

Decomposing the influences of uncertainty on learning

Normative computations, uncertainty biases, and lifespan differences

Dissertation

Zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften
(Dr. rer. nat.)

am Fachbereich Erziehungswissenschaft und Psychologie
der Freien Universität Berlin



Vorgelegt von
Rasmus Bruckner
M.Sc., B.Sc. Psychologie

Berlin, 2020

Erstgutachter: Prof. Dr. Hauke R. Heekeren

Zweitgutachter: Prof. Dr. Gerhard Jocham

Drittgutachter: Prof. Dr. Markus Ullsperger

Datum der Disputation: 27.11.2020

Contents

Danksagung	v
1 Summary	1
2 Glossary	5
3 Introduction	9
3.1 Normative computations	11
3.2 Uncertainty biases	29
3.3 Lifespan differences	34
4 Research questions and hypotheses	39
5 General methodology	43
5.1 Participants	43
5.2 Experimental tasks	44
5.3 Computational modeling	47
5.4 Software and data repositories	53
6 Summary of the dissertation studies	55
6.1 Study I: Perceptual uncertainty	55
6.2 Study II: Uncertainty and aging	56
6.3 Study III: Default beliefs in children and older adults	57
6.4 Study IV: Computational modeling across the lifespan	59
7 General discussion and future directions	61
7.1 Discussion of the research questions	61
7.2 Perceptual uncertainty	72
7.3 Expected uncertainty	74
7.4 Unexpected uncertainty	75
7.5 Resource rationality	76
7.6 Uncertainty in the cycle of adaptive behavior	77
7.7 Conclusion	79
8 Bibliography	81
9 Appendix	95
9.1 Deutsche Zusammenfassung	95

Contents

9.2	List of publications	99
9.3	Talks and colloquium presentations	101
9.4	Eigenanteil	103
9.5	Eidesstattliche Erklärung	105
9.6	Research articles	106

Danksagung

Ich möchte mich bei all denjenigen bedanken, die meine Dissertation ermöglicht, unterstützt und in vielerlei Hinsicht inspiriert haben.

Zuallererst bedanke ich mich sehr herzlich bei meinem Doktorvater *Hauke Heekeren*. Du hast in meinen Augen die seltene Gabe, als Betreuer das große Ganze zu überblicken und an den richtigen Stellen auf die Details zu achten. Du hast mich immer sehr dabei unterstützt, mit unterschiedlichen WissenschaftlerInnen und Arbeitsgruppen zu kooperieren, sodass wir jetzt sogar mit Mäusen forschen. In den vergangenen Jahren bist Du neben all den gemeinsamen wissenschaftlichen Tätigkeiten zu meinem Mentor geworden, der mich hoffentlich auch in Zukunft noch begleiten wird.

Ich bedanke mich auch sehr herzlich bei *Ben Eppinger*. Wir kennen uns schon seit meiner HiWi-Zeit am MPI und haben seitdem an verschiedenen Projekten zusammengearbeitet. Du hast mir schon früh die Möglichkeit gegeben, an Forschungsprojekten mitzuwirken, viel zu lernen und somit die Grundlage für meine Promotion geschaffen. Trotz der großen räumlichen Entfernung haben wir weiterhin im Team mit Matt zusammengearbeitet, ein langjähriges Projekt nun endlich zum Abschluss gebracht und noch viele Projekte in Zukunft geplant.

Dem schon eben erwähnten *Matt Nassar* gilt mein nächster Dank. Dear Matt, I had a great time with you and your family in Providence. Thank you for inspiring me so many times so that I have now completed my own dissertation on learning and uncertainty. Many thanks for teaching me all my modeling basics and the amazing collaboration over the years!

Ebenfalls möchte ich mich vielmals bei *Dirk Ostwald* bedanken. Deine Mathematik hat mich von Anfang an begeistert und mir eine neue Art des wissenschaftlichen Denkens gelehrt. Außerdem werde ich vermutlich immer von Deiner einzigartigen Akribie profitieren.

Danksagung

Ein großer Dank geht an *Ulman Lindenberger*. Du hast mich über die gesamte Zeit der Promotion mit viel Vertrauen unterstützt. Deine wissenschaftliche Begeisterung und Offenheit hat mich immer wieder inspiriert. Außerdem habe ich mich unglaublich über die Einladungen zum Symposium am Tegernsee gefreut und bin jedes Mal mit vielen Ideen nach Berlin zurückgekehrt.

Weiterer Dank gilt der Arbeitsgruppe mit *Felix Molter*, *Julia Rodriguez Buritica*, *Yuan-Wei Yao* und *Adrian Fischer*. Wir haben viele Jahre zusammengearbeitet und jeder von Euch hat einen nicht geringen Teil zu meiner Dissertation beigetragen. Auch sei an dieser Stelle *Daniela Satıcı-Thies* erwähnt. Ohne Deine organisatorische Hilfe wäre ich in bürokratischen Angelegenheiten verloren gewesen. Es ist schön, dass Du unsere Gruppe so gut zusammenhältst. Bei *Timo Schmidt* bedanke ich mich für die tolle Zusammenarbeit bei der Organisation der CCNB-Events und vor allem des 10-jährigen Jubiläums. Außerdem vielen Dank an *Christoph Korn*, dessen Dissertation und Ratschläge mir beim Schreiben viel geholfen haben.

Weiterhin danke ich *Markus Ullsperger*, *Shu-Chen Li* und *Michael Frank*. Working in your labs was a great experience. I learned a lot from you about science and academia, which contributed significantly to my decision to pursue a Ph.D.

Vincent en Erik, jullie hebben mij al in Nijmegen als vrienden begeleid. Wij hebben samen niet alleen veel gestudeerd maar ook ons enthousiasme voor de wetenschap ontwikkeld en uiteindelijk allemaal een promotie gedaan. In Nijmegen en Berlijn hadden we altijd veel plezier en zullen in de toekomst hopelijk nog vaak met elkaar afspreken.

Kai und Marcus, wir haben uns im Psychologie-Master kennengelernt und gemeinsam viel Zeit miteinander und an unseren jeweiligen Laptops in Berliner Bibliotheken, Cafés und später in unseren Büros verbracht. Vielen Dank für Eure Freundschaft.

Bei meinen Eltern, *Rosa und Helmut*, bedanke ich mich für Euer bedingungsloses Vertrauen. Hättet Ihr damals gedacht, dass ich zum Thema Lernen promoviere?

Ohne Eure Unterstützung in jeglicher Hinsicht, das unermüdliche Interesse und ein immer offenes Ohr für meine wissenschaftlichen Ausführungen, hätte ich nicht geschafft, auf was ich jetzt zurückblicken kann. Bei meinem Bruder *Jan* bedanke ich mich ebenfalls. Du hast mich nicht nur immer wieder mit Deinem Informatikwissen inspiriert, sondern ich freue mich vor allem, dass wir uns so gut verstehen und immer eine tolle Zeit miteinander haben.

Zu guter Letzt bedanke ich mich bei *Mareike*. Du hast mich über die gesamte Zeit der Promotion begleitet und bei allen Hürden, die es zu überwinden galt, riesig unterstützt! Außerdem möchte ich Dir vielmals für Deine Ratschläge zur Gestaltung und Verbesserung der Dissertation danken.

1 Summary

Learning often takes place in environments with uncertainty about current and future outcomes. To behave adaptively in these circumstances, people need to learn beliefs from past experiences, based on which they can predict future outcomes.

In my dissertation, I examine:

1. Normative computations that should determine learning under uncertainty.
2. Uncertainty biases that lead to deviations from normative learning.
3. Age-related differences in learning under uncertainty that are characteristic across the lifespan.

Here, the term normative computations from the field of computational neuroscience refers to computations that provide an optimal solution to a learning and decision-making problem. My dissertation studies draw on computational models that implement normative computations and formally define uncertainty. Based on these models, the studies systematically investigated to what degree younger adults and people across the lifespan consider uncertainty when learning from their experiences.

I begin by illustrating that adaptive behavior consists of several related steps, including a representation of the environment, decision making, and learning ([Introduction](#)). Based on this, I present a framework that decomposes uncertainty into three forms: perceptual uncertainty, expected uncertainty, and unexpected uncertainty ([Normative computations](#)). Perceptual uncertainty is related to sensory information processing, expected uncertainty arises from outcome variability, and unexpected uncertainty is the consequence of changes in the environment. For each form, I describe how individuals should learn under uncertainty based on normative computations. I then show that biases, that is, deviations from a normative consideration of uncertainty, are characteristic of human learning behavior ([Uncertainty](#)

1 Summary

biases). Finally, I motivate why capturing these biases in computational models of cognition can improve our understanding of age-related lifespan differences in learning under uncertainty ([Lifespan differences](#)).

The first dissertation study ([Bruckner et al., 2020a](#)) examined which normative computations should guide learning under perceptual uncertainty, to which degree humans regulate learning accordingly, and how past perceptual choices bias this process. The second study ([Nassar et al., 2016](#)) investigated expected and unexpected uncertainty in younger and older adults, particularly how biases in the consideration of uncertainty explain age-related learning differences. The third study ([Bruckner et al., 2020b](#)) built upon this and examined the role of simplified learning strategies across the lifespan. Finally, the fourth study ([Van den Bos et al., 2018](#)) was an opinion paper on how applying computational cognitive models advances our understanding of age-related lifespan differences in learning and decision making.

In the following, I briefly summarize the results of the dissertation studies mentioned above. In [Bruckner et al. \(2020a\)](#), we showed that perceptual uncertainty often corrupts learning because of misinterpreted perceptual information. Learning behavior under perceptual uncertainty should be more cautious than in perceptually clear situations to avoid such misinterpretations. We found that humans consider perceptual uncertainty during learning. However, we also identified learning biases driven by previous perceptual choices, which led to a less cautious regulation of learning.

In [Nassar et al. \(2016\)](#), our results suggested that age-related learning differences are related to the adjustment to expected uncertainty. In particular, we found that older adults (60 to 80 years) exhibit a bias to underestimate uncertainty about their beliefs compared to younger adults (20 to 30 years). This form of uncertainty underestimation leads to less flexible learning behavior compared to younger adults.

In [Bruckner et al. \(2020b\)](#), we found that age-related impairments in learning under uncertainty often arise because children (7 to 11 years) and older adults resort to

simplified learning strategies that lead to more repetitive responding (perseveration) and stronger environmental influences on behavior (environmental control) compared to younger adults.

Finally, in [Van den Bos et al. \(2018\)](#), we argued that computational cognitive models are an essential tool to gain a better understanding of age-related learning and decision-making differences. In particular, we illustrated both promises of the application of computational models to study age-related behavioral differences (concerning risk-taking, strategy selection, and reinforcement learning) and potential pitfalls.

After discussing the implications of these studies ([General discussion and future directions](#)), I propose a cognitive model of learning under uncertainty based on the new insights of my studies and previous work in the literature ([Uncertainty in the cycle of adaptive behavior](#)). In summary, the dissertation highlights that learning is a dynamic process that is influenced by multiple forms of uncertainty. People take uncertainty into account during learning but show inherent uncertainty biases that substantially change across the lifespan.

2 Glossary

agent Something that acts (Russell & Norvig, 2010). Here, mainly referred to as a computationally implemented artificial agent model that perceives its environment and performs actions to achieve a goal. 9, 10, 22, 48, 61

Bayesian inference Mathematical method to update probability distributions based on newly arriving information. In computational neuroscience, Bayesian models are often used to model an agent's beliefs, mathematically define uncertainty, and derive normative computations. See Box 2 for more details. 11, 13, 22, 69

belief In the context of Bayesian inference, the term reflects an agent's knowledge of a variable that is often not directly observable in the environment. Beliefs can be quantified with a belief distribution that assigns a degree of belief to each possible value of the variable. 12, 13, 22

belief state Posterior probability of task states conditional on the current sensory information and computed using Bayesian inference. Often applied to deal with observations when task states are not directly observable, e.g., because of perceptual uncertainty. 12–14, 30, 45, 55

belief uncertainty Variance of an agent's belief distribution, where a wider distribution reflects more belief uncertainty than a narrower distribution. Often emerges from expected uncertainty when beliefs are merely based on a few variable outcomes. Sometimes also referred to as estimation uncertainty. 11–13, 19, 22, 30, 47, 49, 56, 77, 78

bias Systematic deviation from normative behavior. In this dissertation, uncertainty biases are particularly defined as deviations from normative computations. 30, 48, 55, 69, 79

categorical perception Categorical commitment to one interpretation of a stimulus and neglect of alternative interpretations under perceptual uncertainty. Often driven by past perceptual choices. [29](#), [35](#), [55](#), [64](#), [73](#), [79](#)

changepoint Abrupt change in the outcome contingencies of a task. Because changepoints are un signaled, they are a major source of unexpected uncertainty and require strong behavioral adjustments to the new contingencies. [13](#), [21](#), [24](#), [30](#), [46](#)

changepoint probability Posterior probability of a changepoint. Because changepoints are often not directly observable, changepoint probability can guide learning under uncertainty about whether a changepoint occurred. Sometimes also referred to as surprise. [25](#), [47](#), [49](#)

default belief Refers to a default response that can easily be computed, such as relying on previous responses, yielding perseveration; or relying on environmental cues, yielding environmental control. [36](#), [37](#), [42](#), [50](#), [67](#), [78](#)

environmental control Behavior that is guided by environmental cues. It is assumed that environmental control emerges from weakened cognitive representations required for solving a task and is associated with lower task performance. [58](#), [67](#)

expected uncertainty Uncertainty about outcome contingencies in a task that emerges from variability across outcomes. It is assumed that agents expect and tolerate this form of uncertainty during learning. Expected uncertainty leads to belief uncertainty because the outcome variability prevents perfect knowledge of outcome contingencies. [11–13](#), [19](#), [30](#), [77](#), [78](#)

expected value In decision-making tasks involving rewards, the term expresses the expected, average benefits that are likely to result from choosing a particular

option (Rangel et al., 2008). More generally, expected value can refer to an expected outcome that is not necessarily rewarding. 9, 10, 12, 21, 45, 77, 78

hazard rate Parameter in changepoint tasks that controls the frequency of change-points. 24, 41, 56, 77, 78

learning rate In reinforcement-learning models, the learning rate controls the influence of a prediction error on the updated expected value. A high learning rate (near 1) indicates that the prediction error strongly influences the expected-value update. A low learning rate (near 0) reflects a weak influence. 18, 19, 25, 47, 49, 55

normative computations Computations that should be performed to solve a learning or decision-making problem optimally; often derived from Bayesian inference. 13, 22, 36, 55, 69

outcome contingency Mapping between states, actions, and outcomes in a task. 13, 21, 30, 35, 44, 45, 61, 74

perceptual uncertainty Uncertainty over stimuli/states that emerges from the unreliability of perceptual information. Often formalized with the belief state. 11–14, 30, 77

prediction error In reinforcement-learning models, the term refers to the difference between a received and an expected outcome. Prediction errors are used as a teaching signal that drives learning. See Box 1 for more details. 9, 10, 17, 46, 47, 75

reinforcement learning Computational approach to learning based on rewards during interaction with the environment. See Box 1 for more details. 11, 49

reversal learning Common paradigm to study learning and decision making. Often based on two or more choice options where the outcome contingency reverses between two states, one associated with a high expected value, the other associated with a low expected value. High performance requires an adjustment of choice behavior when a reversal is detected. [13](#), [23](#)

state A signal from the environment or an internal process that is relevant for an agent and either fully or partly available, mostly via perceptual processes. [9](#), [10](#), [12](#), [21](#), [44](#), [45](#), [55](#), [61](#), [77](#), [78](#)

surprise Unexpectedness of an outcome. In this dissertation, used in the context of unexpected changepoints and formalized with changepoint probability. Information theory offers additional measures of surprise (MacKay, [2003](#)). [25](#), [28](#)

surprise insensitivity Refers to a bias in the consideration of surprise, which leads to a reduced adjustment after a changepoint. [30](#), [33](#), [41](#), [56](#)

unexpected uncertainty Emerges from changes in the environment. The term denotes that changes lead to strongly unexpected observations beyond the random variability due to expected uncertainty. [11](#), [13](#), [21](#), [30](#), [78](#), [79](#)

volatility A central term in models of learning under unexpected uncertainty that assume drifting changes in the outcome contingencies. Volatility controls the extent to which the outcome probabilities of choice options change across trials. Formally, volatility is often defined as the variance of a distribution that models the drifting outcome contingencies. [13](#), [21](#), [23](#), [24](#), [30](#), [77](#), [78](#)

3 Introduction

Learning is pervasive in humans, animals, and artificial intelligence. It occurs when we acquire knowledge as a result of experience and is the basis for adaptive behavior (Terry, 2015). Examples range from basic phenomena such as classical conditioning to complicated decisions based on assessing the financial market.

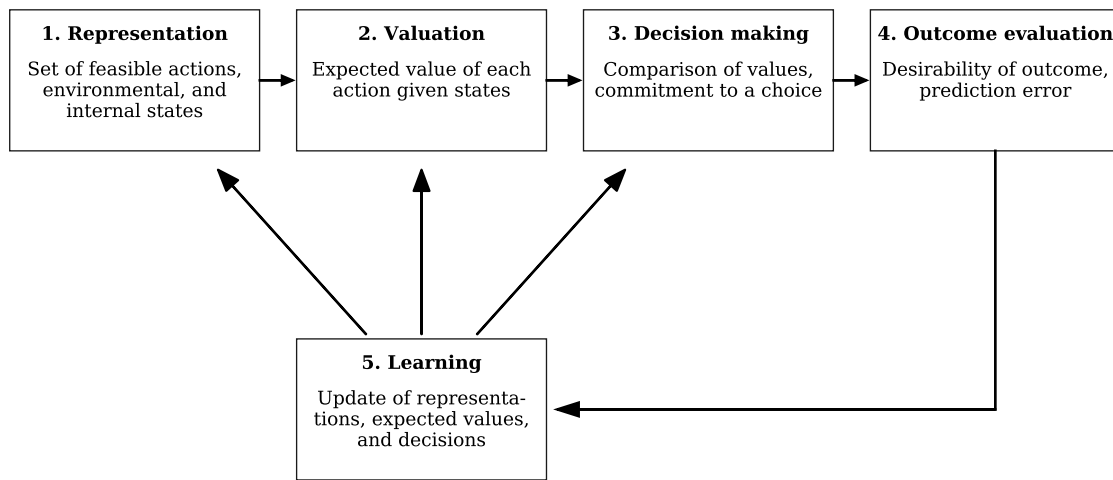


Figure 1. The cycle of adaptive behavior. Adaptive behavior incorporates five related steps. During interaction with the environment, an **agent** sequentially goes through these steps. First, a *representation* of the set of feasible actions (e.g., eating vs. rejecting a berry) as well as environmental and internal states. A **state** signals relevant information for an agent, which can originate in the agent’s environment (e.g., presence of strawberries and cherries) or internally (e.g., feeling hungry). Second, *valuation*, i.e., the assignment of an **expected value** to the feasible actions depending on the environmental and internal states (e.g., high vs. low expected taste value of eating a strawberry). Third, *decision making* based on a comparison of the actions’ expected values (e.g., eating a strawberry because of a higher expected value compared to cherry). Fourth, *outcome evaluation*, expressing the desirability of the received outcome, for example, based on **prediction errors** indicating the difference between an outcome and the expected value (e.g., comparison of the actual taste and the expected value of a strawberry). Fifth, *learning* in response to outcome evaluation, including an update of the representations (e.g., switching from hungry to not hungry), expected values (e.g., increasing the expected strawberry value), and decision-making policy (not further described here). The schematic was adapted from Rangel et al. (2008).

