Optimal bounds, bounded optimality

Böhm, Udo

Chapter 1

Introduction

Decision-making pervades everyday life. Our daily decisions range from complex, life-changing choices such as which study programme to follow at university, to subtle decisions about the meaning of sensory information, for example when trying to understand which platform our train will depart from. Many of these decisions have two features in common. Firstly, they entail uncertainty. Job markets can be volatile and expertise that is sought after now might not be needed anymore in five years time; train stations are notoriously noisy and the loudspeaker always seems to be too far away to be clearly understandable. Secondly, many decisions entail a degree of impatience. Choosing a study programme can be difficult but missing the enrolment deadline is costly. So instead of missing the deadline while trying to find the perfect programme, at some point one might simply have to go with a programme that is sufficiently interesting. Similarly, standing around trying to understand the announcement might mean missing one's train, so at some point a better strategy might be to simply walk to the platform the train usually departs from. Whilst the first of these two factors, uncertainty, has long been acknowledged as an important force shaping human decision-making, the second factor, impatience, seems to have been largely ignored. However, a number of recent publications in neuroscience have suggested that impatience might be a major force that shapes human decision-making at the most fundamental level, namely when interpreting sensory information. In the present work, I will investigate this exciting suggestion theoretically and empirically. In the process, I will develop a number of methodological tools that will not only support this particular line of research but will hopefully also help research efforts in other areas of cognitive science.

Decision-making under uncertainty is traditionally an area of great interest in psychology and efforts to develop testable, quantitative theories span several decades (e.g. Bogacz, Brown, Moehlis, Holmes & Cohen, 2006; Busemeyer & Townsend, 1993; Edwards, 1954; Festinger, 1943; Gigerenzer & Todd, 1999; Kahneman & Tversky, 1979; Laming, 1968; Ratcliff, 1978; Smith, 1995; Tanner & Swets, 1954). Research into human decision-making has largely been carried out along one of two traditions. Economic decision-making, on the one hand, is concerned with the question how the reinforcement history (i.e., rewards and punishments) shapes human decision-making. Perceptual decision-making, on the other hand, is concerned with the question how decision-makers choose between alternative interpretations of sensory information. It might seem intuitive that both factors, perceptual mechanisms as well as economic pressures, play a role in most real decisions. However, perceptual and economic decision-making have for many years enjoyed a surprising degree of separation, with each research tradition developing its own experimental paradigms and quantitative theories (Summerfield & Tsetsos, 2012).

Consider, for example, the random dot motion task, a typical experimental paradigm in the tradition of perceptual decision-making (Britten, Shadlen, Newsome & Movshon, 1992). In this task participants are shown a cloud of pseudo-randomly moving dots on a computer screen. A certain proportion of these dots moves coherently in one direction whilst the remaining dots are randomly displaced from moment to moment, which creates the impression that the cloud is drifting in one direction. The direction of this drift is typically restricted to be either to the left or to the right and participants' task is to press one of two response buttons to indicate as quickly as possible in which direction the cloud is drifting.

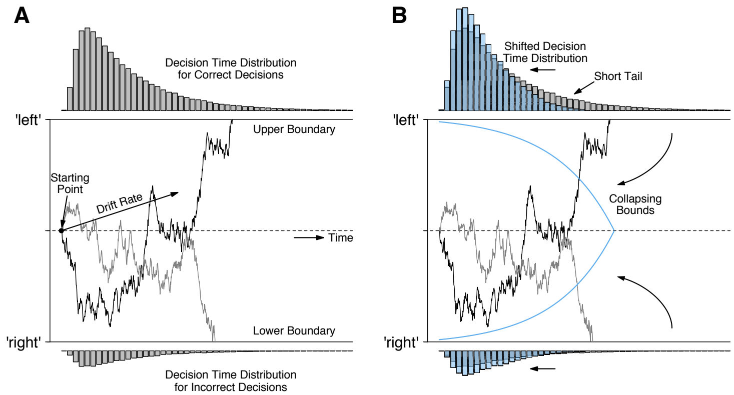

Figure 1.1: Sequential Sampling Model. Panel A shows the basic components of sequential sampling model with a continuous time and evidence scale. Panel B illustrates the effect of collapsing bounds on the predicted decision time distribution (in blue).

