

University of Groningen

Optimal bounds, bounded optimality

Böhm, Udo

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:

2018

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Böhm, U. (2018). Optimal bounds, bounded optimality: Models of impatience in decision-making. [Groningen]: University of Groningen.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Discussion and Conclusions

The present dissertation has investigated recent claims that impatience is a ubiquitous force in human perceptual decision-making. These claims emanate from the assumption that human decision makers are motivated to maximise their reward rates. On the theoretical side of this argument, we have assessed the current evidential support for the impatience hypothesis and have addressed the question why previous empirical studies have not reported systematic discrepancies between the standard model and data. On the empirical side, we have used an expanded judgment task to test how sensitive human decision makers are to the dynamics of their decision environment. Moreover, we have developed a physiological measure for decision makers' boundary setting. On the methodological side, we have developed a regression framework for relating covariates to the parameters of cognitive models, we have assessed different methods for estimating the DDM's between-trial parameters, and we have argued that shortcut methods for fitting cognitive models to hierarchical data can lead to biased conclusions. Despite the apparent diversity of these three methodological points, they are closely linked to the current debate about impatience in perceptual decision-making: a major reason for the many misunderstandings and incorrect assertions in this debate is the lack of rigorous statistical analyses. In what follows I summarise the results of the present work before I turn to a discussion of the broader context and derive a number of recommendations for future research.

8.1 Summary and Conclusions

Chapter 2 reviewed the current empirical support for the impatience hypothesis. The core assumption underlying the impatience hypothesis is that decision makers strive to maximise their reward rates. However, many of the studies that purport to support the impatience hypothesis did neither control nor systematically manipulate reward rates. Moreover, these studies differed systematically from studies that support the traditional standard model with constant decision boundaries in the way stimuli were presented, the duration of practice, and even the species of the participants. Physiological arguments for the impatience hypothesis were sim-

ilarly incomplete. Some of the often-cited single-cell recording studies in monkeys showed good agreement between physiological data and models with constant decision boundaries. Moreover, comparisons of physiological recordings in humans and monkeys, and thus generalisations across species, are complicated by large differences in spatial and temporal scales on which data were recorded. Finally, both behavioural and physiological studies often failed to carry out quantitative comparisons between the standard model with constant boundaries and models with an impatience component. These results led us to conclude that, despite claims to the contrary (Shadlen & Kiani, 2013), the impatience hypothesis is not unequivocally supported by the data. To provide conclusive evidence in the debate about the impatience hypothesis, future studies should systematically manipulate rewards rates within a single experimental setup and carry out rigorous quantitative model comparisons. Moreover, comparative studies should use the same training and testing protocol across species and obtain physiological records at comparable temporal and spatial resolutions.

Chapter 3 presented a theoretical analysis and empirical test of the impatience hypothesis. In our theoretical analysis, we considered two dynamic decision environments, one in which sampling costs increased over time and one in which sampling costs decreased over time. Our results show that in both environments, unless task difficulty was very high, optimal dynamic boundaries and constant boundaries yielded similar reward rates. This suggests that assertions that dynamic decision environments automatically induce collapsing boundaries might be incorrect. Based on these results, we conducted an experimental study in which we tested whether human decision makers adjust their decision boundaries to obtain maximal reward rates. This experiment yielded mixed results. Although participants were sensitive to differences in the dynamics of the decision environment, there were large individual differences in the degree to which participants' decision boundaries approximated the reward rate-optimal boundaries. In dynamic environments in particular, some participants deviated considerably from reward rate optimality, even after extensive practice. These results suggest that reward rate maximisation might not be the main force shaping human decision-making.

Chapter 4 developed a neurophysiological measure of decision makers' boundary setting before the onset of the decision process. Such physiological markers are particularly relevant to tests of the impatience hypothesis as decision makers' decision boundaries are generally not directly observable. We conducted an EEG experiment in which participants performed a random dot motion task either under speed or under accuracy instructions. Estimates of participants' boundary setting were obtained using a linear deterministic accumulator model, a close kin of sequential sampling models. The results of this study showed that the height of participants' decision boundaries correlates with between-trial fluctuations in CNV amplitude under speed but not under accuracy instructions. This led us to conclude that CNV amplitude might reflect short-term adjustments of the decision boundaries that serve to enhance fast decision-making. Consequently, measurements of CNV amplitude might provide a viable method for measuring decision boundaries online during task performance.