Learning takes place during interaction with the environment. To introduce how environment, behavior, and learning are related to each other, I commence with a framework based on earlier work by Rangel et al. (2008), which builds upon theories from psychology, computer science, and economics (Busemeyer & Johnson, 2004; Mas-Colell et al., 1995; Sutton & Barto, 1998). The basic idea is that adaptive behavior involves five related steps (Figure 1). When an animal, a person, or an

3 Introduction

artificial [agent](#) (Hassabis et al., [2017](#); Russell & Norvig, [2010](#)) interacts with the environment, they repeatedly cycle through these steps.

The first step is a representation of the set of feasible actions as well as environmental and internal [states](#) (*representation*). The term state refers to a signal from the environment or an internal process that is relevant for an agent and mostly available via perceptual processes. For example, an animal might be in an environment with different berry varieties such as strawberries and cherries. The set of feasible actions in this environment is eating or rejecting a berry. The two berry varieties comprise the environmental states, and the animal's internal state could be feeling hungry. Therefore, in the representation step, the animal needs to represent those states to identify what is currently relevant for its behavior.

Second, the [expected values](#) assigned to the available actions (*valuation*), referring to the expected benefits of the actions that are likely to result from choosing them. That is, when the animal chooses between eating strawberries or cherries to appease its hunger, it might expect a good taste of strawberries (high expected value). In contrast, if it expects a bad taste of cherries, it is less valuable for the animal to eat cherries (low expected value).

Third, the comparison of expected values to make an informed choice leading to the subjectively preferred outcome (*decision making*). Based on the higher expected value of strawberries compared to cherries, the animal may choose to eat strawberries and refuse cherries.

Fourth, an evaluation of the outcome of the choice (*outcome evaluation*). In the example, this could correspond to an evaluation of the perceived taste of the consumed strawberries, and whether the animal is still hungry or not.

Finally, the evaluated outcome is utilized to update the representations, expected values, and future choices (*learning*). For example, if the animal likes a strawberry better than expected (positive [prediction error](#)), learning leads to an increase of the expected strawberry value, and if a berry is worse than expected (negative prediction

error), learning leads to a decrease. Moreover, after having eaten a sufficient amount of strawberries, the animal's internal state might change from hungry to not hungry.

In summary, the computations sketched in this framework enable the animal to learn about the best course of action to adapt its behavior to internal needs and the current environment. For an example of how the cycle of adaptive behavior can be described within the framework of [reinforcement learning](#), see [Box 1](#).

3.1 Normative computations

The information on which humans and animals rely during learning is often remarkably uncertain. For example, in the representation step, two different types of berries may perceptually be hard to discriminate ([perceptual uncertainty](#)). In the valuation step, the animal may be confronted with uncertainty due to the natural variability of the berries' taste ([expected uncertainty](#)), which leads to uncertainty about the expected taste value ([belief uncertainty](#)). Uncertainty can also emerge from changes in the berries' expected value so that berries taste better at the beginning than at the end of summer ([unexpected uncertainty](#)).

Each of these forms of uncertainty makes it more difficult for an agent to learn from its experienced outcomes. However, although uncertainty can significantly corrupt learning, taking this into account generally improves the learning accuracy. For instance, learning the strawberries' expected value benefits from consideration of expected uncertainty due to natural taste variability. After the consumption of a bad berry, an animal which considers expected uncertainty would not immediately conclude that it dislikes all strawberries and pays attention to other strawberries as well to reduce its belief uncertainty over the value of the berries (see [Box 2](#) on [Bayesian inference](#) for a graphical illustration of this example).

3 Introduction

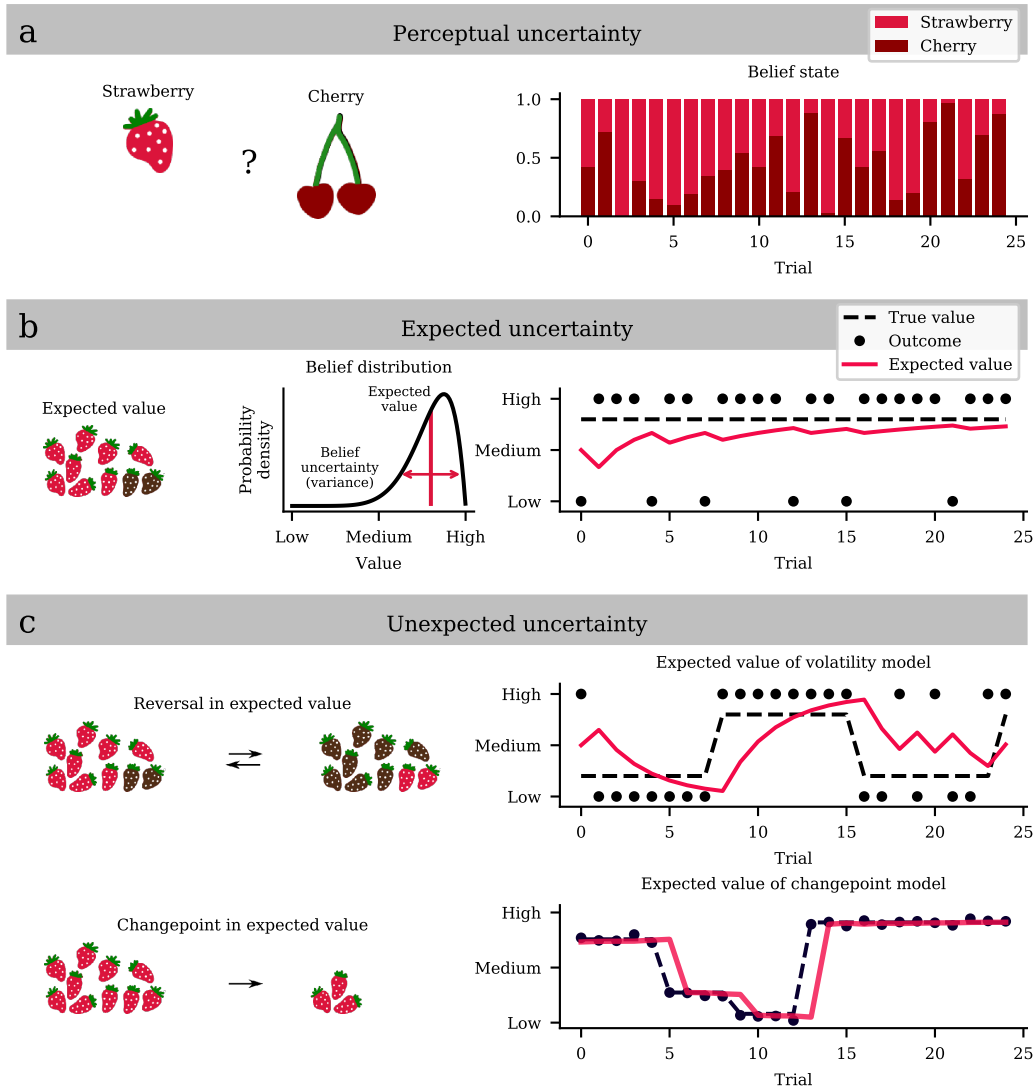


Figure 2. Decomposing uncertainty into three major forms. **a)** First, **perceptual uncertainty** concerns the representation step, particularly the identification of the environmental **states**, because it emerges from the unreliability of perceptual information about those states. For example, perceptual uncertainty due to bad light conditions can impede distinguishing strawberries and cherries when an animal forages for food. One can formalize perceptual uncertainty with the **belief state** that indicates the state probabilities based on the available perceptual information such as the probability of strawberry versus cherry. In the example (right plot), the belief state weakly favors strawberry in trial 0, strongly favors strawberry in trial 2, and strongly favors cherry in trial 13. **b)** Second, **expected uncertainty** primarily concerns the valuation step. Expected uncertainty emerges from variability across outcomes and makes it impossible to know the **expected value** of an action exactly. For example, due to natural variability, only 80% of the strawberries taste well. Using a **belief distribution**, one can model the animal's belief over the strawberry value, as shown in the middle panel. The mean (red vertical line) of the distribution reflects the animal's expected taste value, i.e., the average probability of obtaining a good strawberry. The variance (horizontal arrows) reflects the **belief uncertainty**. Here, a lower variance would indicate less belief uncertainty. As shown in the right panel, based on the belief distribution, one can model the evolution of the animal's expected value. In this example, the animal experiences several good and bad berries, based on which its learned expected value approximates the true value of 80%.

3.1 Normative computations

Figure 2 (Continued). c) Third, **unexpected uncertainty** is the consequence of changes in the environment. Changes specifically affect the **outcome contingency** in a task, i.e., the mapping between states, actions, and outcomes. I distinguish approaches that assume volatile environments (**volatility**, first row) and **change point** environments (second row). A typical example of volatility is **reversal learning** where the outcome contingency switches between two states, for example, associated with a high vs. low expected strawberry value. Volatility models assume that these changes occur smoothly over time (right panel). As indicated by the red line, these models compute drifting changes in the expected values, which approximate the true strawberry value after the experience of multiple berries. In contrast, change point models assume discrete change points in the outcome contingency. In the change point example (second row), the expected value of berry availability changes at discrete time points. The availability of strawberries is initially high (left set of berries) but suddenly changes so that only a few berries are available (right set of berries). As shown in the right plot, a change point model tries to detect those change points and rapidly adjusts the expected value accordingly.

We are now in a position to formulate a framework that dissociates three major forms of uncertainty that plague learning (Figure 2). First, **perceptual uncertainty** concerns perceptual information processing and is primarily related to the representation step of the cycle of adaptive behavior. The second form, **expected uncertainty**, arises from random variability between outcomes and ultimately leads to **belief uncertainty**. Third, **unexpected uncertainty**, which emerges from changes in the environment. Within the cycle of adaptive behavior, expected and unexpected uncertainty primarily affect the valuation step.

The proposed framework gives an overview of **normative computations** that agents should perform to solve the learning problems presented in the following sections optimally. These normative computations are derived from **Bayesian inference**, a method that can be utilized to define uncertainty mathematically (see Box 2). In the following, I will borrow the term **belief** from the language of Bayesian inference (Russell & Norvig, 2010) to refer to an agent’s knowledge of an uncertain variable such as an environmental state (**belief state**) in the representation step or the expected value of an action (**belief distribution**) in the valuation step.¹

¹Please note that other frameworks exist that distinguish between several forms of uncertainty. Yu and Dayan (2005) proposed the terms expected and unexpected uncertainty. Bach and Dolan (2012) distinguish between uncertainty about sensory information, states, rules, and outcomes. This framework is, however, more directly focused on uncertainty during decision making than learning. Soltani and Izquierdo (2019) review the literature on expected and unexpected uncertainty but do not include perceptual uncertainty. Moreover, Ma and Jazayeri (2014) proposed a framework on neural coding of uncertainty. Finally, I will not explicitly consider effects of motor uncertainty during action selection on learning (e.g., Fleming et al., 2013; Frömer et al., 2020).

3 Introduction

3.1.1 Perceptual uncertainty

The form of **perceptual uncertainty** is primarily related to the representation of environmental states. Virtually all learning processes require the processing of perceptual information about environmental states, such as the currently available berry varieties. Previous work traditionally studied learning under conditions in which the perceptual input is clearly interpretable. However, learning is often impeded by perceptual uncertainty because perceptual information is unreliable or distorted by variability inherent in sensory systems (Dayan & Daw, 2008; Gold & Heekeren, 2014; Gold & Shadlen, 2007; Law & Gold, 2009; Summerfield & Tsetsos, 2012).

As illustrated in Figure 2a, an animal could be confronted with perceptual uncertainty whether it found a strawberry or cherry during foraging (e.g., under bad light conditions). Although the animal can not exactly identify the current state of the environment (strawberry vs. cherry), it can rely on a belief about the state. I refer to this as the **belief state**, which expresses the state probability depending on the currently available uncertain sensory information (Daw, 2014; Dayan & Daw, 2008). More distinct belief states (e.g., 95 % in favor of strawberry vs. 5 % of cherry) indicate low perceptual uncertainty and more similar probabilities (e.g., 55 % vs. 45 %), high uncertainty. Formally, these computations are based on Bayes' theorem, which I introduce in Box 2.

Intuitively, the belief state should regulate how much an agent learns from an outcome. If the animal is sure that it has eaten a tasty strawberry and not a tasty cherry (99 % vs. 1 %), it can adjust the strawberries' expected value accordingly. In contrast, if it does not exactly know which berry it consumed (55 % vs. 45 %), the animal should learn more cautiously about the values because it may have confused the berries.

Although it is somewhat surprising that the cognitive-neuroscience literature has not yet proposed a Bayesian-inference model that learns according to this idea

(to the best of my knowledge), perceptual decision-making research provides relevant prior results on perceptual uncertainty. This field of research studies choices based on available perceptual evidence (Heekeren et al., 2008), often using stimuli that are uncertain because they are noisy or weakly presented so that participants can not clearly perceive them (Gold & Stocker, 2017).

Several studies provided rewards for correct choices or punishments for incorrect responses. In these cases, an estimate of the uncertainty with which a stimulus can correctly be detected or classified can be useful to increase the likelihood of obtaining a reward or avoiding punishment. For example, Whiteley and Sahani (2008) asymmetrically penalized participants for providing incorrect responses, i.e., they penalized incorrect answers on the one side more heavily than on the other side. Under high perceptual uncertainty, participants tended to avoid more heavily penalized answers, suggesting that they chose the less negative option to avoid high punishment in trials where they could hardly identify the stimuli due to perceptual uncertainty. In contrast, under low perceptual uncertainty, choices were primarily driven by the available perceptual information and less so by differences in punishment.

There is also compelling evidence that rhesus monkeys consider perceptual uncertainty to similarly maximize collected rewards (Kiani & Shadlen, 2009). In a random-dot task with varying difficulty levels, Kiani and Shadlen rewarded monkeys for correct choices. In half of the trials, the monkeys had the opportunity to opt-out of the decision to obtain a smaller but guaranteed reward. They chose this option particularly in difficult trials, indicating that they utilized an estimate of perceptual uncertainty to determine when the small but guaranteed reward was more valuable than a larger reward that depended on the correctness of the perceptual choice. Here, neurons in the parietal cortex seemed to be involved in the computation of the perceptual choice as such, and the corresponding perceptual uncertainty about the visual input.

3 Introduction

Studies applying sequential-sampling models provided further evidence that humans and animals consider perceptual uncertainty during perceptual decision making (Forstmann et al., 2016). Based on reaction times and choice data, these models describe decision-making processes in terms of an accumulation of uncertain perceptual information integrated over time. According to such models, individuals report their decisions when the accumulated evidence reaches a decision threshold. The study results in monkeys by Kiani and Shadlen (2009) were qualitatively consistent with a sequential-sampling model that framed the computation of the parietal neurons to opt-out of a decision depending on perceptual uncertainty in terms of a Bayesian evidence accumulation process.

Sequential-sampling models further suggested that asymmetric rewards or losses bias evidence accumulation towards the high-reward/low-loss option. Most results support the idea that unequal reward or punishment serves as prior information that shifts the starting point of the evidence accumulation process towards the more desirable option. In effect, especially under high perceptual uncertainty, choices are often biased in the direction of the high-reward/low-punishment option (Diederich & Busemeyer, 2006; Feng et al., 2009; Gold & Shadlen, 2001; Rorie et al., 2010; Simen et al., 2009). Similar studies in humans indicated that neural activity in the parietal and prefrontal cortex (PFC) correlates with perceptual uncertainty during perceptual decision making (Mulder et al., 2012; Summerfield & Koechlin, 2010). Finally, humans seem to utilize perceptual uncertainty in tasks with unobservable and occasionally changing response rules (Purcell & Kiani, 2016). In this study, participants adjusted their responses after feedback about the correct rule more quickly when they suffered from lower perceptual uncertainty. The authors formally described these computations with a sequential-sampling model that dynamically integrated uncertain perceptual evidence and response feedback.

Studies in monkeys and mice that used a reinforcement-learning framework provided more recent evidence that perceptual uncertainty affects learning (Lak

et al., 2017; Lak et al., 2020). These authors suggested that the brain weights reward processing and learning as a function of belief states with a neural mechanism at the level of the reward prediction error in the striatum (see Box 1 on reinforcement learning). Moreover, results in rats demonstrate another potential neural mechanism in the anterior cingulate cortex (ACC) that leads to a similar regulation of learning in response to perceptual uncertainty (Stolyarova et al., 2019).

In summary, perceptual decision making provides evidence that humans and animals represent and utilize perceptual uncertainty to adapt their behavior more cautiously when only weak perceptual information is available, potentially via neural mechanisms including the PFC and parietal cortex. More recent work suggests that perceptual uncertainty guides learning from feedback at the level of the striatum and ACC. However, the field currently lacks Bayesian-inference models indicating how agents should normatively adjust learning under perceptual uncertainty that help us to test to what degree humans learn according to these principles.

3 Introduction

Box 1 | Reinforcement learning

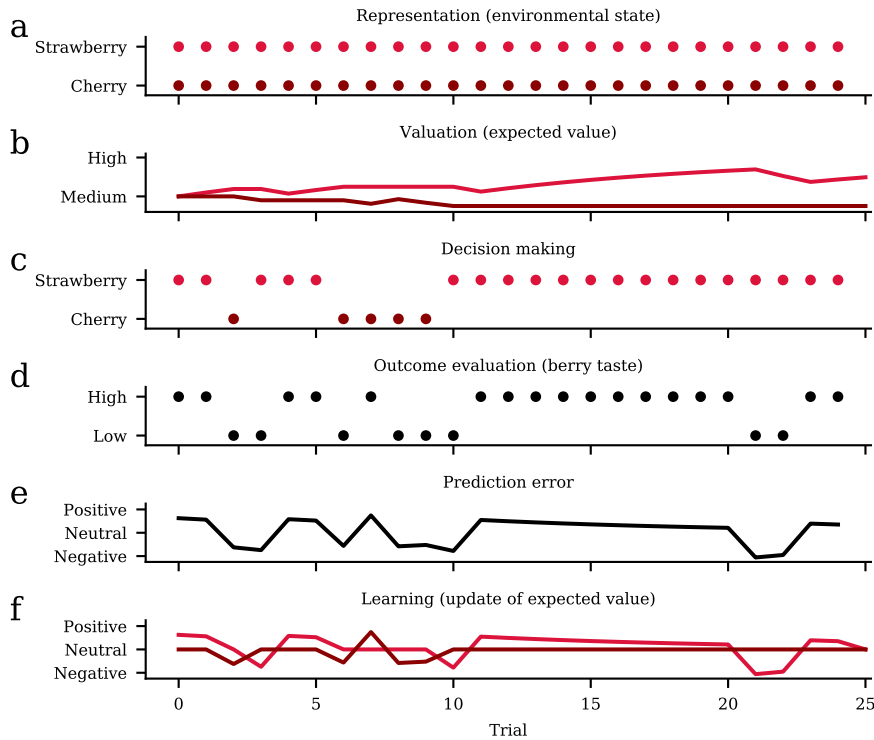
Learning from reward and punishment is a fundamental ability of humans and other animals that is essential for adaptive behavior. In psychology and neuroscience, the reinforcement-learning framework emerged from classical experiments on animal conditioning and learning algorithms in computer science (Daw & Tobler, 2014; Dayan & Daw, 2008; Rescorla & Wagner, 1972; Sutton & Barto, 1998). A popular idea of this approach is that learning results from comparing the reward received at a particular point in time with the expected value. Formally, we can express this as the prediction error

$$\delta_t = o_t - v_t(s_t), \quad (3.1)$$

where the index t denotes a trial, s_t the state of the environment, o_t the outcome, and v_t the expected value. As shown in the example on the [Cycle of adaptive behavior](#) plotted below, the state s_t could refer to the berry variety (strawberry vs. cherry), the outcome o_t to the taste of the consumed berry (high vs. low), and the expected value v_t to the initially expected taste of the berries (e.g., medium). The animal then uses the prediction error to adjust its expected value at the next trial $t + 1$ according to

$$v_{t+1}(s_t) = v_t(s_t) + \alpha \cdot \delta_t, \quad (3.2)$$

which expresses that the current expected value v_t of state s_t is combined with the prediction error δ_t and multiplied by the [learning rate](#) α that determines how strongly δ_t influences the expected value $v_{t+1}(s_t)$. Based on this framework, we can model the five steps of the cycle of adaptive behavior. **a)** In the representation step, the animal has to represent the two environmental states, corresponding to strawberry and cherry (next to internal states, which are ignored for simplicity). **b)** In the valuation step, the animal assigns an expected value to both berry varieties, which differently evolve across time. **c)** In the decision-making step, the animal chooses one of the berries, preferably the one with the higher expected value. **d)** During outcome evaluation, the animal perceives the taste, which is either high or low. **e)** Subsequently, it computes the prediction error (eq. (3.1)). **f)** Finally, the animal updates its expected value according to eq. (3.2). That is, the expected value update shown in (f) indicates to which degree the animal adjusts the expected value in (b) from trial to trial. Thus, the animal learns a higher expected value of strawberries than cherries (b) and tends to choose strawberries more often (c).