A prominent class of mathematical models in perceptual decision-making are sequential sampling models (Vickers, 1979). These models come in various forms but share a few core components that are illustrated in panel A of Figure 1.1. Sequential sampling models assume that decision makers arrive at a decision by integrating noisy information over time until the accumulated information exceeds one of two decision bounds. In the case of the random dot motion task, the participant has two response options, one for a 'leftward' drift and one for a 'rightward' drift. Each of these response options is associated with a decision bound. In the figure, the upper black line represents the bound for a 'leftward' decision and the lower black line represents the bound for a 'rightward' decision. The jagged black line represents the integration of information as the participant observes the cloud stimulus. At any given time a few dots move coherently in one direction. However, because the remaining dots move randomly, the participant will sometimes perceive the movement to be more coherent and sometimes less coherent, and might at times even perceive a movement in the opposite direction. This leads to the up-and-down movement of the black line. The black dot labelled 'starting point' indicates the starting point of the integration process. This point is typically located at the midpoint between the two boundaries. A starting point that lies above the midpoint results in a bias towards the 'leftward' decision because the decision maker would need to integrate less information to reach the upper boundary than to reach the lower boundary. The black arrow labelled 'drift rate' indicates the average rate at which the black line climbs towards the upper boundary. The point at which the jagged black line touches the upper decision bound is the time when the participant decides that the cloud is drifting to the left. If this decision process is repeated several times, the black line will hit the upper boundary at different times, because the integrated information is noisy. This gives rise to the distribution of decision times shown on top of the upper boundary. Moreover, on some trials the integrated information might hit the lower boundary first, as illustrated by the jagged grey line, which leads the participant to incorrectly conclude that the cloud is drifting to the right. The resulting distribution of decision times for incorrect decisions is shown below the lower boundary.

One interesting aspect of the sequential sampling model just described is that it assumes constant decision boundaries. This assumption, which is shared by most standard models (e.g., Laming, 1968; Ratcliff, 1978; Ratcliff & Smith, 2004; Smith & Vickers, 1988), implies that decision makers should continue accumulating information until a decision boundary is hit, irrespective of how much time they have already spent on the decision. This results in the long right tail of the decision time distributions in panel A of Figure 1.1. However, it seems intuitively clear that in any realistic context decision makers will not be willing or able to devote vast resources to a single decision; participants in experimental studies typically leave the laboratory after one or two hours.

Recently, this criticism has been expressed more formally in terms of the economic argument of reward rate optimality. Assuming that decision makers are motivated to maximise their reward rate, that is, the average number of rewards per unit time, researchers have argued that decision makers should become increasingly impatient as they spend more time on a decision (Shadlen & Kiani, 2013; Cisek, Puskas & El-Murr, 2009; Thura, Beauregard-Racine, Fradet & Cisek, 2012; Hanks, Mazurek, Kiani, Hopp & Shadlen, 2011). This growing impatience should lead decision makers to forego overly long decision times by picking whichever option is currently favoured by the accumulated information. One way to implement impatience in sequential sampling models is by lowering the decision boundaries over time, which is illustrated in panel B of Figure 1.1. The blue collapsing boundaries intersect the jagged black and grey lines earlier than either of them reaches the constant black decision boundaries, thus short-cutting the decision process. This results in a shift of the blue decision time distribution and markedly reduces the right tail compared to the grey decision time distribution predicted by the standard model with constant boundaries.