Chapter 5 considered the problem of relating covariates to parameter estimates in cognitive models. Based on previous work in statistics, we suggested a

Bayesian regression framework that allows researchers to simultaneously fit a cognitive model to behavioural data and relate the model parameters to covariates. To illustrate the superiority of our method to popular categorisation-based analysis approaches, we fitted a reinforcement-learning model to simulated data using our regression method and a categorisation-based approach. These simulations showed that, in the case of multiple uncorrelated covariates, a categorisation-based approach can miss veridical relationships between model parameters and covariates, whilst in the case of multiple correlated covariates, a categorisation-based approach can suggest spurious relationships between model parameters and covariates. Our regression framework avoids these statistical fallacies and allows for the computation of Bayes factors as a measure of the evidential support for relationships between model parameters and covariates.

Chapter 6 compared different methods for estimating the DDM's between-trial variability parameters. We invited experts from the DDM community to apply their estimation methods to three simulated data sets. The results of this study showed that all estimation methods could reliably recover between-trial variability in non-decision time. Estimates of variability in starting point and drift rate, on the other hand, were associated with considerable uncertainty across estimation methods and tended to miss the true parameter value by a wide margin. However, uncertainty for estimates of drift rate variability was markedly lower for methods that pooled data across participants. Moreover, imposing a priori constraints on variability in drift rate and starting point resulted in estimates that were closer to the true value. Both of these measures are naturally implemented by hierarchical Bayesian estimation methods. These results led us to conclude that, firstly, researchers should give careful consideration to whether between-trial parameters are needed to fit a particular data set. Secondly, if between-trial parameters are to be estimated, researchers should use hierarchical Bayesian estimation methods.

Chapter 7 scrutinised a number of popular shortcut strategies in cognitive modelling. Cognitive models are often applied to hierarchical experimental data. However, two popular modelling approaches do not properly accommodate this hierarchical structure, which can lead to biased conclusions. To gauge the severity of these biases we conducted a simulation study for a two-group experiment. One popular shortcut strategy ignores the hierarchical data structure. In line with theoretical results from statistics, our simulations showed that this strategy biases Bayesian and frequentist tests towards the null hypothesis. Another popular shortcut strategy takes a two-step approach by first obtaining participant-level estimates from a hierarchical cognitive model and subsequently using these estimates in a follow-up statistical test. In line with theoretical results, our simulations showed that this strategy leads to a bias towards the alternative hypothesis. The conclusion from this study was that only hierarchical models of the multilevel data guarantee correct conclusions.

8.2 Discussion and Future Directions

The goal of this present work was to investigate the impatience hypothesis. The results of the theoretical and empirical analyses presented in Chapters 2 and 3 suggested that impatience hypothesis does not hold in the strong form it was originally stated in. Exposing decision makers to a dynamic decision environment does not automatically induce collapsing boundaries. The present work shares this finding with a few recent publications that tested human decision makers in dynamic environments and found no clear evidence for collapsing boundaries (Hawkins, Forstmann et al., 2015; Voskuilen et al., 2016). Moreover, in situations where reward rate maximisation requires expanding instead of collapsing boundaries, decision makers seem to invariable fail to maximise their reward rates, as evidenced by the findings in Chapter 3 and Experiment 3 in Malhotra et al. (2017).

These results raise the question whether the impatience hypothesis is incorrect and whether human decision-makers ignore reward rates. I would argue that it is not the impatience hypothesis that is incorrect but rather the way it is phrased. Decision makers certainly become impatient, participants in experimental studies typically press a button in no more than a few seconds. However, human decision makers never have full knowledge of the probabilistic structure of the environment but rather need to infer parameters such as the distribution of task difficulties or sampling costs from repeated interactions with the environment. Moreover, the neurophysiological processes that implement human decision-making are noisy, adding to the uncertainty about the environment (Deneve, 2012). Consequently, decision makers are unlikely to find the exact shape of the reward rate optimal boundaries, especially when this shape is complex. Instead, decision makers will need to settle for simple boundaries that require only minor adjustments to yield high reward rates in most decision environments; constant boundaries are a good candidate for such a simple shape.

A more adequate statement of the impatience hypothesis seems to be ‘human decision-makers maximise reward rates within the limits of their cognitive apparatus’. Progress in testing this version of the impatience hypothesis will hinge critically on the development of quantitative models that account explicitly for the limited and uncertain knowledge available to human decision makers. A candidate class of models for this task are reinforcement learning models (Busemeyer & Stout, 2002; R. S. Sutton & Barton, 1998), which have been used successfully to explain the acquisition of decision policies in value-based decision-making (Ahn et al., 2008; Fridberg et al., 2010; Steingroever et al., 2014). In fact, Khodadadi, Fakhariand and Busemeyer (2014) have proposed a model that combines reinforcement learning mechanisms with a sequential sampling model to explain the acquisition of reward rate optimal constant decision boundaries. An extension of this model to dynamic decision environments might allow researchers to quantify the degree of reward rate optimality that decision makers can realistically achieve.