3.1.2 Expected uncertainty

Within the cycle of adaptive behavior, **expected uncertainty** primarily affects the valuation step. This form of uncertainty arises from random variability across outcomes. When the learner is aware that outcomes are plagued by uncertainty, the term expected uncertainty is often used to indicate that some degree of randomness is expected and tolerated during learning (Yu & Dayan, 2005).

As shown in Figure 2b, when the animal forages for strawberries, it suffers from expected uncertainty due to outcome variability (80 % probability of good berry and 20 % bad berry). Consequently, the animal can not exactly know the true strawberry value (i.e., 80 % probability of good berry). However, it can rely on a belief about the value, which we can express with a belief distribution over the strawberry value (Figure 2b, middle panel). The mean of the belief distribution denotes the animal's expected value, i.e., the average probability of obtaining a good strawberry. The variance characterizes the animal's **belief uncertainty** over the value (sometimes also referred to as estimation uncertainty), where a lower belief uncertainty (lower variance) indicates that the animal has learned a more reliable expected value of the strawberry.

Therefore, when learning the strawberries' expected value, the animal has to average out the outcome variability by combining multiple experienced outcomes. As shown in Figure 2b (right panel), doing so in response to experiencing good and bad berries results in an expected value that converges towards the true value. The key prediction of Bayesian learning under expected uncertainty is that the influence of a new outcome on the expected value decays over time. This decay in the **learning rate** is related to **belief uncertainty**. At the beginning of learning, the belief about the value is highly uncertain and, therefore, strongly influenced by a new outcome. Conversely, after multiple outcomes, the agent is more confident, and a new observation changes the belief distribution merely weakly (see also Box 2 on Bayesian inference for illustration).

3 Introduction

Several studies suggest that humans adjust their learning rate under expected uncertainty according to the principles of Bayesian inference. For example, Payzan-LeNestour and Bossaerts (2011) used a six-armed-bandits task with considerable expected uncertainty over the bandits' payoffs. The authors compared a Bayesian learning model that regulates learning according to belief uncertainty with classical reinforcement-learning algorithms lacking this ability. The results of this study favored Bayesian learning, suggesting that human participants can flexibly determine how much they learn from new observations under expected uncertainty. Follow-up neuroimaging work showed expected-uncertainty-related activity in the ACC and PFC (Payzan-LeNestour et al., 2013).

Meyniel and colleagues presented further evidence that belief uncertainty impacts learning (Meyniel & Dehaene, 2017; Meyniel, Schlunegger, et al., 2015). In a learning task, they instructed participants to provide explicit confidence ratings about their learned beliefs. These confidence ratings were in line with "belief certainty" (i.e., the inverse of belief uncertainty) of a Bayesian-inference model, and related to brain activity in the PFC. Thus, this result also suggests that human learners use an estimate of belief uncertainty to regulate learning but extends previous findings to show that this estimate can be explicitly reported and related to brain activity.

Nassar and colleagues (Nassar et al., 2010) showed that participants report more substantial belief updates after observing only a few compared to multiple uncertain outcomes across an extended period. The authors quantitatively captured this result in a Bayesian belief-updating model suggesting that participants regulated their belief updates depending on the amount of belief uncertainty. Functional neuroimaging work using this model revealed that activity in the medial PFC and parietal areas correlates with the uncertainty-dependent learning regulation, suggesting that these areas support such computations (McGuire et al., 2014).

Finally, an influential theory states that the neurotransmitter acetylcholine signals the level of expected uncertainty (Yu & Dayan, 2005). Several studies that

tested this idea offer support along these lines, for example, pharmacological work in humans (Marshall et al., 2016; Moran et al., 2013; Vossel et al., 2014), and optogenetics in mice (Hangya et al., 2015).

In summary, expected uncertainty comes from variability between outcomes. Therefore, learning under expected uncertainty leads to uncertain beliefs about the expected value of choice options, mainly when an agent has only observed a few uncertain outcomes. Several studies suggest that humans flexibly adjust learning under expected uncertainty, which is associated with brain activity in the ACC, PFC, parietal cortex, and the neurotransmitter acetylcholine.

3.1.3 Unexpected uncertainty

The form of **unexpected uncertainty** also primarily affects the valuation step. Unexpected uncertainty emerges from changes in the environment and implies such changes lead to strongly unexpected observations beyond the random variability through expected uncertainty (Yu & Dayan, 2005). The literature on learning in these environments can broadly be divided into approaches using models that assume drifting outcome contingencies modulated by **volatility** and abruptly-changing contingencies related to a **change point**.

Volatility

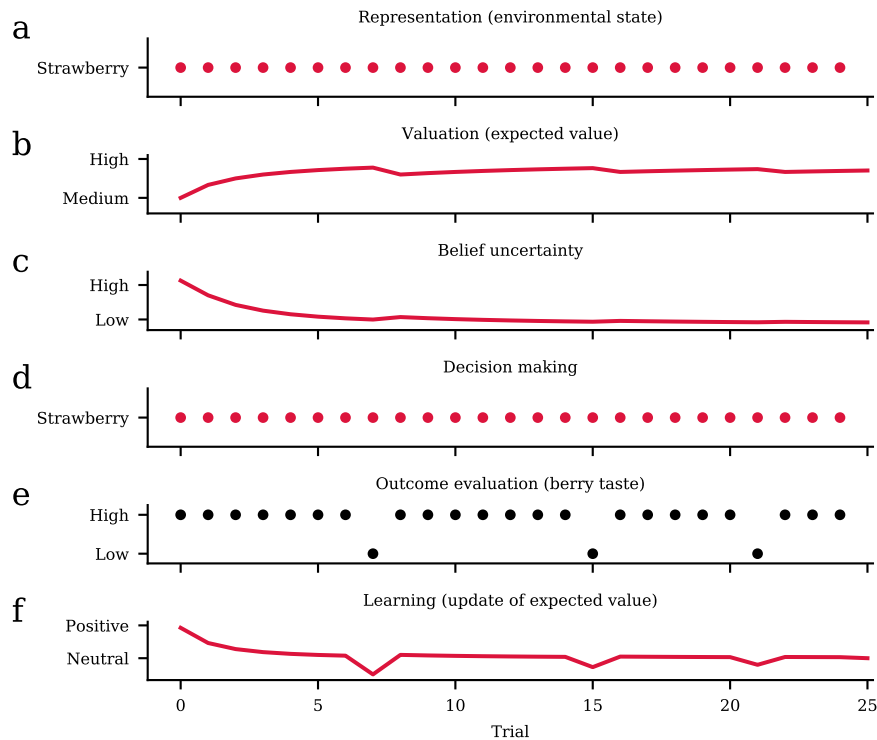
Changes in the environment affect the **outcome contingency** in a task, referring to the mapping between **states**, actions, and outcomes. A typical task is reversal learning, where the **expected value** of choosing an option reverses between two states (Figure 2c, upper row). In one state, the expected strawberry value is high because it is likely that the animal picks a good berry. In the other state, the expected value is low because the animal is unlikely to find a good berry. To adjust to these changes in the outcome contingencies, the animal has to update its learned expected value of the strawberries accordingly.

Box 2 | Bayesian inference

Bayesian inference is a mathematical method to update probability distributions based on newly arriving information. In psychology and neuroscience, the application of this method has proven to be suitable for the development of computational models of cognition. An appealing characteristic is that it offers insights into the **normative computations** indicating how a learning or decision-making problem under uncertainty should optimally be solved. An important term in the language of Bayesian inference is **belief**, referring to an **agent's** knowledge of an uncertain variable. In the **Cycle of adaptive behavior**, one could, for example, use this method to learn a belief about the expected value v_t of the strawberries' taste. The central component is Bayes' theorem, which describes how the probability of an event (v_t , indicating the probability of obtaining a tasty strawberry) can be computed after a new outcome o_t when we have prior knowledge about the event, i.e.,

$$p(v_t|o_t) = \frac{p(o_t|v_t)p(v_t)}{p(o_t)}, \quad (3.3)$$

where $p(v_t|o_t)$ refers to the posterior probability after incorporating the new outcome information, which is obtained by consideration of the prior information about the event $p(v_t)$ and the likelihood of the new outcome o_t , conditional on v_t being true, i.e., $p(o_t|v_t)$. $p(o_t)$ in the denominator refers to the marginal probability of o_t and acts as a normalization factor. As shown in the example on the cycle of adaptive behavior in the plot below, v_t can be interpreted as the strawberries' expected value, and the outcome o_t denotes the taste of the strawberry (high vs. low). The prior $p(v_t)$ could denote the initially expected value of the strawberries. **a)** To illustrate the evolution of the animal's expected strawberry value according to Bayesian inference, the animal is always in the strawberry state. **b)** In the valuation step, the animal tracks the strawberries' expected value, which develops from medium to high. **c)** Because Bayesian inference is concerned with updating a probability distribution, we can explicitly model an agent's uncertainty regarding the belief about the expected value. I refer to this form of uncertainty as **belief uncertainty**, which formally denotes the variance of the belief distribution (see also Figure 2b, middle panel). As shown in the plot below, belief uncertainty slowly decays across trials, reflecting that the agent's belief over the strawberry value is increasingly certain. **d)** Because the animal is constantly in the strawberry state, it only eats strawberries and subsequently evaluates the perceived taste. **e)** Thus, during outcome evaluation, the animal experiences mostly good berries. **f)** Finally, in the learning step, the animal updates the expected value under consideration of belief uncertainty using eq. (3.3).



3.1 Normative computations

A common idea about the computational underpinnings of learning under unexpected uncertainty is that humans and animals assume continuous drifts instead of abrupt changes in the outcome contingencies. In this scenario, [volatility](#) controls to which degree the outcome contingency is assumed to change across time. Therefore, environments with more unexpected uncertainty are assumed to be more volatile; that is, outcome contingencies drift more heavily. It is important to keep in mind that although the reversals occur at discrete time points, volatility models assume that behavioral adjustments to such changes occur continuously. As illustrated in the figure, the animal's expected value of the strawberries (red line) smoothly develops across time.

An influential study proposed a Bayesian-inference model for [reversal learning](#) in such environments (Behrens et al., 2007). To optimally respond to reversals in the task, the model by Behrens et al. (2007) not only learns the choice options' expected value, but also the volatility itself. In effect, the model can flexibly adjust to environments with different levels of volatility such that in more volatile environments, the model more quickly adjusts to the contingency changes in the task. The empirical results on learning under volatility in humans were in line with this idea and suggested that subjects who performed better were better calibrated to volatility. Moreover, participants showing more ACC activity achieved a higher learning performance, suggesting that the ACC might be involved in regulating learning in response to volatility.

While the model by Behrens et al. (2007) is suitable for simulating learning behavior, it is hard to estimate it based on participant data, primarily because this would take too much time due to a large number of required numerical calculations (Mathys et al., 2011). However, model estimation offers the appealing possibility to uncover individual differences in the latent processes governing learning that might deviate from a normative model (see [Model estimation](#)).

3 Introduction

Therefore, more recent studies proposed computational models that approximate optimal Bayesian inference for a fraction of the computational complexity, allowing for subject-specific parameter estimates reflecting the computational underpinnings of learning behavior. A prominent example is the hierarchical Gaussian filter (HGF), which performs hierarchical Bayesian inference (Mathys et al., 2011; Mathys et al., 2014). One can apply the HGF to a broad array of learning tasks, which offers the possibility to extract general principles of learning under uncertainty. A previous application of this model, for example, provided evidence for an involvement of the basal forebrain in belief updating of volatility (Iglesias et al., 2013). However, the generalizability of the HGF may come at the cost of missing learning dynamics that are specific to particular tasks (Bröker et al., 2018). More recent studies, therefore, offered simpler, alternative models that are less able to generalize across various tasks but might better capture uncertainty-driven learning in the tasks they were designed for (Bröker et al., 2018; Piray & Daw, 2020).

In summary, in volatile environments, outcome contingencies smoothly change, and the volatility controls the extent of this change during each trial. The different volatility models shown in this section have their strengths and weaknesses. However, the bigger picture that this line of research sketches is that humans can utilize unexpected uncertainty in volatile reward environments to regulate learning flexibly. Finally, at the neural level, volatility-related activity enabling such inferences was found in the ACC and the basal forebrain.

Changepoints

An alternative way of responding to unexpected uncertainty is to assume discrete **changepoints** during learning. Thus, in contrast to the assumption in **volatility** models that outcome contingencies change smoothly, changepoint models assume discrete changes or jumps. In this scenario, the **hazard rate** controls how often changes occur. In Figure 2c (lower row), the changepoint environment depicts a situation in which

3.1 Normative computations

the expected value of strawberry availability changes systematically at particular points in time. For example, in some periods, strawberries are widely available, while they are rare in other periods. Adaptive behavior in this environment requires keeping track of the expected value of the strawberry availability by detecting and responding to the changepoints. If a changepoint in the strawberry availability occurs, the animal should update the expected value to the new outcome contingencies. In contrast, if no change occurs, there is no need for a strong adjustment because previous strawberry-availability values are most predictive for future outcomes. Please note the evolution of the animal's expected strawberry value (red line), which similarly reflects such discrete changes in the environment.

Nassar and colleagues (Nassar et al., 2010) have conducted an ongoing line of research on learning in changing environments. They developed a Bayesian learning model that captures how humans and animals learn in changing environments by detecting the changepoints themselves, expressed as [changepoint probability](#) or [surprise](#). When changepoint probability is high, the model increases its [learning rate](#) to adjust to the potential changepoint. When changepoint probability is low, the model more strongly considers expected uncertainty similar as described above. This model describes human learning in changing environments better than a reinforcement-learning model with a fixed learning rate, i.e., a learning rate that is not modulated by the probability of a change in the environment (Nassar et al., 2010). Neuroimaging work showed that learning in response to changepoint probability is reflected in visual areas and the medial PFC (McGuire et al., 2014).

Furthermore, as shown in Box 3, I collaborated with Nassar and Frank to study surprise-driven learning in different statistical contexts using electroencephalography (EEG) (Nassar, Bruckner, & Frank, 2019). In this study, we demonstrated that the event-related P300 component (Ullsperger, Danielmeier, et al., 2014; Ullsperger, Fischer, et al., 2014) signals surprising outcomes, which subsequently leads to an adjustment to a changepoint through a high [learning rate](#). In contrast, in a different

3 Introduction

statistical context where surprise indicates occasional oddballs, this component is associated with lower learning rates to suppress the influence of irrelevant information on learning.

In line with this idea, O'Reilly et al. (2013) showed that belief updating in response to surprising outcomes is related to ACC activity, while the perception of an unexpected outcome independent of a belief update has neural signatures in the posterior parietal cortex. Furthermore, Meyniel and Dehaene (2017) compared a learning task in the visual domain with a task with auditory stimuli to test which learning signals might be independent of the sensory modality of the stimuli. They found such modality-independent surprise representations in frontal and temporal areas. Finally, in the six-armed-bandits task described above, Payzan-LeNestour et al. (2013) reported unexpected-uncertainty-related signals in posterior areas and the temporal cortex.

In summary, in abruptly-changing environments, outcome contingencies occasionally jump from one state to the other, and the hazard rate controls the frequency of changepoints. Model-based studies assuming these types of changes provide evidence that humans can respond to unexpected uncertainty by up-regulating learning after surprise signals that indicate the probability of a changepoint. Changepoint-related brain activity is reflected in sensory and higher-order areas, including the posterior, temporal, and prefrontal cortex.

Finally, Yu and Dayan (2005) argued that norepinephrine transmitters might signal unexpected uncertainty. One way to test this idea is pupillometry, which could reflect noradrenergic brainstem activity (Joshi et al., 2016). Several studies using this approach support the idea that norepinephrine signals unexpected uncertainty (Krishnamurthy et al., 2017; Muller et al., 2019; Nassar et al., 2012). Pharmacological work (Marshall et al., 2016) and the neuroimaging findings by Payzan-LeNestour et al. (2013) that showed that unexpected uncertainty is associated with noradrenergic brainstem activity provide further support in favor of this hypothesis.

3.1.4 Summary

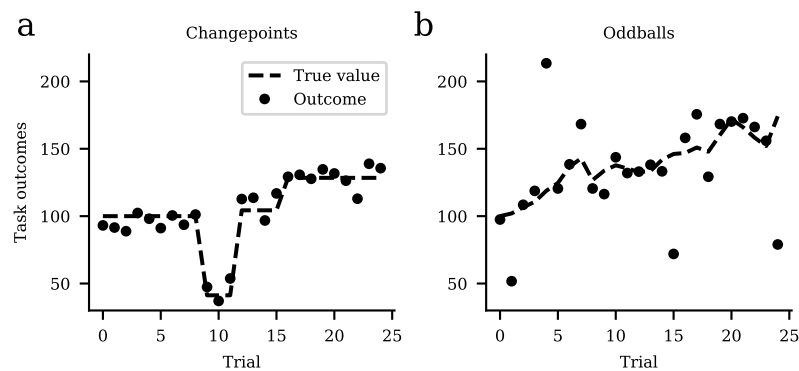
The framework shows multiple factors that should optimally influence learning under uncertainty. Perceptual uncertainty is the consequence of uncertain sensory information. Here, the belief state expresses the probability of an uncertain state conditional on the available perceptual information. The belief state should regulate the influence of new information on learning so that an agent can take into account that it might misinterpret stimuli under high perceptual uncertainty. Expected uncertainty emerges from variability across outcomes. Here, belief uncertainty indicates how reliably beliefs about expected values have been learned so far and should regulate the degree of learning to average out the variability. Unexpected uncertainty arises from changes in outcome contingencies. I have distinguished between approaches assuming continuously drifting changes related to volatility and changepoints related to the hazard rate. While these approaches differ in their assumptions of how changes are perceived, they have in common that unexpected uncertainty should lead to a belief update to adjust to such changes. Across the different forms, the framework highlights that multiple brain areas respond to uncertainty, particularly parietal and frontal areas, including the medial PFC and ACC.

3 Introduction

Box 3 | Surprise

Unexpected uncertainty emerges from unsignaled changes in the environment. In changepoint models, these unexpected events are assumed to elicit [surprise](#) to initiate an adjustment to the new outcome contingencies. In Nassar, Bruckner, and Frank (2019), we dissociated surprising outcomes from the degree to which they demand learning using a changepoint task, Bayesian modeling, and electroencephalography (EEG). As shown in the figure, the task comprised two statistical contexts. **a)** The first one required learning under expected and unexpected uncertainty, particularly the regulation of learning after changepoints. The dashed line ("True value") refers to the most likely trial-by-trial outcome, which corresponded to the mean of a Gaussian distribution. Outcomes were drawn from this distribution and deviated from the mean due to considerable variance inducing expected uncertainty. Unexpected uncertainty was due to the occasional changepoints that led to a change in the mean of the outcome-generating distribution. **b)** The second context did not feature changepoints but instead oddballs that were uninformative for future outcomes. Here "mean" similarly refers to the most likely outcome, which, however, slowly drifted across trials in contrast to the changepoint context. In this context, oddballs were perceptually similarly surprising as changepoints. Crucially, oddballs should not lead to a behavioral adjustment because they do not contain any information about the mean of the outcome-generating distribution.