Despite the intuitive appeal of the impatience hypothesis, there are several theoretical, empirical, and methodological aspects that remain unaddressed. Firstly, for the last 40 years sequential sampling models with constant decision boundaries have successfully accounted for behavioural and neurophysiological data from a wide range of experimental paradigms (e.g., Ratcliff & McKoon, 2008; Smith, Ratcliff & Wolfgang, 2004; Forstmann, Ratcliff & Wagenmakers, 2016). If impatience is a ubiquitous influence on human decision-making, the question arises why earlier studies did not report systematic discrepancies between models and data. I argue here that in most decision environments constant decision boundaries yield near-maximal reward rates. Therefore, empirical tests of the impatience hypothesis need to be based on thorough quantitative analyses to identify the stochastic and economic structure of the decision environment that ought to result in a detectable degree of impatience.

Secondly, testing the impatience hypothesis requires assessing the shape of human decision maker's decision boundaries. Although these boundaries are not directly observable, different experimental methods can be employed to obtain proxy measurements of the decision boundaries. In particular, I advocate the use of expanded judgment tasks (Irwin, Smith & Mayfield, 1956; Vickers, Smith, Burt & Brown, 1985), which allow researchers to record the stimulus information presented to decision makers and thus infer the amount of information decision makers have observed at the time of decision commitment. Moreover, I argue that recordings of the Contingent Negative Variation (Walter, 1964), an EEG potential, can be used as a physiological marker of the decision boundaries.

Finally, quantitatively testing competing sequential sampling models requires fitting the models to behavioural data and comparing their fit. As I discuss here, hierarchical Bayesian methods are currently the best available tools for both tasks. In fitting models to data, hierarchical Bayesian methods allow parameter estimates for individual participants and parameter estimates for the group of participants as a whole to mutually inform each other. This optimal use of all available information results in the smallest estimation error for individual participants (Efron & Morris, 1977). In comparing the relative fit of competing models, the Bayesian standard approach – the Bayes factor – not only allows researchers to select the best-fitting model but also to quantify the relative support the data lend to each model.

In the remainder of this section I will give an overview of the problems that will be addressed in each chapter of this thesis.

1.1 Chapter Outline

Chapters 2 to 4 focus on a theoretical and experimental assessment of the impatience hypothesis. Chapter 2 provides a review of the literature on the impatience hypothesis. The first part of the chapter summarises the theoretical underpinnings of the impatience hypothesis and gives an overview of the different implementations of impatience in computational models of decision-making. The second part of the chapter discusses the empirical support for the impatience hypothesis from behavioural and neurophysiological studies. According to some accounts, the empirical support for the impatience hypothesis is so overwhelming that collapsing decision boundaries should replace the current default assumption of constant decision boundaries (e.g., Shadlen & Kiani, 2013). However, our review of the existing literature reveals a number of methodological shortcomings in studies supporting the impatience hypothesis. Moreover, systematic discrepancies between studies with monkeys, which are often cited in support of the impatience hypothesis, and traditional studies with human participants impede generalisations across species. Based on our review, we suggest that a combination of rigorous quantitative model comparisons, suitable experimental tasks and EEG recordings are needed to obtain more decisive evidence in the debate.

Chapter 3 explores the quantitative predictions of the impatience hypothesis and presents an experimental test of these predictions. An often-made assertion in the debate about the impatience hypothesis is that, in order to maximise reward rates in dynamic environments, decision makers need to rely on decision boundaries that change over the course of the decision process (Cisek et al., 2009; Shadlen & Kiani, 2013; Thura et al., 2012). However, in light of the considerable success of sequential sampling models with constant decision boundaries, the question arises whether and under what circumstances dynamic decision boundaries yield substantially higher reward rates than constant boundaries. To address this question, we use dynamic programming and simulation methods to quantify the reward rates obtained by constant and dynamic decision boundaries in different decision environments. Our results suggest that, in most situations, constant boundaries yield near-maximal reward rates. Based on these results we conducted an experiment in which we tested whether decision makers adjust their decision boundaries to maximise reward rates. We exposed decision makers to different decision environments that should reliably induce different shapes of the optimal decision boundaries. This experiment yielded mixed results. Whilst participants were sensitive to the environmental dynamics, there were large individual differences in the degree to which participants' decision boundaries approximated the reward rate optimal boundaries. In complex dynamic environments in particular, participants deviated considerably from reward rate optimality, even after extensive practice. These results draw further doubt on claims that human decision makers rely on a dynamic decision criterion by default.

