
Revised August 2003

References

[1] I. Ajzen. *Attitudes, Personality, and Behavior*. Open University Press, Stony Stratford, Milton Keynes, UK, 1988.

[2] I. Ajzen. *Attitudes, Personality, and Behavior*. Open University Press, Maidenhead, Berkshire, England, second edition, 2005.

[3] B.D.O. Anderson and J.B. Moore. *Optimal Filtering*. Prentice Hall, Englewood Cliffs, NJ, 1979.

[4] R.F. Baumeister and K.D. Vohs. *Encyclopedia of Social Psychology*, volume 1. Sage Publications, Thousand Oaks, CA, 2007.

[5] R.F. Baumeister and K.D. Vohs. *Encyclopedia of Social Psychology*, volume 2. Sage Publications, Thousand Oaks, CA, 2007.

[6] N.B. Cottrell, D.L. Wack, G.J. Sekerak, and R.H. Rittle. Social facilitation of dominant responses by the presence of an audience and the mere presence of others. *Journal of Personality and Social Psychology*, 9(3):245, 1968.

[7] D.E. Davison. ECE 481: Digital control systems. University of Waterloo, 2014.

[8] D.E. Davison, R. Vanderwater, and K. Zhou. A control-theory reward-based approach to behavior modification in the presence of social-norm pressure and conformity pressure. In *American Control Conference (ACC), 2012*, pages 4076–4052, June 2012.

[9] E.L. Deci. Effects of externally mediated rewards on intrinsic motivation. *Journal of Personality and Social Psychology*, 18(1):105–115, 1971.

[10] L. Festinger. A theory of social comparison processes. *Human Relations*, 7(2):117–140, 1954.

[11] L. Festinger. *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, CA, 1957.

[12] G.C. Goodwin and R.H. Middleton. The class of all stable unbiased state estimators. *Systems & Control Letters*, 13:161–163, 1989.

[13] G.C. Goodwin and R.L. Payne. *Dynamic System Identification: Experiment Design and Data Analysis*. Academic Press, New York, NY, 1977.

[14] M.S. Grewal and A.P. Andrews. *Kalman Filtering: Theory and Practice Using MATLAB*. John Wiley & Sons, Hoboken, NJ, third edition, 2008.

[15] T. Henchy and D.C. Glass. Evaluation apprehension and the social facilitation of dominant and subordinate responses. *Journal of Personality and Social Psychology*, 10(4):446, 1968.

[16] I.L. Janis. Groupthink. *Psychology Today*, 5(6):43–46, 1971.

[17] R. Johansson. *System Modeling and Identification*. Prentice Hall, Englewood Cliffs, NJ, 1993.

[18] R.E. Kalman. A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 82(1):35–45, 1960.

[19] D.A. Kravitz and B. Martin. Ringelmann rediscovered: The original article. *Journal of Personality and Social Psychology*, 50(5):936–941, May 1986.

[20] L. Ljung. *System Identification: Theory for the User*. Prentice Hall, Upper Saddle River, NJ, second edition, 1999.

[21] L. Ljung and T. Glad. *Modeling of Dynamic Systems*. Prentice Hall, Englewood Cliffs, NJ, 1994.

[22] S. Moscovici and M. Zavalloni. The group as a polarizer of attitudes. *Journal of Personality and Social Psychology*, 12(2):125–135, 1969.

[23] D. Myers and S. Spencer. *Social Psychology*. McGraw Hill Ryerson, 3rd Canadian edition, 2006.

[24] J. Ni, D. Kulic, and D.E. Davison. A model-based feedback-control approach to behavior modification through reward-induced attitude change. In *American Control Conference (ACC), 2013*, pages 1956–1963, Washington, DC, June 2013.

[25] H.V. Poor. *An Introduction to Signal Detection and Estimation*. Springer, New York, second edition, 1994.

[26] W.T. Powers. *Behavior: The Control of Perception*. Aldine Publishing Company, Chicago, IL, 1973.

[27] D.E. Rivera, M.D. Pew, and L.M. Collins. Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction. *Drug and Alcohol Dependence*, 88(Suppl2):S31–S40, May 2007.

[28] J.A.F. Stoner. A comparison of individual and group decisions involving risk. Unpublished masters thesis, Massachusetts Institute of Technology, 1961.

[29] R. Vanderwater and D.E. Davison. A dynamic control approach to studying the effectiveness of rewards in inducing behavior and attitude change. In *Control and Automation, 2009. ICCA 2009. IEEE International Conference on Control and Automation*, pages 1062–1067, Dec 2009.

[30] R. Vanderwater and D.E. Davison. Using rewards to change a person's behavior: A double-integrator output-feedback dynamic control approach. In *American Control Conference*, San Francisco, 2011.

[31] W. Wilson. ECE 682: System identification. University of Waterloo. course notes available at https://ece.uwaterloo.ca/~ece683/Lecture_Slides/index.html, 2008.

[32] R.B. Zajonc. Social facilitation. *Science (New York, NY)*, 149(3681):269–274, 1965.

[33] P. Zarchan and H. Musoff. *Fundamentals of Kalman Filtering: A Practical Approach*, volume 232 of *Progress in Astronautics and Aeronautics*. American Institute of Aeronautics and Astronautics, Reston, VA, third edition, 2009.