This contextual manipulation allowed us to examine the regulation of learning in response to comparably surprising outcomes (changepoints vs. oddballs) that demand opposing learning-rate adjustments. While changepoints should lead to more learning, oddballs should reduce the learning rate because they contain no useful information. The results showed that the P300, a well-known event-related potential linked to learning (Ullsperger, Danielmeier, et al., 2014; Ullsperger, Fischer, et al., 2014), signals the presence of surprise. In particular, we found that this signal is associated with greater learning in the changepoint condition but less learning in the oddball condition. This finding suggests that the P300 component signals unexpected outcomes, leading to an up- or down-regulation of learning depending on the current environment.



3.2 Uncertainty biases

In the previous section on normative computations that should determine learning under uncertainty, I presented findings showing that humans and animals take into consideration different forms of uncertainty during learning and decision making. However, there is no guarantee that uncertainty shapes learning to the same extent as in a normative model that performs at the theoretical upper limit. Instead, a common finding is that humans and animals show specific deviations from these models (Ma, 2019).

3.2.1 Perceptual uncertainty

Biases in considering perceptual uncertainty can lead to belief states that are too certain. For instance, although an animal acts under bad light conditions in which the perception of the environment is uncertain, it might conclude with 100% confidence that a stimulus is a strawberry and not a cherry (Figure 3a). Similar to the example, research by Fleming et al. (2013) showed that humans sometimes consider perceptual uncertainty insufficiently for optimizing perceptual choices. Participants reported the predominant motion direction of moving dots in a random-dot kinematogram by rapidly pointing to targets on a touch screen with varying target sizes. The randomly moving dots induced perceptual uncertainty about the predominant motion direction while the requirement to report the decision quickly induced motor uncertainty about the chances of missing the target. Bayes-optimal decisions would consider both forms of uncertainty to maximize the probability of obtaining a reward. However, the results of the study indicated that participants often ignored motor uncertainty. They placed too much weight on the perceptual information and too little on the probability of missing the target. Fleming and colleagues suggested that this finding reveals a phenomenon referred to as *categorical perception*. When confronted with uncertain perceptual information, the perceptual system may categorically commit to one interpretation of a stimulus while disregarding others.

3 Introduction

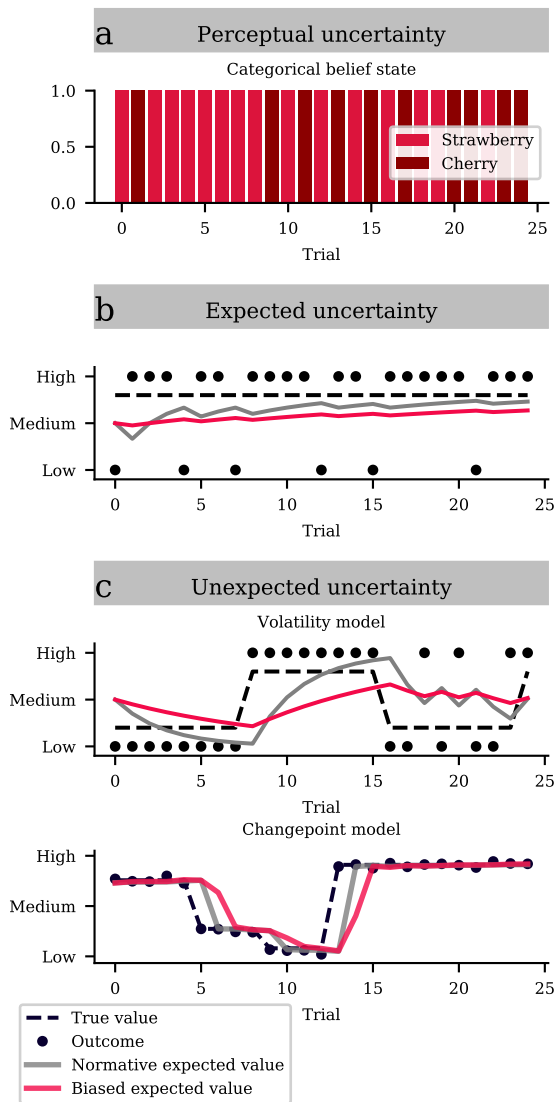


Figure 3. Uncertainty biases. Due to inherent biases, individuals might sub-optimally consider uncertainty during learning. **a)** Regarding perceptual uncertainty, they might represent belief states that are too certain. In the example plot showing the agent’s belief states of strawberries and cherries, belief states are categorically in favor of strawberry (e.g., trial 0) or cherry (e.g., trial 1). Thus, compared to normative belief states (see Figure 2a) based on the currently available perceptual information, categorical belief states discard perceptual uncertainty about the environmental state. **b)** Learning under expected uncertainty can be corrupted by sub-optimal consideration of belief uncertainty during learning, referred to as uncertainty underestimation. In this example, underestimation of belief uncertainty leads to inaccurate learning of the strawberries’ expected value. In comparison to the normative model that optimally considers belief uncertainty (gray line), the agent learns at a much slower pace and fails to learn an expected value that approximates the true probability of obtaining a good strawberry (80%). **c)** Concerning unexpected uncertainty, individuals might show biases in the adjustment to changes in the environment. One possibility is that they underestimate the level of volatility; that is, they might assume smaller changes in the environment. This bias leads to a slower adjustment to changes in the outcome contingency. For example, the animal might assume that the strawberries’ expected value changes only marginally across the year (although it changes relatively frequently). Therefore, its adjustment to changes in the environment is substantially slower than in the normative model (gray line). Finally, according to changepoint models, individuals might suffer from a reduced ability to respond to changepoints in the environment (surprise insensitivity). In this case, an animal would fail to detect that a change has taken place. Consequently and in contrast to the normative model (gray line), the animal only slowly adjusts its expected value after a change in the outcome contingencies.

A related line of work suggests that categorical perception can, in some circumstances, emerge from previous perceptual choices (Luu & Stocker, 2018; Stocker & Simoncelli, 2007). In the study by Luu and Stocker (2018), participants performed a perceptual two-stage task. In the first stage, they categorically judged the orientation of a stimulus (e.g., left vs. right). In the second stage, they indicated the actual orientation of the same stimulus on a continuous scale. The results showed that categorical judgments biased continuous orientation judgments. Luu and Stocker captured this categorical bias in a Bayesian-inference model that considered not

only perceptual evidence for the orientation judgments but also the prior categorical perceptual choice. According to this modeling framework, the categorical-choice bias arose from self-consistency during perceptual inference. In contrast to the true task structure where the categorical and continuous judgments are independent of each other, participants might have assumed a causal dependency between the two stages as a result of which the continuous orientation judgment was dependent on the initial categorical choice. More generally, this study illustrates that Bayesian models can also be applied to study biases in uncertainty computations by examining participants' prior assumptions about the task structure.

Together, these findings could suggest that learning under perceptual uncertainty is similarly affected by categorical perception. However, because previous work only examined categorical perception with respect to perceptual decision making, it remains an open question of how learning under perceptual uncertainty depends on previous categorical perceptual choices.

3.2.2 Expected uncertainty

Regarding expected uncertainty, individuals can deviate from an optimal consideration by assuming wrong levels of variability across outcomes. For example, when an animal has experienced one or two bad berries and wrongly assumes minimal expected uncertainty, it might conclude that all berries in the current context are inedible. In line with this idea, studies on the effects of acetylcholine on learning under expected uncertainty provided indirect evidence for human biases regarding this form of uncertainty (Marshall et al., 2016; Moran et al., 2013; Vossel et al., 2014). These results hint that differences in this neurotransmitter's availability can change the individual levels of expected uncertainty during learning. More direct evidence suggests that expected uncertainty is substantially heterogeneous between participants; that is, some participants tend to underestimate outcome variability while others overestimate it (Nassar et al., 2010).

3 Introduction

An additional, not mutually exclusive possibility is that individuals show learning difficulties due to biases in the computation of belief uncertainty that reflects the reliability of the learned belief. In this scenario, the animal might have a realistic expected-uncertainty estimate but compute too much or too little uncertainty over the berries' taste (Figure 3b). One line of work on pupil responses during learning under uncertainty found that belief uncertainty is reflected in baseline pupil diameter and varies substantially between participants (Krishnamurthy et al., 2017; Nassar et al., 2012), suggesting that some individuals might hold more uncertain beliefs despite the experience of similar outcome sequences. However, previous work failed to directly estimate the influence of belief uncertainty on learning based on computational modeling. Estimating the impact of belief uncertainty could more directly reveal subject-specific biases in the computation of this crucial factor.

3.2.3 Unexpected uncertainty

Volatility

Concerning unexpected uncertainty, humans show biases in responding to changes in the environment. In volatility models that assume drifting changes, individual differences in the assumptions about volatility might lead to such biases. For example, an animal that underestimates that the berry taste changes across time might draw wrong conclusions about such changes (Figure 3c). In line with this idea, Behrens et al. (2007) reported a relationship between volatility and ACC activity, suggesting that activity in this area modulated individual differences in the regulation of learning. Moreover, this finding indicates that beliefs about volatility itself are an important determinant of learning differences between people, which could be related to individual differences in levels of norepinephrine (Marshall et al., 2016).

Another possible deviation from an optimal consideration of volatility can be ascribed to approximate algorithms that are simpler than normative inference. Thus, in the example, the animal might not explicitly represent the concept of volatility but

use computationally cheaper solutions that lead to similar behavioral adjustments. Bröker et al. (2018) provided evidence for this perspective. They showed that a simple model in which the regulation of learning depends on adaptive forgetting of past trials sometimes captures human behavior better than the HGF. As described above, the HGF regulates learning according to normative computations by inferring the environmental volatility using hierarchical Bayesian inference. This finding suggests that humans do not necessarily form a mental model of volatility that they update throughout the task, but they might rely on computationally cheaper alternatives to normative computations, potentially at the cost of biases in adapting to volatility.

Changepoints

Changepoint models offer an alternative account to capture biases in responding to changes in the environment with differences in the assumptions about the hazard rate. Nassar et al. (2010) provided evidence for this idea based on the finding that individual learning differences in changing environments are related to subjective hazard rates. In particular, subjects who responded more strongly to changepoints seemed to assume higher hazard rates than subjects who responded more weakly.

Finally, individuals might also differ in their ability to detect changes, referred to as [surprise insensitivity](#). That is, an animal might suffer from an impaired ability to attribute surprising outcomes to a changepoint and instead assume that prediction errors are due to variability between outcomes (Figure 3c). Consistent with this idea, both Nassar et al. (2010) and Payzan-LeNestour and Bossaerts (2011) showed that participants often insufficiently update beliefs after changepoints, which could also be mediated by inter-individual differences in the norepinephrine system (Krishnamurthy et al., 2017; Nassar et al., 2012; Payzan-LeNestour et al., 2013).

3.3 Lifespan differences

Throughout the lifespan, age-related changes of brain mechanisms at a functional, anatomical, and chemical level shape behavior (Lindenberger et al., 2006). Concerning their sensory abilities, children show improvements across development while older adults suffer from a decline (Bradley & Freeman, 1982; Burton et al., 1993; Dully et al., 2018; Ellemberg et al., 1999; Elliott et al., 1990; Leat et al., 2009). Moreover, it is well known that lifespan brain development affects a broad range of higher-level cognitive functions related to learning and decision making (Eppinger et al., 2011; Manning et al., 2020; Mather & Harley, 2016; Nyberg et al., 2010; Raz & Rodrigue, 2006; Toga et al., 2006). Therefore, it should come as no surprise that learning biases and performance impairments are not only present in younger adults but also in children and older adults, and importantly often to a different degree than in younger adults. Previous work only rarely investigated age-related differences in learning under uncertainty with normative computational models (De Boer et al., 2017; Hämmerer et al., 2019). However, as described in the following sections, some results hint that children and older adults show characteristic age-related impairments, potentially related to each form of uncertainty and beyond.

3.3.1 Perceptual uncertainty

There is some evidence to suggest that the processing of perceptual uncertainty in perceptual decision making changes across the lifespan. The few studies that investigated age-related influences of perceptual uncertainty on perceptual decision making using computational models suggest that both children and older adults choose more cautiously, that is, take more time to accumulate perceptual evidence compared to younger adults (Manning et al., 2020; for review focusing on older adults, see Dully et al., 2018). Moreover, some findings suggest a lower speed of evidence accumulation in both age groups, i.e., a lower drift rate (Forstmann et al., 2011; Manning et al., 2020). While these results suggest age-related lifespan differences in

perceptual uncertainty, it remains an open question of how these differences affect learning. One hypothesis is that learning under perceptual uncertainty in both age groups is worse compared to younger adults. In particular, it could be the case that they show stronger biases due to [categorical perception](#), as described above.

3.3.2 Expected uncertainty

Children and older adults suffer from learning impairments in the presence of expected uncertainty compared to younger adults. Many studies applied probabilistic learning tasks similar to the example in Figure 2b in which participants have to learn the reward probabilities of two or more options from uncertain outcomes. A frequent finding is that children (Cohen et al., 2010; Crone, Jennings, et al., 2004; Eppinger et al., 2009; Hämmerer et al., 2011; Van den Bos et al., 2012; Van Duijvenvoorde et al., 2013) and older adults (Chowdhury et al., 2013; Eppinger et al., 2008; Eppinger et al., 2013; Hämmerer et al., 2011; Samanez-Larkin et al., 2010; Samanez-Larkin et al., 2012; Samanez-Larkin et al., 2014) perform particularly worse than younger adults when outcomes are more variable, suggesting that the underlying impairment is related to the presence of expected uncertainty. However, from these studies, the computations explaining these impairments remain unclear. One shortcoming is that many studies relied on descriptive summary statistics and neither of them used a computational model that directly captured the effects of expected uncertainty on learning.

3.3.3 Unexpected uncertainty

Finally, several studies indicated that children and older adults show age-related learning impairments in tasks involving unexpected uncertainty. Most studies applied reversal-learning tasks, as illustrated in Figure 2c, where the [outcome contingency](#) occasionally reverses. These studies showed reduced learning accuracy in children (Van der Schaaf et al., 2011) and older adults (Mell et al., 2005; Mell et al., 2009;

3 Introduction

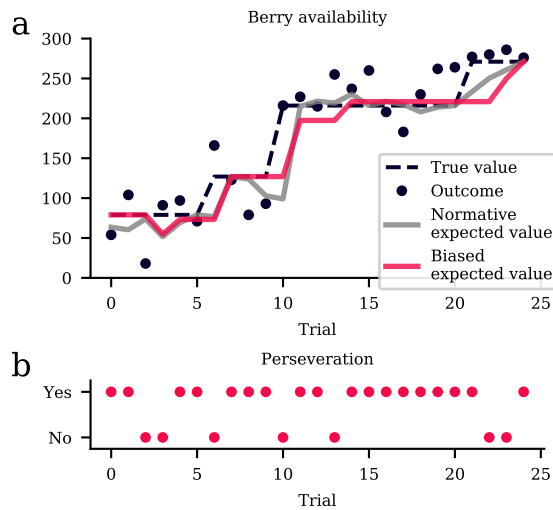


Figure 4. Simplified learning strategy in a changepoint task. In a changing environment, expected values might be reported by relying on a **default belief** about the value. **a)** For example, in an environment where the strawberries' availability changes occasionally, the animal might rely on a default belief about the value during stable periods and primarily update the belief after a surprising outcome that signals a changepoint. **b)** Relying on a default belief yields learning biases due to perseverative behavior. Because this learning strategy requires less flexible belief updating compared to belief updating based on **normative computations** (gray line in (a)), it might be simpler, especially for children and older adults.

Rutledge et al., 2009) compared to younger adults. However, it is an open question if these impairments are directly related to the computation of unexpected uncertainty. One possible explanation for these deficits is that children and older adults systematically differ in their volatility or hazard-rate estimates. However, to the best of my knowledge, such tasks, combined with appropriate computational modeling to capture unexpected uncertainty, were previously not applied in developmental or aging studies.

3.3.4 Simplified learning strategies

In the previous sections, I have collected evidence for potential learning impairments characterized by biases on normative influences of uncertainty (e.g., different hazard rates). However, an additional explanation for age-related learning differences is the reliance on simplified learning strategies that are computationally cheaper and potentially more age-appropriate given limited cognitive capacities. For example, from a normative perspective, an animal should strongly update beliefs in response to changepoints in the berry availability. In contrast, in the absence of a change, it should subtly regulate learning to average out fluctuations in the berry availability due to expected uncertainty. However, a computationally cheaper strategy would

be that the animal relies on a [default belief](#) about the availability of berries until the availability changes dramatically. As shown in [Figure 4](#), according to this simple strategy, the animal would show perseveration on a default response until the occurrence of an obvious changepoint. Thus, this strategy partly ignores expected uncertainty but might be sufficient to adjust to varying levels of berry availability.

As would be expected when children and older adults relied on such defaults more strongly, they often perseverate on previous behavior ([Blackwell et al., 2014](#); [Blackwell & Munakata, 2014](#); [Carroll et al., 2016](#); [Cepeda & Munakata, 2007](#); [Chatham et al., 2009](#); [Crone, Jennings, et al., 2004](#); [Crone, Ridderinkhof, et al., 2004](#); [Head et al., 2009](#); [Munakata et al., 2012](#); [Ridderinkhof et al., 2002](#); [Rutledge et al., 2009](#)). For instance, during reversal learning, older adults often tend to stick to an option (perseveration) despite the occurrence of a reversal in the outcome contingencies ([Ridderinkhof et al., 2002](#); [Rutledge et al., 2009](#)).

However, despite this suggestive evidence, the link between normative computations and simplified learning strategies is currently unclear. In particular, previous work did not examine whether age-related learning differences between children and older adults compared to younger adults emerge from differences in resorting to such simplified strategies.

4 Research questions and hypotheses

The general aim of this dissertation is to advance our understanding of (1) how humans should learn under uncertainty from a normative perspective, (2) which uncertainty biases lead to deviations from normative models, and (3) the lifespan differences in learning under uncertainty. Regarding the aim to study normative influences, I specifically examine the questions:

1. *How should perceptual uncertainty determine learning from a Bayes-optimal perspective?*
2. *Do younger adults regulate learning according to perceptual uncertainty?*

Regarding the aim to study biases in the consideration of uncertainty, I investigate:

3. *Which uncertainty biases do younger adults show during learning under perceptual, expected, and unexpected uncertainty?*

Finally, concerning lifespan differences in learning under uncertainty, I examine the questions:

4. *What is the role of biases while learning under expected and unexpected uncertainty across the lifespan?*
5. *Do children and older adults rely more strongly on simplified learning strategies than younger adults?*
6. *How can computational cognitive models advance our understanding of age-related lifespan differences in learning?*

Rationale and hypotheses

Question 1. *How should perceptual uncertainty determine learning from a Bayes-optimal perspective?* Previous work on perceptual uncertainty used reinforcement-learning models (Lak et al., 2017; Lak et al., 2020, see [Perceptual uncertainty](#)), but it is currently unclear how perceptual uncertainty should control learning from a Bayes-optimal perspective. In [Bruckner et al. \(2020a\)](#), we developed a task combining elements of perceptual and economic decision making where participants were required to learn outcome contingencies under perceptual uncertainty. To study learning from a normative perspective, we developed a Bayes-optimal model that relied on belief states to regulate learning under perceptual uncertainty.

Question 2. *Do younger adults regulate learning according to perceptual uncertainty?* It is currently unclear if humans consider perceptual uncertainty during reward-based learning and decision making. Based on findings in perceptual decision making (Kiani & Shadlen, 2009; Whiteley & Sahani, 2008), we hypothesized that humans would consider perceptual uncertainty for reward-based learning. In [Bruckner et al. \(2020a\)](#), we formalized this hypothesis in our Bayes-optimal model mentioned above and examined human learning behavior based on our integrated perceptual and economic decision-making task.