Chapter 4 presents a neurophysiological method for measuring the setting of decision makers' boundaries before the onset of the decision process. A complication in experimental tests of the impatience hypothesis is that decision makers' decision boundaries are generally not directly observable. However, neurophysiological recordings can provide a measure of the activity of brain areas that are responsible for setting decision makers' boundaries before the onset of a decision task. A current theoretical framework from cognitive neuroscience suggests that the basal ganglia control the trade-off between fast and accurate decision-making, that is, the setting of the decision boundaries, by modulating the excitability of cortical areas (Forstmann et al., 2008, 2010). We propose that the Contingent Negative Variation (CNV; Walter, 1964), a slow cortical EEG potential, reflects fluctuations in cortical excitability, and thus the setting of the decision boundary. We tested this hypothesis in an EEG experiment in which we instructed participants to either respond quickly or accurately. Our results show that trial-by-trial fluctuations in participants decision boundary correlate with single-trial CNV amplitude under conditions of speed but not accuracy stress. This leads us to conclude that the CNV might serve as a measure of short-term adjustments of the decision boundaries.

Chapters 5 to 7 discuss statistical and methodological problems related to the assessment of the impatience hypothesis that are also of interest in cognitive modelling in general. In Chapter 5 we develop a Bayesian regression framework for relating the parameters of cognitive models to covariates. Testing the impatience hypothesis often requires relating covariates that are thought to reflect participants' decision boundary to the boundary parameters in sequential sampling models. Similarly, testing other types of cognitive models often involves the evaluation of hypotheses about relationships between covariates and model parameters. However, many models do not come equipped with the statistical framework needed to relate model parameters to covariates. Instead, researchers often revert to classifying participants into groups depending on their values on the covariates, and subsequently comparing the estimated model parameters between these groups. This classification-based approach can severely bias statistical inference. We develop a Bayesian regression framework for hierarchical cognitive models that allows researchers to compute Bayes factors for relationships between covariates and model parameters. Using a simulation study, we demonstrate how our regression framework overcomes the statistical biases associated with the classification-based approach.

In Chapter 6 we present a comprehensive comparison of fitting methods for the Drift Diffusion Model (DDM, Ratcliff, 1978), one of the most popular sequential sampling models. The DDM describes decision-making in terms of seven model parameters, four main parameters that account for the general shape of participants' response time distributions and three between-trial variability parameters that allow the model to capture more subtle aspects of response time distributions. Several researchers have reported difficulties estimating the between-trial parameters yet reliable parameter estimates are a prerequisite for evaluating hypotheses about sequential sampling models. This situation is further complicated by the availability of numerous estimation methods for the DDM. To assess how reliably the between-trial parameters can be estimated, we invited experts from the DDM community to apply their various fitting methods to simulated data and provide guidance on estimating the DDM's between-trial parameters. Our results show that some between-trial parameters can be estimated more reliably than others across fitting methods. Nevertheless, estimation performance can be improved by putting a priori constraints on these parameters and by pooling data across participants, both of which is naturally achieved by hierarchical Bayesian methods.

Finally, Chapter 7 discusses a number of popular shortcut analysis strategies in cognitive modelling that can lead to biased conclusions. Cognitive models are often applied to experimental data that are hierarchically structured. However, two popular modelling strategies do not properly accommodate this hierarchical structure. We review some established theoretical results from statistics that suggest that these shortcut strategies can result in biased conclusions. To gauge the severity of these biases we conducted a simulation study for a two-group experiment. Our results show that one shortcut strategy biases statistical tests towards the null hypothesis whilst the other strategy results in a bias towards the alternative hypothesis. We conclude that only hierarchical models of the multilevel data guarantee correct conclusions.