A further important implication for future work is that researchers need to explore the full range of quantitative predictions generated by their theories. An intuitively appealing argument for the impatience hypothesis is that, in order to maximise the ratio of average rewards to average decision time, decision makers should cut short decisions that take long (i.e., task difficulty is high) or yield

smaller rewards as time passes (i.e., mounting sampling costs). To achieve this, decision makers should adopt collapsing boundaries. However, the quantitative analysis presented in Chapter 2 showed that this intuition only applies to a surprisingly limited set of decision environments; despite mounting sampling costs, constant boundaries lead to near-optimal reward rates in most decision environments. Similarly, the quantitative analysis in Malhotra et al. (in press) shows that constant boundaries are nearly optimal across a wide range of mixed task difficulties. Another instance where intuitions about the predictions of sequential sampling models have led to incorrect conclusions is discussed in Evans et al. (2017). These examples demonstrate that the complexity of contemporary quantitative models in psychology exceeds the range of what can be grasped through simple intuition. A main motivation for the development of quantitative models is that they can generate precise and testable predictions. However, to be able to take advantage of these precise predictions, researchers need to understand what exactly a model does or does not predict. In light of the increasing popularity of complex, high-dimensional models of decision-making, such as reinforcement learning models (Busemeyer & Stout, 2002; R. S. Sutton & Barton, 1998) and neural networks (Huang & Rao, 2013; Rao, 2010; Standage et al., 2011), methods for systematic exploration of a model's parameter space, such as parameter space partitioning (Pitt, Kim, Navarro & Myung, 2006; Pitt, Myung, Montenegro & Pooley, 2008), will become increasingly important.

The analyses in Chapter 4 used the single-trial LBA model (Van Maanen et al., 2011; Ho et al., 2012) to obtain estimates of participants' drift rate and boundary setting on a single-trial basis. One drawback of this model is that it relies on a two-step estimation procedure; in a first step the LBA model is fitted to the participant-level data and in a second step, these maximum-likelihood estimates are used to compute single-trial maximum-likelihood estimates of drift rate and boundary. The reason for this procedure is that the likelihood function is not identified on the single-trial level. However, the resulting single-trial estimates are conditional on the participant-level estimates, and might therefore be biased. Recently developed hierarchical Bayesian methods avoid this problem by either directly modelling the joint distribution of neurophysiological measurements and model parameters on the single-trial level (Turner, Van Maanen & Forstmann, 2015), or by defining a regression model that relates single-trial neurophysiological measurements to model parameters (Wiecki et al., 2013; Frank et al., 2015). Consequently, future applications of sequential sampling models to single-trial data should forego two-step estimation and instead rely on a statistically optimal joint modelling approach.

Throughout this thesis, I have advocated the use of hierarchical Bayesian models in combination with Bayes factors to test scientific hypotheses. Most applications of these methods presented here were limited to relatively simple hypotheses that often involved nested models. In Chapter 5, for example, we tested whether individual regression weights deviated from zero, and in Chapter 7 we tested whether standardised group differences were zero. In these instances, Bayes factors could be computed analytically or could easily be approximated using MCMC-sampling in combination with Savage-Dickey density ratios (Dickey & Lientz, 1970). However, many research questions involve complex, non-nested

models. Recent applications of models of choice response times, for example, involve the comparison of competing, non-nested stochastic and ballistic accumulator models (Donkin, Brown, Heathcote & Wagenmakers, 2011; Heathcote & Hayes, 2012; Osth, Bora, Dennis & Heathcote, 2017). In these cases, Bayes factors cannot be computed analytically and Savage-Dickey density ratios are not applicable. A possible alternative solution is offered by Monte Carlo sampling methods that approximate the marginal likelihood of the competing models. One particularly attractive method is bridge sampling (Bennett, 1976; Gronau et al., 2017; Meng & Wong, 1996), which allows for the approximation of Bayes factors under relatively mild conditions (e.g., Frühwirth-Schnatter, 2004). Combining this method with Bayesian implementations of sequential sampling models and the regression framework developed in Chapter 5 will yield a powerful tool for testing complex hypotheses about complex models.