Question 3. *Which uncertainty biases do younger adults show during learning under perceptual, expected, and unexpected uncertainty?* Biases can affect each form of uncertainty (see [Uncertainty biases](#)). Based on results in perceptual decision making (Fleming et al., 2013; Luu & Stocker, 2018), we hypothesized that categorical-choice biases lead to a lower consideration of perceptual uncertainty during learning in human participants than in a normative Bayesian model ([Bruckner et al., 2020a](#)).

Regarding expected uncertainty, it is unclear to which degree younger adults utilize their belief uncertainty during learning. Previous work suggests that they take belief uncertainty into consideration (Meyniel, Schlunegger, et al., 2015; Nassar

et al., 2010; Payzan-LeNestour & Bossaerts, 2011), but it is possible that they show systematic biases during learning because of an over- or underestimation of belief uncertainty (Krishnamurthy et al., 2017; Nassar et al., 2012). In Nassar et al. (2016) and Bruckner et al. (2020b), we applied a predictive-inference task (McGuire et al., 2014; Nassar et al., 2010) and computational modeling to systematically assess the degree to which biases affect the regulation of learning according to belief uncertainty.

With respect to unexpected uncertainty, we built upon the predictive-inference task mentioned above to examine two potential biases. First, whether younger adults show biases in responding to changepoints due to [surprise insensitivity](#). Based on prior results (Nassar et al., 2010; Payzan-LeNestour & Bossaerts, 2011), we expected to find an under-adjustment to changepoints compared to normative Bayesian models. Second, whether younger adults over- or underestimate the frequency of changepoints in the environment ([hazard rate](#)). In this case, we had no specific hypothesis about the direction of the effect.

Question 4. *What is the role of biases while learning under expected and unexpected uncertainty across the lifespan?* Previous work suggests that children and older adults show characteristic age-related biases in learning under uncertainty (see [Lifespan differences](#)). Concerning expected uncertainty, children and older adults might underestimate belief uncertainty during learning compared to younger adults. Using computational modeling and the predictive-inference task by McGuire et al. (2014), we systematically compared learning based on belief uncertainty between younger and older adults (Nassar et al., 2016) and across the lifespan (Bruckner et al., 2020b).

Moreover, we hypothesized that age-related learning differences could be related to biases in responding to unexpected uncertainty, particularly because of an insensitivity to surprise or a misestimation of the hazard rate leading to too strong or too weak responses to changepoints. In Nassar et al. (2016), we tested these hypotheses in older compared to younger adults and in Bruckner et al. (2020b) across the lifespan.

Question 5. *Do children and older adults rely more strongly on simplified learning strategies than younger adults?* In [Bruckner et al. \(2020b\)](#), our initial aim was to build upon [Nassar et al. \(2016\)](#) to study how learning under uncertainty differs across the lifespan. However, consistent with previous developmental and aging work ([Blackwell et al., 2014](#); [Blackwell & Munakata, 2014](#); [Carroll et al., 2016](#); [Cepeda & Munakata, 2007](#); [Chatham et al., 2009](#); [Crone, Jennings, et al., 2004](#); [Crone, Ridderinkhof, et al., 2004](#); [Head et al., 2009](#); [Munakata et al., 2012](#); [Ridderinkhof et al., 2002](#); [Rutledge et al., 2009](#)), the results showed that children and older adults frequently perseverated on previous predictions, which was initially not explained by our Bayesian model. This finding pointed to the reliance on simplified learning strategies in these age groups, specifically, an over-dependence on **default beliefs** that they only updated after particularly unexpected outcomes. To investigate the role of such simplified learning strategies, we conducted two age-comparative behavioral experiments and developed a novel computational model to formalize age-related differences in the reliance on default beliefs across the lifespan.

Question 6. *How can computational cognitive models advance our understanding of age-related lifespan differences in learning?* The final question of this dissertation combines insights from [Van den Bos et al. \(2018\)](#), which was an opinion article on computational neuroscience across the lifespan, [Nassar et al. \(2016\)](#), and [Bruckner et al. \(2020b\)](#) on age-related learning differences. Based on these studies, I will discuss why computational models are an important tool to better understand age-related differences in learning but also consider some of the pitfalls of applying computational modeling to study lifespan differences.

5 General methodology

In this chapter, I will give a general overview of the methods of the three empirical studies ([Bruckner et al., 2020a](#), [Nassar et al., 2016](#), and [Bruckner et al., 2020b](#)). These include participant samples, experimental tasks, and computational modeling.

5.1 Participants

Study I: Perceptual uncertainty

We conducted the experiment in [Bruckner et al. \(2020a\)](#) at Freie Universität Berlin and analyzed data from 52 younger adults (18 to 33 years).

Study II: Uncertainty and aging

We conducted the experiment in [Nassar et al. \(2016\)](#) at the Max Planck Institute for Human Development Berlin. The effective study sample comprised 57 younger adults (20 to 30 years) and 57 older adults (56 to 80 years).

Study III: Default beliefs in children and older adults

[Bruckner et al. \(2020b\)](#) consists of two experiments. We conducted the first experiment at Technische Universität Dresden. Here, we analyzed data of 33 children (8 to 10 years), 29 adolescents (13 to 17 years), 32 younger adults (20 to 28 years), and 35 older adults (62 to 80 years). We conducted the second experiment at the Max Planck Institute for Human Development Berlin with an effective sample size of 31 children (7 to 11 years), 25 younger adults (20 to 28 years), and 34 older adults (61 to 76 years).

5.2 Experimental tasks

The studies of the dissertation build upon two experimental tasks. In [Bruckner et al. \(2020a\)](#), we developed the Gabor-bandit task, a novel paradigm to study the influence of perceptual and expected uncertainty on reward-guided learning (Figure 5a). In [Nassar et al. \(2016\)](#) and [Bruckner et al. \(2020b\)](#), we used the helicopter task, which was developed by Nassar and colleagues ([McGuire et al., 2014](#); [Nassar et al., 2010](#)) to examine learning under expected and unexpected uncertainty. The aim of this task is to maximize the number of caught bags that are dropped by a helicopter hovering in the sky. Because of this intuitive cover story, the helicopter task is particularly suited for participants across the lifespan (Figure 6a).

5.2.1 Gabor-bandit task

As shown in Figure 5a, the Gabor-bandit task is a combination of a perceptual (Gabor patches) and an economic-choice task (one-armed bandit). The goal of the task is to collect as much reward as possible, which requires reward-guided learning under perceptual and expected uncertainty. The Gabor-bandit task consists of three stages. In the first stage, participants report a perceptual choice about which of the two Gabor patches displays the higher contrast. The second stage requires an economic choice of the fractal with the higher reward probability. In the third stage, the task delivers a probabilistic reward of either zero or one point that depends on the participant's economic choice.

As shown in Figure 5b, the learning problem is to infer the true but unknown probability that governs the [outcome contingency](#) in the task. The task [state](#) ("State") is unobservable for participants and determines both which Gabor patch has a higher contrast and which fractal has a higher reward probability (e.g., state 0: negative contrast differences and blue fractal better; state 1: positive contrast difference and red fractal better). Random noise in the contrast differences induces trial-by-trial differences in the perceptual uncertainty ("Contrast difference"). Based on the

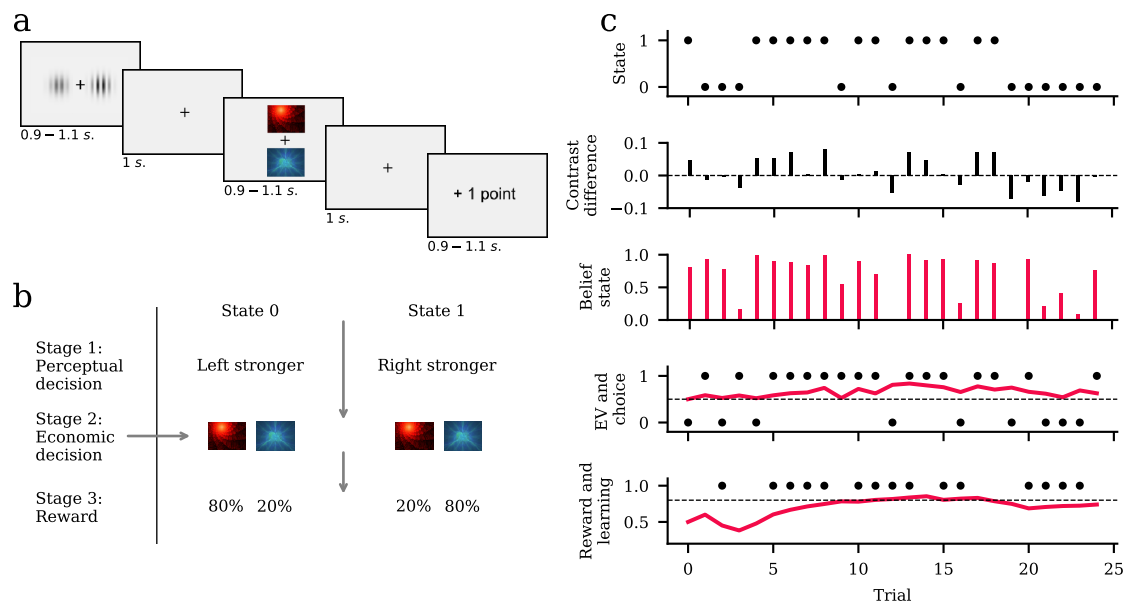


Figure 5. Gabor-bandit task and Bayesian inference under perceptual uncertainty. In the Gabor-bandit task, the aim is to collect as much reward as possible, which requires reward-based learning under perceptual and expected uncertainty. **a)** Task stages: 1) In the first stage, a perceptual choice about the contrast difference between two Gabor patches has to be reported. In this example, the right patch has a stronger contrast. 2) In the second stage, an economic choice about the fractal with the higher reward probability has to be made. The fractals' reward probabilities are initially unknown to the participants and must be learned throughout a block. 3) In the third stage, participants receive a probabilistic reward for their economic choice (either 0 or 1 point). **b)** Task structure. At each trial, the **state** is either 0 or 1. The state determines the **outcome contingency** in a trial, i.e., the position of the high-contrast Gabor patch and the reward probabilities of the fractals. In this example, state 0 leads to a negative contrast difference between the Gabor patches, i.e., the left patch has a higher contrast and the blue fractal a higher reward probability (80%) than the red fractal (20%). In state 1, contrast differences are positive; that is, the right patch has a higher contrast and the red fractal a higher reward probability (80%) than the blue fractal (20%). Note that the participants' perceptual decisions have no consequences for the reward delivery. **c)** Trial-by-trial example and normative computational model that we developed to examine reward-based learning under perceptual uncertainty. "State": Task state governing the outcome contingency. "Contrast difference": The exact contrast differences are drawn from a uniform distribution, which induces varying perceptual uncertainty in the task. For example, in trial 0, the model observes an intermediate positive contrast difference (right patch stronger), while in trials 1 and 2, a weakly negative difference (left patch stronger). "Belief state": Based on the perceived contrast differences, the model computes the belief about the task state of the current trial (**belief state**). Moreover, based on the belief state, the model reports its perceptual choice about the contrast difference (not shown here). "EV (Expected Value) and choice": In the economic-choice stage, the model computes its expected value (red line) about the high-reward fractal as a function of the belief state and the learned outcome contingency, and reports its economic choice (black dots, 0 indicates red choice, 1 indicates blue choice). "Reward and learning": During reward-based learning, the model regulates learning of the outcome contingency (probability of reward if high-reward fractal given state is chosen) according to the belief state and tends to learn more from an outcome when the belief state signals low perceptual uncertainty about the contrast differences (red line).

contrast differences, participants can compute **belief states** reflecting the probability of the current task state ("Belief state"). During the presentation of the fractals, they compute their **expected value** (red line) and indicate their economic choice about the high-reward fractal (black dots) ("EV and choice"). Here, the fractals' reward probabilities lead to expected uncertainty because they only deliver a reward in 80%

5 General methodology

of the correct economic choices (black dots) ("Reward and learning"). Based on these variable rewards, participants are required to learn the contingency between states, choices, and rewards (red line).

5.2.2 Helicopter task

As shown in Figure 6a, the helicopter task is a sequential learning task embedded in a cover story in which a hidden helicopter drops bags that participants try to catch. The goal of the task is to catch as many bags as possible to maximize the collected amount of reward, which requires inferring the helicopter location under expected and unexpected uncertainty. Participants first place a bucket at the location where they predict the next bag to be dropped by the helicopter ("Prediction"). Next, the helicopter drops the bag ("Outcome"), and participants see their [prediction error](#), indicated by a red line between the actual location of the bag and the predicted location ("Prediction error"). Finally, participants can update their buckets to predict the next trial's outcome ("Update").

As shown in Figure 6b, the learning problem is to infer the true but unknown location of the helicopter (first panel). The task takes place under expected uncertainty because the location of the dropped bags varies across trials even if the helicopter does not move, which is introduced as "wind" in the cover story. Unexpected uncertainty is present because the helicopter occasionally changes its location ([change point](#)). Thus, the central problem in the helicopter task is to distinguish variability across the outcomes (expected uncertainty) from change points in the helicopter position (unexpected uncertainty). Because the helicopter is not visible, participants have to consider both sources of uncertainty during learning and more strongly adjust their predictions after change points compared to outcomes that more likely lead to prediction errors because of random variability.

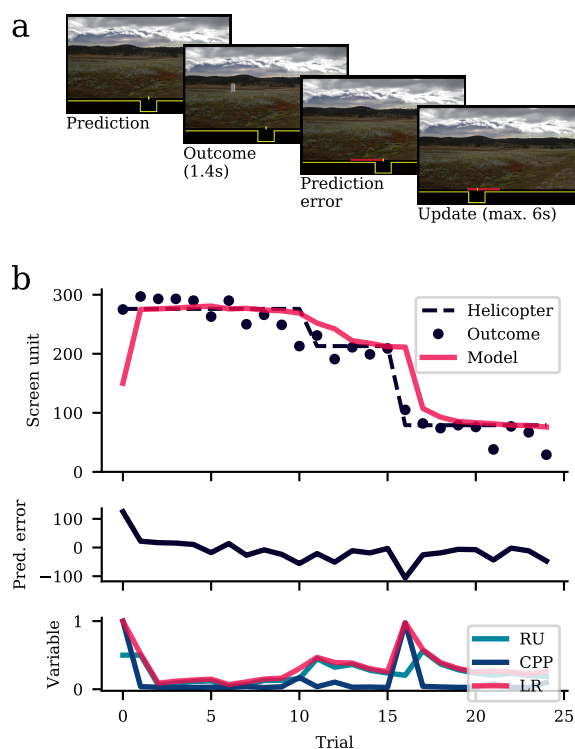


Figure 6. Helicopter task and reduced Bayesian model. a) The aim of the task is to catch as many bags as possible dropped by a hidden helicopter, which requires inferring the helicopter’s location under expected and unexpected uncertainty. Task stages: 1) Participants predict the location of a bag that is dropped by a helicopter hidden behind the clouds. 2) The helicopter drops the bag, which is distorted by "wind" in the environment. 3) After the bag is dropped, a red line indicates the **prediction error** (the difference between the bag location and the prediction). 4) Participants update their prediction of the next bag location. b) Trial-by-trial example and normative model that we used to investigate learning under expected and unexpected uncertainty. First panel: The location of the dropped bags varies across trials (expected uncertainty), and after an occasional changepoint, the helicopter changes its position (unexpected uncertainty). The model learns the position of the hidden helicopter based on the observed outcomes. To do so, the model considers its prediction errors (second panel) and two normative factors—**changepoint probability** (CPP) and relative uncertainty (RU) (third panel). Changepoint probability reflects the probability of a changepoint in the helicopter location. Relative uncertainty reflects how well the model has learned the helicopter location (**belief uncertainty**) relative to the variability in the environment (expected uncertainty). The combination of changepoint probability and relative uncertainty controls the model’s **learning rate** (LR), which determines how strongly a prediction error drives the belief update about the helicopter location.

5.3 Computational modeling

The essential approach to analyzing the data of the three studies is computational modeling. As done throughout the [Introduction](#) of this dissertation, computational models can be used for simulating predictions of human and animal behavior. Moreover, based on the estimation of computational models, one can examine the algorithms underlying cognition and behavior and uncover latent variables that are not directly observable in behavioral data (e.g., expected values, perceptual uncertainty) but are essential elements of such algorithms (Wilson & Collins, 2019).

As outlined in the following sections, in all studies, we started with a normative Bayesian computational model that learned the tasks (near-)optimally ([Normative models](#)). We used these normative models as a theoretical upper-benchmark of how participants would ideally behave in the task. Note that we did not expect participants

5 General methodology

to perform at the same level as these models, but used the normative models to better understand which task factors should modulate learning in the tasks. Based on these insights, we could generate testable hypotheses about how human participants might regulate learning according to these factors. We also added uncertainty [biases](#) to these models based on which we examined if humans systematically deviate from the normative computational models ([Uncertainty biases](#)). To illustrate the different hypothesized learning mechanisms, we first simulated behavioral predictions with the models ([Simulated model hypotheses](#)). We then parameterized the models and estimated several free parameters, allowing us to examine how well a model accounted for the participants' behavior ([Model estimation](#)). To test the reliability of our models for data analysis, we also performed several model validation checks ([Model validation](#)). We additionally compared different computational models to find out which model of our model space explained the empirical data best ([Model comparison](#)). In order to examine if the models accurately described the participants' behavior, we tested whether the best-fitting models could qualitatively reproduce the empirical data ([Post-hoc simulations](#)). Finally, we used a comparative computational approach to investigate age-related learning differences, where we systematically compared the estimated parameters between different age groups ([Computational modeling across the lifespan](#)).

5.3.1 Normative models

Gabor-bandit task

Optimal learning in the Gabor-bandit task requires considering perceptual uncertainty arising from noisy contrast differences between the Gabor patches and expected uncertainty related to variability across outcomes (Figure 5b). Our newly developed Bayes-optimal [agent](#) model considered both forms of uncertainty. Perceptual uncertainty rendered learning more cautiously. Under high perceptual uncertainty, the model learned less from rewards because it could not reliably discriminate the

Gabor patches from each other. In the case of maximal perceptual uncertainty, the model did not learn at all. As a consequence of expected uncertainty, the obtained rewards were noisy, and the model additionally suffered from [belief uncertainty](#) over the true but unknown reward probability. The agent explicitly considered its current belief uncertainty and tended to learn more from reward feedback when its belief uncertainty was high (given the same belief state).

Helicopter task

Optimal learning in the helicopter task relies on considering expected uncertainty due to variability across outcomes and unexpected uncertainty resulting from changepoints (Figure 6b). As shown in the third panel, the model combined its trial-by-trial uncertainty over the helicopter’s actual location (belief uncertainty) with expected uncertainty to compute a factor called relative uncertainty (RU). Moreover, during the observation of an outcome, the model computed [changepoint probability](#) (CPP). To regulate how strongly a prediction error (second panel) influenced the belief update about the hidden helicopter location, the model’s [learning rate](#) (LR) was controlled by a combination of the trial-by-trial relative uncertainty and changepoint probability (third panel).

5.3.2 Uncertainty biases

Gabor-bandit task

In the Gabor-bandit task, we assumed that humans might deviate from the normative model in two principled ways. First, we tested if participants considered perceptual uncertainty to a reduced degree due to categorical-choice biases induced by the perceptual choice about the location of the high-contrast Gabor patch. Second, we considered that humans might not learn according to the principles of Bayesian inference but rather [reinforcement learning](#). Therefore, we additionally tested several reinforcement-learning models.

Helicopter task

In the helicopter task, we considered three potential deviations from the normative model. In [Nassar et al. \(2016\)](#), we showed that differences in uncertainty underestimation, surprise insensitivity, and the assumed hazard rate could lead to dissociable learning impairments between age groups. Uncertainty underestimation specifically arises from low belief uncertainty. Low surprise sensitivity leads to an underestimation of changepoint probability, and differences in hazard rate reflect an under- or overestimation of the prior probability of a changepoint. In [Bruckner et al. \(2020b\)](#), we considered the same factors and additionally examined the role of simplified learning strategies, particularly [default beliefs](#), leading to perseveration and an over-reliance on environmental information.

5.3.3 Simulated model hypotheses

In all studies, we used the computational models to illustrate our hypotheses with simulated predictions. For example, in [Nassar et al. \(2016\)](#), we simulated the effects of age-related differences between younger and older adults in uncertainty underestimation, surprise insensitivity, and hazard rate on learning behavior in the helicopter task. That way, we visualized our hypotheses about the underlying computational mechanisms of age-related learning differences. Similarly, in [Bruckner et al. \(2020a\)](#), we showed that Bayesian-inference and reinforcement-learning models make dissociable predictions about learning under uncertainty.

5.3.4 Model estimation

Model estimation was the core of the computational analyses in all projects. The term refers to the analysis step in which computational models equipped with free parameters are estimated based on the participants' behavioral data. An example of such a freely estimated parameter is uncertainty underestimation ([Nassar et al. \(2016\)](#)). Here, we estimated for each participant to which degree he or she

underestimated belief uncertainty. Another example from [Bruckner et al. \(2020a\)](#) is a parameter that modeled the degree of categorical influences of the past perceptual choice on learning. In summary, model estimation allowed us to test for differences in key parameters of the computational models that provided new insights into the computational mechanisms that guide learning in human participants.

5.3.5 Model validation

An important part of computational analyses of behavior is the validation of several aspects of the computational models, most importantly, to test if the models serve as a reliable data-analysis tool. A critical part of model validation is parameter recovery, testing how reliably the free parameters can be estimated based on simulated experimental data. Another essential step is model recovery, which is especially important if one compares multiple models. This refers to examining whether the models of the model space make dissociable predictions about participants' learning behavior. For example, in [Bruckner et al. \(2020a\)](#), we tested how well we could dissociate the predicted choices of a Bayesian-inference model from the predicted choices of a reinforcement-learning model. Both types of analyses are essential to ensure that the conclusions drawn conditional on model-based analyses are statistically plausible.

5.3.6 Model comparison

During model comparison, one systematically compares the models of the model space concerning how well they capture the participants' behavior. One important aspect of model comparison is how likely the participants' choices are according to a computational model. For example, a random-choice model applied to a two-alternative forced-choice task would assign a choice probability equal to 0.5 to both options, thereby only explain choices at chance level. In contrast, a better model would predict choices above chance, for example, with an average probability of

0.7. In this case, model comparison would favor the second model. However, one issue with model comparison based on choice probabilities is that models with more free parameters tend to explain more variance of the data at the risk of over-fitting. Over-fitting refers to the phenomenon that "in general, a more complex model will fit data better than a simpler model, by capturing noise in the data" (Daw, 2011). To account for this possibility, we used the Bayesian information criterion (BIC), which considered the predicted choice likelihoods and penalized for complexity to avoid over-fitting. Taken together, in the projects of this dissertation, we compared different models to test which models described the data best, in order to draw conclusions about the computational mechanisms potentially underlying participants' learning behavior in the different tasks.

5.3.7 Post-hoc simulations

Moreover, we performed post-hoc simulations to check if our models accurately described the behavioral data. This step refers to simulations based on the freely estimated parameters to test if the estimated models qualitatively reflect the learning and decision-making behavior of the participants (Palminteri et al., 2017). For example, in Bruckner et al. (2020a), we plotted the participants' learning curves in the Gabor-bandit task together with the simulated learning curves. If the actual and simulated learning curves are comparable, post-hoc analyses offer support for a computational model, while substantial deviations between them indicate that a model does not accurately capture learning behavior.

5.3.8 Computational modeling across the lifespan

Finally, computational modeling offers a principled approach to study age-related differences in learning under uncertainty. As illustrated in Figure 7, a comparative computational approach is based on behavioral experiments in a target age group (e.g., older adults) and a comparison group (e.g., younger adults). The experimental

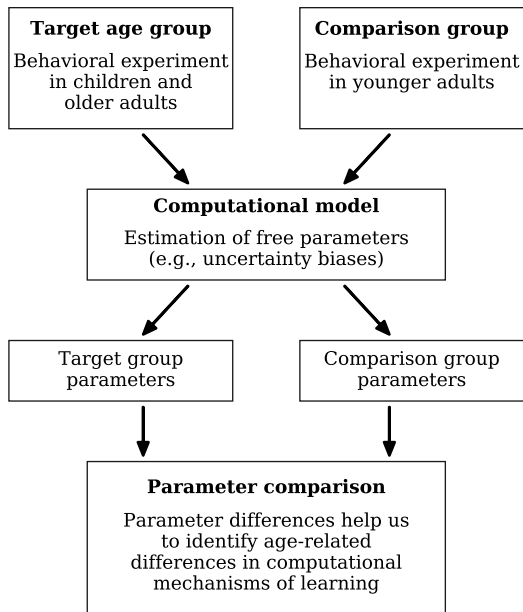


Figure 7. Computational modeling across the lifespan. Age-related differences in learning under uncertainty can be analyzed using a combination of behavioral experiments and computational modeling. A common approach is to conduct behavioral experiments with a target age group (e.g., children, older adults) and a comparison group (e.g., younger adults). The corresponding data are subsequently analyzed using computational cognitive models equipped with free parameters that capture individual differences, for example, concerning uncertainty biases. Parameters of the target group are systematically compared to parameters of the comparison group to examine age-related differences in the computational mechanisms underlying learning. This schematic is based on Maia and Frank (2011).

data of each participant are analyzed based on computational models. This approach particularly affords the estimation of latent participant-specific parameters indicative of individual differences in key factors involved in such computations (e.g., uncertainty biases) that can not directly be extracted from behavioral data (e.g., based on the percentage of correct responses or reaction times). The estimated parameters are subsequently compared across the target and comparison groups, which can help us identify age-related differences in the computational mechanisms of learning.

5.4 Software and data repositories

The experimental data were analyzed in Python (Python Software Foundation) and Matlab (The Mathworks Inc., USA). All scripts and data of Bruckner et al. (2020a) can be found at https://github.com/rasmusbruckner/gaborbandit_analysis. The Gabor-bandit task code is available at https://github.com/rasmusbruckner/gaborbandit_task. All scripts and data of Bruckner et al. (2020b) can be found at https://github.com/rasmusbruckner/adaptivelearning_analysis.

6 Summary of the dissertation studies

In this chapter, I summarize the four studies of the dissertation.

6.1 Study I: Perceptual uncertainty

In [Bruckner et al. \(2020a\)](#), we examined the role of [belief states](#) and categorical-choice biases during reward-guided learning under perceptual uncertainty. Based on previous work in perceptual decision making ([Fleming et al., 2013](#); [Luu & Stocker, 2018](#); [Stocker & Simoncelli, 2007](#)), we hypothesized that normative belief states modulate learning under perceptual uncertainty but that [categorical perception](#) driven by categorical perceptual choices about the current [state](#) of the environment additionally [biases](#) learning in the direction of this decision. We examined this question based on a Bayesian-inference and reinforcement-learning framework.

We first developed a Bayes-optimal agent model to derive the [normative computations](#) that should guide reward-based learning under perceptual uncertainty. This analysis showed that consistent with the intuition that the regulation of learning should be more cautious under perceptual uncertainty, the belief state optimally scales down to which degree an agent learns from reward feedback. In the case of equal belief states indicating maximal perceptual uncertainty, the agent should ignore feedback because it can not meaningfully interpret it. We also applied a reinforcement-learning model that weighted the [learning rate](#) as a function of the belief state. The comparison to the Bayes-optimal model showed that although both models down-scaled the impact of feedback on belief updating according to the belief state, a Bayes-optimal model learns the reward probabilities of the economic-choice options more accurately than a reinforcement-learning model.

Next, we tested to which degree learning under perceptual uncertainty in younger adults is modulated by the belief state and categorical-choice biases. The results showed that both factors modulate economic decision making and learning. We

6 Summary of the dissertation studies

identified that the contribution of each factor was substantially heterogeneous between participants. That is, some participants almost entirely regulated learning according to the belief state, while others almost exclusively relied on their previous categorical perceptual choice during learning.

Finally, we found that a Bayes-optimal agent model that included the categorical-choice bias described the participants' learning data better than a belief-state reinforcement-learning agent with a similar bias. The worse model fit of the reinforcement-learning model was primarily the consequence of an additional free learning-rate parameter, which we penalized in the BIC-based model comparison (see [Model comparison](#)). Both models explained a similar amount of variance in the participants' behavior, and we, therefore, concluded that they similarly described human learning under perceptual uncertainty.

In summary, [Bruckner et al. \(2020a\)](#) provides evidence that human participants consider belief states to regulate learning under perceptual uncertainty. However, learning is categorically biased due to prior perceptual choices, leading to systematic deviations from the normative computations according to optimal Bayesian inference.

6.2 Study II: Uncertainty and aging

In [Nassar et al. \(2016\)](#), we aimed to examine which computational factors account for age-related learning differences between younger and older adults. Previous work paid attention to this question yet either without using a computational-modeling approach or using traditional reinforcement-learning models that did not explicitly capture uncertainty ([Chowdhury et al., 2013](#); [Eppinger et al., 2013](#)). Therefore, we tested whether a learning deficit in older adults is related to the three computational factors [belief uncertainty](#), [surprise insensitivity](#), and [hazard rate](#).

We first used a regression approach to test these three hypotheses. This analysis showed that the idea of belief-uncertainty underestimation in older adults explained our empirical data best. This result was consistent with a model that simulated

6.3 Study III: Default beliefs in children and older adults

uncertainty underestimation in older adults and correctly indicated that older adults would show weaker influences of small prediction errors and relative uncertainty on learning.

Finally, we directly estimated several computational models to dissociate the computational impairments related to belief uncertainty, surprise insensitivity, and hazard rate in more detail. We equipped a Bayesian learning model (Nassar et al., 2010) with multiple free parameters to quantitatively estimate the degree of uncertainty underestimation, surprise insensitivity, and the individual assumption about the hazard rate. Consistent with our regression results, individual parameter estimates showed a more substantial degree of belief-uncertainty underestimation in older adults compared to younger participants.

In summary, Nassar et al. (2016) was the first study comparing learning under uncertainty using normative computational models between younger and older adults. The study provides evidence that older adults underestimate belief uncertainty yielding lower learning flexibility than in younger adults.

6.3 Study III: Default beliefs in children and older adults

In Bruckner et al. (2020b), our initial aim was to build upon the results of Nassar et al. (2016) to examine if children show similar computational learning impairments as older adults. Consistent with our previous results, we found some evidence for uncertainty underestimation in older adults. However, these differences insufficiently explained the performance differences that we identified across the lifespan, particularly less accurate learning in children and older adults than in adolescents and younger adults.

We found that these performance differences were primarily related to age-related differences in perseveration on previous predictions. To account for these differences in our computational analyses, we developed a novel mixture model that captured both the regulation of learning according to principles of Bayesian belief updating and perseveration on previous predictions. Specifically, this model indicated that

6 Summary of the dissertation studies

perseveration in children and older adults was particularly frequent in the range of small and intermediate prediction errors.

These results strongly suggested that we might rescue performance differences between the age groups if perseveration is reduced or eliminated. To test this idea, we conducted a follow-up experiment with children, younger adults, and older adults. We found that a manipulation that randomly replaced the participants' default prediction about the helicopter location eliminated perseveration in all age groups. However, the lower perseveration levels came at the cost of increased performance differences between the age groups, and in particular, lower performance in children and older adults. These performance impairments resulted from a stronger impact of the randomly initialized default predictions on learning in these age groups. We referred to this effect as **environmental control** (Lindenberger & Mayr, 2014) of learning because the externally determined default predictions partly controlled the learning behavior, particularly in children and older adults. Thus, children and older adults showed not only more perseveration (first experiment) but also environmental control of learning (follow-up experiment), depending on the exact task condition.

Therefore, the results of the two experiments indicated an increased reliance on default-belief settings in children and older adults. In the first experiment, the participants' previous predictions internally generated the default belief. In the follow-up experiment, the randomized default predictions generated it externally. Thus, relying on the default belief in an environment where the default emerges from past predictions leads to perseveration, whereas in the presence of random influences on the default, it leads to stronger environmental control of learning.

We formally captured this relationship between perseveration and environmental control in a new computational model that reproduced our empirical results across the two experiments by assuming age-related differences in how strongly default-belief settings influence learning behavior across the lifespan. In particular, this model assumed age-related differences in how quickly participants are satisfied with relying

6.4 Study IV: Computational modeling across the lifespan

on the default belief. Accordingly, beliefs are adjusted away from default values until reaching a criterion value of plausibility. Lower criterion values captured both perseveration and environmental control along with age differences therein, providing the first comprehensive understanding of how developmental groups can rely more on internal and external information under different conditions.

In summary, [Bruckner et al. \(2020b\)](#) shows that perseveration, as well as an over-reliance on externally provided information during learning under uncertainty in children and older adults, can both be explained in terms of insufficient updating from a default belief yielding lower learning performance in these age groups.

6.4 Study IV: Computational modeling across the lifespan

In [Van den Bos et al. \(2018\)](#), we discussed how computational cognitive models could improve our understanding of learning and decision making across the lifespan. Although developmental cognitive neuroscience has generated a considerable amount of data and theories about developmental differences and aging, the field suffers from two explanatory problems — the specificity and identity problems — that hinder progress in understanding the link between (brain-)development and behavior.

First, neuro-cognitive theories are often not specific enough to be translated into behavioral and neuroscientific predictions. One reason for this problem is that previous work almost exclusively formulated those theories verbally. Consequently, in some circumstances, the same data set can support or reject a verbally-formulated theory. Formalized computational models can increase the specificity of developmental theories because they can strongly constrain the predicted effects and, importantly, indicate which effects falsify a theory.

Second, it is often hard to identify the cognitive and neural processes that underlie developmental differences in behavior. One reason for this problem is that many variables required to explain developmental differences are latent, i.e., they are not directly observable, such as learned beliefs about outcome probabilities or

6 Summary of the dissertation studies

expected uncertainty. Computational modeling can increase our ability to identify those underlying processes because they allow us to define and measure latent variables.

We illustrated the potential of computational modeling to deal with these issues using three popular domains of developmental and aging research as an example — risk-taking, strategy selection, and reinforcement learning. Next to demonstrating the advantages of computational models for studying lifespan development, we offered a general overview of how researchers can tackle potential pitfalls. This discussion addressed how one should deal with group differences in how well a model fits the data, potential solutions to parameter estimation based on sparse data using hierarchical estimation approaches, and the proper link of parameter estimates and neuroimaging data.

In summary, in [Van den Bos et al. \(2018\)](#), we argued that the specificity of model-based predictions and the possibility to define and measure latent variables in a computational model substantially improves research on age-related differences in learning and decision making across the lifespan.

7 General discussion and future directions

I will begin this chapter by discussing the six research questions of my dissertation. I will then link perceptual, expected, and unexpected uncertainty to a broader literature in psychology and cognitive neuroscience. Subsequently, I will argue that one can interpret the findings of my studies regarding categorical-choice biases and default beliefs as rational processes in the face of limited cognitive resources. Finally, I will embed the three forms of uncertainty into the cycle of adaptive behavior with which I introduced the dissertation.

7.1 Discussion of the research questions

Question 1. *How should perceptual uncertainty determine learning from a Bayes-optimal perspective?* In [Bruckner et al. \(2020a\)](#), we developed a Bayes-optimal [agent](#) model that performed belief-state inference based on perceptual information ([Daw, 2014](#); [Dayan & Daw, 2008](#)) and subsequently regulated reward-based learning according to the belief state. This optimal model demonstrates that more perceptual uncertainty, i.e., weaker belief states, should lead to more cautious learning behavior by scaling down the impact of new observations on updated beliefs about the [outcome contingency](#). The model highlights that maximal levels of perceptual uncertainty should lead an agent to ignore reward feedback completely. Therefore, when the learner has no access to relevant information about the current [state](#) of the environment, reward feedback is useless and therefore discarded. We showed that these normative computations are dissociable from the computations of belief-state reinforcement-learning agents ([Chrisman, 1992](#)) used in past work on learning under perceptual and other forms of uncertainty about the state of the environment ([Babayan et al., 2018](#); [Lak et al., 2017](#); [Lak et al., 2020](#); [Starkweather et al., 2017](#)). The reinforcement-learning agent scaled-down learning differently according to the belief state and in particular failed to ignore reward feedback when perceptual uncertainty

7 General discussion and future directions

was at the maximum level. These differences resulted in worse learning accuracy compared to the Bayes-optimal agent, suggesting that reinforcement-learning agents integrate belief states and uncertain outcome feedback less efficiently.

Future studies could build upon this result to examine belief-state-based learning with Bayes-optimal models compared to reinforcement-learning approaches. It is possible that, depending on the task, Bayes-optimal models capture learning in humans and other animals better than reinforcement learning, which could lead to different conclusions about the computational and neural underpinnings of belief-state-guided learning.

Question 2. *Do younger adults regulate learning according to perceptual uncertainty?* In [Bruckner et al. \(2020a\)](#), we tested to which degree younger adults follow the normative computations of our optimal agent model to adjust reward-guided learning as a function of perceptual uncertainty. These model-based analyses showed that younger adults regulate learning similarly, although, to a lower degree than dictated by our normative agent model. Therefore, younger adults sub-optimally consider perceptual uncertainty because their learning is biased by categorical choices about the most likely state of the environment (as further discussed below).

The finding that humans take into account perceptual uncertainty agrees with previous work on perceptual decision making in animals ([Lak et al., 2017](#); [Lak et al., 2020](#)). These studies suggest that belief states modulate the activity of dopamine neurons in the striatum, signaling reward prediction errors, and activity in the medial prefrontal cortex, signaling the expected value of a choice option during decision making. Thus, at a neural level, belief states could reduce reward expectations and prediction errors used for learning to take into consideration when perceptual information is uncertain. Similar evidence was provided by [Stolyarova et al. \(2019\)](#) in a reversal-learning task under perceptual uncertainty, where activity in the basolateral amygdala and ACC covaried with perceptual uncertainty levels. Analogously to our study, these authors suggested that this mechanism might improve an animal's

ability to integrate reward feedback more strongly when perceptual information is more reliable.

While the studies mentioned above provide evidence for effects of perceptual uncertainty on prediction errors and expected values, they did not address how the brain computes perceptual uncertainty as such and how these signals ultimately reach higher-order regions such as the striatum and medial PFC (Lak et al., 2017; Lak et al., 2020; Stolyarova et al., 2019). As shown in [Introduction](#), perceptual decision-making studies in humans and animals indicate that sensory, parietal, and frontal areas play a role in representing perceptual uncertainty (Kiani & Shadlen, 2009; Mulder et al., 2012; Summerfield & Koechlin, 2010).

Moreover, a recent theoretical (Ma et al., 2006) and empirical (Van Bergen & Jehee, 2017; Van Bergen et al., 2015) line of research investigates the neural representations of perceptual uncertainty in visual brain areas. Based on fMRI-data from a perceptual decision-making task, these authors successfully predicted perceptual decision-making performance as a function of a decoded perceptual uncertainty measure. Under the assumption that participants internally represent perceptual uncertainty in visual brain areas, this decoding approach explicitly modeled subjective probability distributions over the presented stimuli. Based on this approach, future work could aim at decoding belief-state representations during learning in the Gabor-bandit task to examine the effect of belief states on the striatum and medial PFC.

Question 3. *Which uncertainty biases do younger adults show during learning under perceptual, expected, and unexpected uncertainty?* Regarding perceptual uncertainty, [Bruckner et al. \(2020a\)](#) identified a substantial bias on learning and economic decision making under perceptual uncertainty, driven by a categorical choice about the most likely yet uncertain state of the environment. Both the Bayesian-inference and reinforcement-learning agents showed that younger adults not only take into account perceptual uncertainty. The commitment to a categorical perceptual choice

7 General discussion and future directions

considerably biases learning and economic decision making into the direction of this choice at the cost of lower learning accuracy.

While our results are the first to demonstrate this categorical-choice bias during reward-guided learning, previous work has hinted that this might be the case. In a perceptual choice task, Fleming et al. (2013) demonstrated that the commitment to a categorical interpretation of uncertain sensory information could lead to a substantially lower consideration of uncertainty due to [categorical perception](#). Research suggesting such categorical biases can emerge from previous perceptual choices supplemented these results (Luu & Stocker, 2018; Stocker & Simoncelli, 2007; Urai et al., 2019). From this perspective, our results are the consequence of self-consistent information processing of the brain. This means that categorical choices about the Gabor patches nudged learning from reward feedback into the direction of the chosen interpretation of the stimuli. Consequently, behavior across the perceptual, economic-choice, and learning stage of the Gabor-bandit task was self-consistent at the cost of a sub-optimal consideration of perceptual uncertainty.

In light of the discussion on the neural representations of perceptual uncertainty, future work should also examine the nature of categorical representations. Research in monkeys suggests that such representations might emerge already in visual areas (Nienborg & Cumming, 2009). Moreover, the study by Akrami et al. (2018) could indicate that parietal areas are similarly involved in the implementation of choice-history biases. This study examined neural representations and behavioral effects of the stimulus history on perceptual decision making and provided compelling evidence for a causal involvement of the posterior parietal cortex. Although it remains unclear whether this finding generalizes to effects of the choice history, it hints that parietal areas could play a similar role in our categorical-choice bias on learning.

With respect to expected uncertainty resulting from variability across outcomes, Nassar et al. (2016) and Bruckner et al. (2020b) found evidence for biases in belief uncertainty. Belief uncertainty reflects how well the expected value of a variable, e.g.,

the probability of obtaining a tasty berry or the helicopter location, has been learned. This form of uncertainty emerges from the presence of variability across outcomes through expected uncertainty; that is, when the learner has only experienced a few outcomes, belief uncertainty is higher than after a greater amount of experiences. Importantly, more belief uncertainty should lead to more learning to reduce uncertainty by incorporating new outcome information into the belief. Both [Nassar et al. \(2016\)](#) and [Bruckner et al. \(2020b\)](#) showed that younger adults underestimate belief uncertainty, i.e., they learned less during times of higher belief uncertainty than they should have, compared to a normative Bayesian model. Explicitly modeling an underestimation of belief uncertainty extends previous computational models ([Behrens et al., 2007](#); [Meyniel, Schlunegger, et al., 2015](#); [Payzan-LeNestour & Bossaerts, 2011](#)) because it allowed us to quantify the deviation from normative belief updating. Moreover, our approach extends previous models assuming a dynamic regulation of learning depending on the amount of learning experiences that, unlike our model, did not explicitly link these adjustments to changes in belief uncertainty ([Fischer & Ullsperger, 2013](#); [Krugel et al., 2009](#)).

Finally, concerning unexpected uncertainty that emerges from surprising changes in the environment, [Nassar et al. \(2016\)](#) and [Bruckner et al. \(2020b\)](#) replicated important findings in younger adults. The hazard rate reflecting the assumption about the prior probability of changepoints is crucial for appropriately adjusting to changing environments. Consistent with previous work ([Nassar et al., 2010](#); [Wilson et al., 2010](#)), we found substantial heterogeneity between younger adults and a tendency to overestimate hazard rates. Another critical factor is the actual response to changepoints as such. We estimated individual levels of surprise insensitivity, which showed that younger adults are partly insensitive to surprise, that is, adjust learning less strongly after a changepoint than a normative model ([Nassar et al., 2010](#); [Payzan-LeNestour & Bossaerts, 2011](#)).

Question 4. *What is the role of biases while learning under expected and unexpected uncertainty across the lifespan?* In Nassar et al. (2016) and Bruckner et al. (2020b), we directly examined this question in order to increase our understanding of age-related learning deficits in children and older adults. In particular, we investigated age-related biases in belief uncertainty, surprise insensitivity, and hazard rate. Nassar et al. (2016) indicated that deficits to utilize belief uncertainty in the service of learning provided a better explanation for age-related learning differences between younger and older adults than differences in surprise insensitivity or individual assumptions about the hazard rate. In Bruckner et al. (2020b), we found similar evidence for reduced consideration of uncertainty in older compared to younger adults. However, after accounting for the observed perseveration that was especially evident in children and older adults, these uncertainty differences were considerably smaller. The perseveration parameters of our model in Bruckner et al. (2020b) most likely captured variance that we previously ascribed to uncertainty underestimation and therefore led to lower age-related differences in the parameter estimates than in Nassar et al. (2016).

Importantly, these results are in line with each other but differ in their interpretation of the specific learning impairments in older adults. Both belief-uncertainty underestimation and perseveration capture a lower regulation of learning in the range of small and intermediate prediction errors. However, while underestimating belief uncertainty results in a constant under-adjustment of outcome predictions in this range, perseveration leads to the exact repetition of the previous prediction on a subset of trials, resulting in a similar under-adjustment of predictions on average. Thus, uncertainty underestimation and perseveration are both related to belief updating in the lower range of prediction errors, make similar predictions about learning behavior on average, but differ in their predictions at the trial-by-trial level. Our most recent results in Bruckner et al. (2020b) more strongly support the idea that perseveration underlies the observed performance differences between younger and older adults.

Taken together, [Nassar et al. \(2016\)](#) and [Bruckner et al. \(2020b\)](#) indicate that uncertainty underestimation plays a role in older adults. However, across the lifespan, differences in learning according to normative factors (belief uncertainty, surprise, hazard rate) do not sufficiently explain age-related performance differences. As further discussed below, age-related performance differences emerge from relying on simplified learning strategies, particularly an over-dependence on [default beliefs](#), in children and older adults.

Question 5. *Do children and older adults rely more strongly on simplified learning strategies than younger adults?* In [Bruckner et al. \(2020b\)](#), we found that perseveration and influences of random environmental information on learning ([environmental control](#)) were responsible for a lower learning performance in children and older adults compared to younger adults. This result shows that age-related learning differences are not necessarily related to a biased consideration of expected or unexpected uncertainty but primarily to simplified learning strategies in children and older adults.

In particular, across our two experiments, we demonstrated that both perseveration and environmental control are the consequence of a more substantial reliance on default beliefs during learning in children and older adults compared to younger adults. Relying on default beliefs is a computationally cheap strategy to regulate learning in changing and uncertain environments but leads to lower performance than Bayes-optimal learning according to environmental statistics. Past choices can lead to the establishment of a default belief, yielding perseveration. In contrast, environments that hinder such internally generated default beliefs can externally establish a default belief. Externally generated default beliefs, conversely, lead to environmental control of learning.

7 *General discussion and future directions*

The default-belief model not only explains our findings but also provides a unifying framework for similar lifespan findings across multiple tasks. In both children and older adults, perseveration frequently occurs in tasks such as reversal-learning or task switching (Blackwell et al., 2014; Blackwell & Munakata, 2014; Carroll et al., 2016; Cepeda & Munakata, 2007; Crone, Jennings, et al., 2004; Crone, Ridderinkhof, et al., 2004; Head et al., 2009; Munakata et al., 2012; Ridderinkhof et al., 2002; Rutledge et al., 2009). While prior work primarily linked these deficits to executive control impairments, our model provides a formal and mechanistic interpretation of perseveration in terms of default beliefs. Accordingly, default beliefs could emerge from the repetition of a response (e.g., before a reversal or a task switch) and lead to perseveration when children and elderly participants are satisfied with the accuracy of the default response.

Similarly, earlier work demonstrated environmental control in task-switching, working-memory, and attention tasks in both children and older adults (Alarcón & Bonardi, 2020; Craik & Bialystok, 2006; Lindenberger & Mayr, 2014). One such example is a cognitive-control task requiring flexible responses to different types of stimuli (AX-CPT, see Chatham et al. (2009)). A frequently observed pattern in this task is that children behave reactively, that is, they tend to provide their answers in response to external task cues instead of based on internal task representations that are independent of such external cues (Chatham et al., 2009; Gonthier et al., 2019; Troller-Renfree et al., 2020). These studies also primarily explained those differences between children and younger adults with deficits in executive functions. Yet, a testable prediction of our model is that externally provided cues might generate a default belief and therefore yield stronger environmental control of behavior.

Thus, while existing accounts provide explanations either for perseveration (Gershman, 2020; Miller et al., 2019; Urai et al., 2019) or a more substantial reliance on environmental information in children and younger adults (Chatham et al., 2009; Lindenberger & Mayr, 2014), they have not yet uncovered the relationship between

them. Our default-belief model might provide this missing link because it can explain both of these ubiquitous effects based on a common cause within a single framework. In summary, our study re-characterizes age-related learning differences in terms of default beliefs. It serves as a proof-of-concept for the idea that apparent deficits in learning under uncertainty can result from computationally cheaper strategies that might be more appropriate for children and older adults in the face of lower cognitive resources.

Question 6. *How can computational cognitive models advance our understanding of age-related lifespan differences in learning?* In [Van den Bos et al. \(2018\)](#), we identified two important explanatory problems for developmental cognitive neuroscience that one can address with the application of computational cognitive models. First, the specificity problem indicates that many current theories about development and aging are verbally formulated, leading to a lack of specificity in empirical predictions to test those theories. Second, the identity problem states that the mechanisms underlying age-related differences are often hard to identify without an explicit computational model of the cognitive processes that lead to these differences.

In order to address these problems, the dissertation studies followed a principled computational approach. We first analyzed the [normative computations](#) that an agent should perform to solve a task according to principles of [Bayesian inference](#). We then considered potential [biases](#), that is, deviations from the normative solution, such as uncertainty underestimation or surprise insensitivity. That way, we could derive specific predictions about age-related learning differences that could be falsified based on experimental data. Furthermore, based on the specific, model-based predictions, we addressed the identity problem, particularly by testing which of these predictions provided the best explanation for the behavior of our target age groups. As similarly argued in [Hauser et al. \(2019\)](#), a computational approach thus affords a mechanistic understanding of development and aging that goes beyond analyses building upon descriptive summary statistics such as the percentage of correct responses.

7 General discussion and future directions

Nevertheless, a computational approach does not come without potential pitfalls. The development of our conclusions about belief uncertainty versus default beliefs as the primary contributor to age-related learning limitations highlights the importance of the model space considered in a study (Nassar & Frank, 2016). While we considered quantitative deviations from the normative models in Nassar et al. (2016) (i.e., uncertainty biases), we only tested qualitative deviations in Bruckner et al. (2020b) (i.e., perseveration and environmental control). The conclusions drawn in these studies were thus strongly contingent on the set of specific models included in the analysis, which clearly shows that it is essential to ask whether other models could account for behavioral data even better than the best-fitting model.

A related issue is that children and older adults often use different task strategies compared to younger adults. For example, in a working-memory task comparing younger and older adults longitudinally, older adults tended to use different working-memory updating strategies depending on the exact task demands (Shing et al., 2012). This finding suggests that older adults might switch to simpler strategies when task demands exceed their cognitive capacities. In Bruckner et al. (2020b), we probably encountered a similar situation because the helicopter task was potentially more difficult for children and older adults compared to younger adults. Consequently, both age groups relied more strongly on default beliefs that either led to perseveration or environmental control of belief updating depending on the task condition.

Concerning the application of computational modeling to samples of different ages, this demonstrates that although it is reasonable to use a computational model validated in younger adults for lifespan studies, one should consider the possibility that children or older adults might use different strategies than younger adults. A potential way of accounting for such differences is to include age-related behavioral effects in a computational model documented in the literature, such as perseveration and environmental control *a priori*. That way, one can begin a study of age-related differences based on a computational model that already captures such effects instead

of modifying the model space *a posteriori*. As further illustrated in [Van den Bos et al. \(2018\)](#), there is a branch of computational modeling that studies strategy differences between age groups, which could provide a starting point for tackling this issue in future research.

Another important lesson of the dissertation and as also argued in [Van den Bos et al. \(2018\)](#) is that proper experimental design is at least as important as the model space to dissociate competing theories. In [Bruckner et al. \(2020b\)](#), the results of the first experiment suggested that perseveration was an important factor for understanding age-related learning differences. However, it remained an open question of why it occurred exactly. One possibility was that participants perseverated on the default prediction irrespective of its exact location, potentially to finish the task quickly. However, an alternative possibility was that the acceptance of the default depended on the default's accuracy with respect to the hidden helicopter location. Only the additional experimental manipulations of the default belief in the follow-up experiment made it possible to dissociate these two possibilities. That way, we found evidence for a stronger reliance on default-belief settings in children and older adults, indicating that they only perseverated when default beliefs provided an acceptable level of accuracy.

Finally, lifespan results have implications for computational approaches that primarily concern healthy younger adults. For example, one potential reason why previous work on learning in an uncertain and changing environment has not systematically examined perseveration is that younger adults show this effect considerably less frequently than children and older adults. Nevertheless, our results clearly showed perseveration in younger adults too, and by adjusting the model primarily to understand development and aging better, our default-belief model offered a new perspective on habitual and goal-directed behavior ([Miller et al., 2019](#); [Wood & R nger, 2016](#)) that provides a mechanism to explain within-trial dynamics of both default beliefs and goal-directed influences on behavior.

After discussing the research questions, I now turn to a more general discussion of perceptual, expected, and unexpected uncertainty in light of related topics such as confidence and risk, as well as potential questions for future research.

7.2 Perceptual uncertainty

An obvious question concerning perceptual uncertainty and, in particular, our formalization in terms of the belief state is its relationship to confidence. Many approaches define confidence as a subjective probability of a choice being correct (Pouget et al., 2016), a definition that stresses that confidence is a particular form of uncertainty about a decision. In our work, the belief state (at the time of the perceptual choice) is computed independently of the choice itself. Our definition of the belief state is, therefore, more in line with a confidence definition in terms of Bayesian probability independent of decision making (Meyniel, Sigman, et al., 2015). Importantly, however, many previous studies found that subjective confidence reports reflect not only stimulus probabilities that directly follow from Bayesian inference, but also higher-order factors such as error monitoring, past choices, attentional processes, and seemingly irrelevant stimulus properties (Fleming & Daw, 2017; Fleming & Lau, 2014; Frömer et al., 2020; Yeung & Summerfield, 2012). The belief state of our Bayes-optimal agent is agnostic to such higher-order influences, but, likely, they were also at play in our study.

In particular, our finding that younger adults did not only consider the optimal belief state during reward-guided learning but also took into account their perceptual choice suggests that a choice-congruent sense of confidence instead of the optimal belief state guides learning under perceptual uncertainty (Peters et al., 2017). This perspective stresses the need to build upon our work by explicitly asking participants to report their confidence levels. A related extension could compare confidence reports in our current task version to a version without an explicit categorical perceptual choice. The absence of a categorical commitment in such a version could

eliminate or reduce the impact of categorical biases on learning through [categorical perception](#).

Another interesting link to previous work is that perceptual uncertainty requires inferences about hidden states of the environment to deal with limited and ambiguous information. Previous work not only examined the role of belief states in learning under perceptual but also state uncertainty due to uncertain temporal or reward-magnitude information (Babayán et al., 2018; Starkweather et al., 2017). Both of these studies suggest that belief states modulate dopaminergic prediction-error signals to regulate learning in the face of state uncertainty. More generally, theoretical (Wilson et al., 2014) and empirical (Nassar, McGuire, et al., 2019; Schuck et al., 2016) results provide evidence that learning significantly depends on the ability to infer hidden environmental states.

Therefore, it could be relevant to study if learning under perceptual uncertainty can be interpreted as a particular form of state inference that is comparable across tasks requiring hidden state inference. Conversely, our results on categorical-choice biases suggest that humans and animals exhibit categorical-choice biases previously ignored in the literature.

Finally, an exciting avenue for future research is to investigate learning under perceptual uncertainty in children and older adults. As argued in [Lifespan differences](#), children and older adults likely show stronger signatures of categorical learning under perceptual uncertainty than younger adults. This hypothesis is plausible because previous research in these age groups showed age-related differences in decision making under perceptual uncertainty (Dully et al., 2018; Manning et al., 2020), which, however, did not further examine the relationship to learning processes as in [Bruckner et al. \(2020a\)](#).

7.3 Expected uncertainty

A crucial difference between expected uncertainty in the Gabor-bandit and helicopter task is that expected uncertainty is directly related to the parameter governing the [outcome contingency](#) in the first case (state-dependent fractal reward probability) but not in the second (helicopter location). In the Gabor-bandit task, the reward probability of the fractals directly relates to expected uncertainty. If the reward probability is equal to 0.8, then the expected probability of not receiving a reward is $1 - 0.8 = 0.2$. This simple example shows that a model that accurately learns the outcome contingency directly knows the level of expected uncertainty in the Gabor-bandit task. In contrast, in the helicopter task, the helicopter location followed from the mean of a Gaussian distribution, and the variance of this distribution governed the variability across outcomes.

Therefore, a more sophisticated computational model of the helicopter task also learns the magnitude of unexpected uncertainty (variance of the outcome-generating distribution) next to the location of the helicopter itself (mean of the distribution). To address this problem, we tested several models in [Nassar et al. \(2016\)](#) that assumed incomplete knowledge of the amount of expected uncertainty. These models quantitatively described the data better than models with fixed expected-uncertainty parameters indicating that the participants' representations of expected uncertainty deviated from our assumed values. However, the better fit of these models came at the cost of worse parameter recoverability (see [Model validation](#)), and the main conclusions were still in line with the simpler models with better recoverability. Relying on simpler models assuming accurate knowledge of expected uncertainty was thus a reasonable compromise between model fit and parameter recoverability. Nevertheless, future work should develop more reliable models that allow us to examine age-related differences in representations of expected uncertainty.

Finally, expected uncertainty is related to the economic concept of *risk*, referring to the known probability of outcomes (Dayan, [2012](#); Mohr et al., [2010](#)). Both

concepts share the idea that one can not predict outcomes corrupted by variability perfectly. However, it might be more sensible to use the term expected uncertainty when studying learning problems. Primarily because in learning research, the exact representation of outcome variability in the environment often has to be learned across time, whereas economic risk is assumed to be known.

7.4 Unexpected uncertainty

With respect to unexpected uncertainty, an interesting follow-up question concerns the representation of environmental changes in humans and animals. This dissertation has shown that volatility and changepoint models differ in their assumptions about whether changes occur slowly and continuously (volatility models) or abruptly at discrete time points (changepoint models). There is evidence that both types of models can be dissociated based on simulated experimental data (Marković & Kiebel, 2016). Therefore, an important direction for future research is to systematically assess if empirical, experimental data in humans generally favor one type of model. One possibility is that the response to changes depends on the exact environment such that in some circumstances (e.g., small and/or slow changes in the presence of high outcome variability), individuals behave more like volatility models but in other circumstances (e.g., large changes in the presence of low outcome variability) more similar to changepoint models.

Finally, one might criticize the term unexpected uncertainty as such. In particular, one potential concern is that unexpected uncertainty is not unexpected as soon as an agent expects changes in the environment (e.g., assuming a hazard rate of changes). However, in this scenario, a change is "more" unexpected than small fluctuations in the outcomes due to variability. Therefore, unexpected uncertainty primarily emerges when [prediction errors](#) are larger than expected conditional on the amount of expected uncertainty in the environment (Yu & Dayan, 2005).

7.5 Resource rationality

Across the two learning tasks used in the dissertation, the results highlight that past choices and environmental cues strongly influence human behavior (categorical-choice bias, default beliefs), although, from the perspective of our normative belief-updating models, they should not play a considerable role (Bruckner et al., 2020a; Bruckner et al., 2020b). As a consequence of limited cognitive resources, both the effects of categorical perceptual choices and default beliefs could be similar manifestations of a general tendency of the brain to reduce the cognitive resources that are required to solve a learning and decision-making problem.

In Bruckner et al. (2020a), participants who relied more strongly on their perceptual choices during learning potentially experienced considerably more difficulties in holding the relevant sensory information about the Gabor patches in working memory. Therefore, they partly replaced this representation with the less complex representation of the perceptual choice as such (Luu & Stocker, 2018). Our Bayes-optimal agent did not suffer from similar difficulties to maintain information in working memory, but it would be possible that it shows similar biases on learning under limited working-memory resources (Qiu et al., 2020).

In line with the idea that normative computations require higher resources than simplified learning strategies, the default-belief model that we proposed in Bruckner et al. (2020b) regarded normative belief updating as a costly deviation from relying on the default belief. Maybe children and older adults relied more strongly on their defaults to adapt their behavioral strategies to their lower cognitive resources.

Therefore, another future direction is to directly assume limited cognitive resources in the computational cognitive models used in the dissertation projects. An appealing approach could be to formulate resource-rational models that, given their limited resources (e.g., working-memory capacity), demonstrate how a normative agent would behave in a task (Gershman, 2020; Griffiths et al., 2015; Lieder et al., 2018). Especially from a lifespan perspective, this approach could help us to

understand learning differences in the face of resource limitations in children and older adults better. In fact, studying lifespan development in terms of resource rationality could highlight that age-related differences reflect not only behavioral impairments but also reasonable adaptations to changes in cognitive capacities.

7.6 Uncertainty in the cycle of adaptive behavior

I started this dissertation by showing how learning occurs during interaction with the environment (Rangel et al., 2008) but stressed that different forms of uncertainty often plague adaptive behavior during this interaction. Now I conclude by explicitly connecting the cycle of adaptive behavior with uncertainty based on the dissertation studies and the literature discussed throughout the dissertation (Figure 8).

An agent needs a *representation* of the set of possible actions and the environmental and internal *states*. In this step, the dissertation highlights the importance of *perceptual uncertainty* about environmental states. Extending the framework by Rangel et al. (2008), we found that the amount of perceptual uncertainty is a critical factor that directly influences learning. Another important topic of the studies was *expected uncertainty*, emerging from variability across outcomes. Because the learner is assumed to expect and tolerate the presence of outcome variability (Yu & Dayan, 2005), I propose to interpret expected uncertainty as a part of the agent's set of representations about the environment. Similarly, subjective representations of the *hazard rate* and *volatility* determine how strongly an agent responds to changes in the environment and are, therefore, part of the representation of the environment.

The *valuation* step refers to the assignment of *expected values* to actions at the time of a choice (Rangel et al., 2008). The dissertation indicates the importance of *belief uncertainty* about the learned expected values (Bruckner et al., 2020a, Nassar et al., 2016, and Bruckner et al., 2020b). Belief uncertainty emerges from variability across outcomes through expected uncertainty in the environment and should, from a normative perspective, determine how strongly a new outcome affects

7 General discussion and future directions

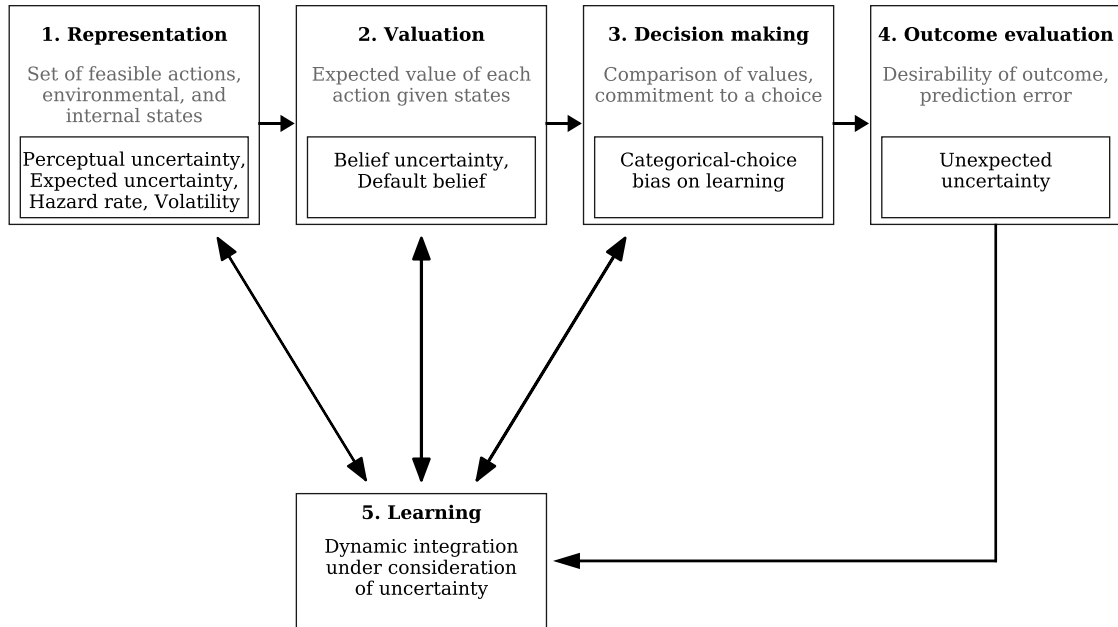


Figure 8. Uncertainty in the cycle of adaptive behavior. In the first step, *representations of states* required for adaptive behavior are often uncertain. This dissertation has particularly focused on perceptual uncertainty about environmental states and shows that weak and unreliable perceptual information about the environment is often plagued by uncertainty, which directly controls the regulation of learning in response to new outcomes. Moreover, adaptive behavior in uncertain environments requires representations of the variability of outcomes (*expected uncertainty*) and the degree of change in the environment (*hazard rate, volatility*). Second, in the *valuation* step, humans and animals have *belief uncertainty* about their *expected values* as a consequence of expected uncertainty. From a normative computational perspective, belief uncertainty should regulate how strongly new outcomes influence an update of the expected value. Moreover, next to learned expected values, *default beliefs* exert strong influences on human behavior across the lifespan, which, depending on the environment, either yield perseverative or environmentally-controlled behavior. Third, *decision making* is not only based on these expected values and default beliefs but prior perceptual choices themselves influence learning through categorical-choice biases. Fourth, during *outcome evaluation*, humans assess whether the environment has changed or not (*unexpected uncertainty*), next to whether an outcome is desirable or not. Taken together, the dissertation highlights that multiple forms of uncertainty dynamically evolving across the various steps of the cycle of adaptive behavior shape *learning*.

an agent’s update of the expected value. Furthermore, in Bruckner et al. (2020b), we demonstrated the importance of *default beliefs* during learning. The study indicates that default beliefs can be generated internally through a commitment to a previous choice, or externally, through environmental information. Adaptive behavior thus not only depends on learned expected values but also on default beliefs independent of experienced outcomes.

Decision making refers to the step in which the learner selects an action based on the expected values of the choice options. An essential insight of Bruckner et al. (2020a) is that previous perceptual choices directly influence reward-based learning under perceptual uncertainty as a consequence of *categorical perception*, showing that decision making not only depends on previously learned expected values, but past choices themselves exert clear influences on learning.

During *outcome evaluation*, an agent evaluates the desirability of a received outcome, which subsequently guides learning. The dissertation primarily contributes to a better understanding of outcome evaluation concerning *unexpected uncertainty* arising during the evaluation of a surprising outcome. We showed that unexpected uncertainty strongly modulates the amount of learning to adjust behavior to changes in the environment (Nassar et al., 2016; Bruckner et al., 2020b).

Finally, *learning* is not merely driven by the evaluation of the latest outcome but also the result of integrating prior representations, expected values, and choices. The dissertation highlights the multiple forms of uncertainty that determine how much humans learn from an outcome and indicates that learning behavior changes across the lifespan substantially, mainly because of age-related *biases* regarding the consideration of uncertainty.

7.7 Conclusion

To conclude, this dissertation offers a theoretical framework of the normative computations that should determine learning under several dissociable forms of uncertainty. The studies show that people consider uncertainty during learning but are affected by biases amplified at both ends of the lifespan.

8 Bibliography

- Akrami, A., Kopec, C. D., Diamond, M. E., & Brody, C. D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature*, *554*(7692), 368–372. <https://doi.org/10.1038/nature25510>
- Alarcón, D. E., & Bonardi, C. (2020). Under the influence of the environment: Children’s responding invigorated and biased by predictive cues. *Journal of Experimental Child Psychology*, *191*, 104741. <https://doi.org/10.1016/j.jecp.2019.104741>
- Babayan, B. M., Uchida, N., & Gershman, S. J. (2018). Belief state representation in the dopamine system. *Nature Communications*, *9*(1), 1891. <https://doi.org/10.1038/s41467-018-04397-0>
- Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don’t know: A neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, *13*(8), 572–586. <https://doi.org/10.1038/nrn3289>
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221. <https://doi.org/10.1038/nn1954>
- Blackwell, K. A., Chatham, C. H., Wiseheart, M., & Munakata, Y. (2014). A developmental window into trade-offs in executive function: The case of task switching versus response inhibition in 6-year-olds. *Neuropsychologia*, *62*, 356–364. <https://doi.org/10.1016/j.neuropsychologia.2014.04.016>
- Blackwell, K. A., & Munakata, Y. (2014). Costs and benefits linked to developments in cognitive control. *Developmental Science*, *17*(2), 203–211. <https://doi.org/10.1111/desc.12113>
- Bradley, A., & Freeman, R. D. (1982). Contrast sensitivity in children. *Vision Research*, *22*(8), 953–959. [https://doi.org/10.1016/0042-6989\(82\)90031-1](https://doi.org/10.1016/0042-6989(82)90031-1)
- Bröker, F., Marshall, L., Bestmann, S., & Dayan, P. (2018). Forget-me-some: General versus special purpose models in a hierarchical probabilistic task. *PloS ONE*, *13*(10), e0205974. <https://doi.org/10.1371/journal.pone.0205974>
- Bruckner, R., Heekeren, H. R., & Ostwald, D. (2020). Belief states and categorical-choice biases determine reward-based learning under perceptual uncertainty. *bioRxiv*. <https://doi.org/10.1101/2020.09.18.303495>
- Bruckner, R., Nassar, M. R., Li, S.-C., & Eppinger, B. (2020). Default beliefs guide learning under uncertainty in children and older adults. *PsyArXiv*. <https://doi.org/10.31234/osf.io/nh9bq>
- Burton, K. B., Owsley, C., & Sloane, M. E. (1993). Aging and neural spatial contrast sensitivity: Photopic vision. *Vision Research*, *33*(7), 939–946. [https://doi.org/10.1016/0042-6989\(93\)90077-A](https://doi.org/10.1016/0042-6989(93)90077-A)

- Busemeyer, J. R., & Johnson, J. G. (2004). Computational models of decision making (D. J. Koehler & N. Harvey, Eds.). In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making*. Blackwell Publishing. <https://doi.org/10.1002/9780470752937.ch7>
- Carroll, D. J., Blakey, E., & FitzGibbon, L. (2016). Cognitive flexibility in young children: Beyond perseveration. *Child Development Perspectives*, 10(4), 211–215. <https://doi.org/10.1111/cdep.12192>
- Cepeda, N. J., & Munakata, Y. (2007). Why do children perseverate when they seem to know better: Graded working memory, or directed inhibition? *Psychonomic Bulletin & Review*, 14(6), 1058–1065. <https://doi.org/10.3758/BF03193091>
- Chatham, C. H., Frank, M. J., & Munakata, Y. (2009). Pupillometric and behavioral markers of a developmental shift in the temporal dynamics of cognitive control. *Proceedings of the National Academy of Sciences*, 106(14), 5529–5533. <https://doi.org/10.1073/pnas.0810002106>
- Chowdhury, R., Guitart-Masip, M., Lambert, C., Dayan, P., Huys, Q. J. M., Düzel, E., & Dolan, R. J. (2013). Dopamine restores reward prediction errors in old age. *Nature Neuroscience*, 16(5), 648–653. <https://doi.org/10.1038/nn.3364>
- Chrisman, L. (1992). Reinforcement learning with perceptual aliasing: The perceptual distinctions approach. *AAAI-92 Proceedings*, 183–188.
- Cohen, J. R., Asarnow, R. F., Sabb, F. W., Bilder, R. M., Bookheimer, S. Y., Knowlton, B. J., & Poldrack, R. A. (2010). A unique adolescent response to reward prediction errors. *Nature Neuroscience*, 13(6), 669–671. <https://doi.org/10.1038/nn.2558>
- Craik, F. I. M., & Bialystok, E. (2006). Cognition through the lifespan: Mechanisms of change. *Trends in Cognitive Sciences*, 10(3), 131–138. <https://doi.org/10.1016/j.tics.2006.01.007>
- Crone, E. A., Jennings, J. R., & van der Molen, M. W. (2004). Developmental change in feedback processing as reflected by phasic heart rate changes. *Developmental Psychology*, 40(6), 1228–1238. <https://doi.org/10.1037/0012-1649.40.6.1228>
- Crone, E. A., Ridderinkhof, R. K., Worm, M., Somsen, R. J. M., & van der Molen, M. W. (2004). Switching between spatial stimulus–response mappings: A developmental study of cognitive flexibility. *Developmental Science*, 7(4), 443–455. <https://doi.org/10.1111/j.1467-7687.2004.00365.x>
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models (M. R. Delgado, E. A. Phelps, & T. W. Robbins, Eds.). In M. R. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Decision making, affect, and learning: Attention and performance xviii*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199600434.003.0001>

- Daw, N. D. (2014). Advanced reinforcement learning (P. W. Glimcher & E. Fehr, Eds.; 2nd edition). In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: Decision making and the brain* (2nd edition). Academic Press. <https://doi.org/10.1016/B978-0-12-416008-8.00016-4>
- Daw, N. D., & Tobler, P. N. (2014). Value learning through reinforcement: The basics of dopamine and reinforcement learning (P. W. Glimcher & E. Fehr, Eds.; 2nd edition). In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: Decision making and the brain* (2nd edition). Academic Press. <https://doi.org/10.1016/B978-0-12-416008-8.00015-2>
- Dayan, P. (2012). Twenty-five lessons from computational neuromodulation. *Neuron*, 76(1), 240–256. <https://doi.org/10.1016/j.neuron.2012.09.027>
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 429–453. <https://doi.org/10.3758/CABN.8.4.429>
- De Boer, L., Axelsson, J., Riklund, K., Nyberg, L., Dayan, P., Bäckman, L., & Guitart-Masip, M. (2017). Attenuation of dopamine-modulated prefrontal value signals underlies probabilistic reward learning deficits in old age. *eLife*, 6, e26424. <https://doi.org/10.7554/eLife.26424.001>
- Diederich, A., & Busemeyer, J. R. (2006). Modeling the effects of payoff on response bias in a perceptual discrimination task: Bound-change, drift-rate-change, or two-stage-processing hypothesis. *Perception & Psychophysics*, 68(2), 194–207. <https://doi.org/10.3758/BF03193669>
- Dully, J., McGovern, D. P., & O’Connell, R. G. (2018). The impact of natural aging on computational and neural indices of perceptual decision making: A review. *Behavioural Brain Research*, 355, 48–55. <https://doi.org/10.1016/j.bbr.2018.02.001>
- Elleberg, D., Lewis, T. L., Liu, C. H., & Maurer, D. (1999). Development of spatial and temporal vision during childhood. *Vision Research*, 39(14), 2325–2333. [https://doi.org/10.1016/S0042-6989\(98\)00280-6](https://doi.org/10.1016/S0042-6989(98)00280-6)
- Elliott, D., Whitaker, D., & MacVeigh, D. (1990). Neural contribution to spatiotemporal contrast sensitivity decline in healthy ageing eyes. *Vision Research*, 30(4), 541–547. [https://doi.org/10.1016/0042-6989\(90\)90066-T](https://doi.org/10.1016/0042-6989(90)90066-T)
- Eppinger, B., Hämmerer, D., & Li, S. C. (2011). Neuromodulation of reward-based learning and decision making in human aging. *Annals of the New York Academy of Sciences*, 1235, 1–17. <https://doi.org/10.1111/j.1749-6632.2011.06230.x>

- Eppinger, B., Kray, J., Mock, B., & Mecklinger, A. (2008). Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia*, *46*(2), 521–539. <https://doi.org/10.1016/j.neuropsychologia.2007.09.001>
- Eppinger, B., Mock, B., & Kray, J. (2009). Developmental differences in learning and error processing: Evidence from ERPs. *Psychophysiology*, *46*(5), 1043–1053. <https://doi.org/10.1111/j.1469-8986.2009.00838.x>
- Eppinger, B., Schuck, N. W., Nystrom, L. E., & Cohen, J. D. (2013). Reduced striatal responses to reward prediction errors in older compared with younger adults. *Journal of Neuroscience*, *33*(24), 9905–9912. <https://doi.org/10.1523/JNEUROSCI.2942-12.2013>
- Feng, S., Holmes, P., Rorie, A., & Newsome, W. T. (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Computational Biology*, *5*(2), e1000284. <https://doi.org/10.1371/journal.pcbi.1000284>
- Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron*, *79*(6), 1243–1255. <https://doi.org/10.1016/j.neuron.2013.07.006>
- Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-Making: A general Bayesian framework for metacognitive computation. *Psychological Review*, *124*(1), 91–114. <https://doi.org/10.1037/rev0000045>
- Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*, 443. <https://doi.org/10.3389/fnhum.2014.00443>
- Fleming, S. M., Maloney, L. T., & Daw, N. D. (2013). The irrationality of categorical perception. *Journal of Neuroscience*, *33*(49), 19060–19070. <https://doi.org/10.1523/JNEUROSCI.1263-13.2013>
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, *67*, 641–666. <https://doi.org/10.1146/annurev-psych-122414-033645>
- Forstmann, B. U., Tittgemeyer, M., Wagenmakers, E.-J., Derrfuss, J., Imperati, D., & Brown, S. (2011). The speed-accuracy tradeoff in the elderly brain: A structural model-based approach. *Journal of Neuroscience*, *31*(47), 17242–17249. <https://doi.org/10.1523/JNEUROSCI.0309-11.2011>
- Frömer, R., Nassar, M. R., Bruckner, R., Stürmer, B., Sommer, W., & Yeung, N. (2020). I knew that! Response-based outcome predictions and confidence regulate feedback processing and learning. *bioRxiv*. <https://doi.org/10.1101/442822>

- Gershman, S. J. (2020). Origin of perseveration in the trade-off between reward and complexity. *Cognition*, 104394. <https://doi.org/10.1016/j.cognition.2020.104394>
- Gold, J. I., & Heekeren, H. R. (2014). Neural mechanisms for perceptual decision making (P. W. Glimcher & E. Fehr, Eds.; 2nd edition). In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: Decision making and the brain* (2nd edition). Academic Press. <https://doi.org/10.1016/B978-0-12-416008-8.00019-X>
- Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5(1), 10–16. [https://doi.org/10.1016/S1364-6613\(00\)01567-9](https://doi.org/10.1016/S1364-6613(00)01567-9)
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>
- Gold, J. I., & Stocker, A. A. (2017). Visual decision-making in an uncertain and dynamic world, 3, 227–250. <https://doi.org/10.1146/annurev-vision-111815-114511>
- Gonthier, C., Zira, M., Colé, P., & Blaye, A. (2019). Evidencing the developmental shift from reactive to proactive control in early childhood and its relationship to working memory. *Journal of Experimental Child Psychology*, 177, 1–16. <https://doi.org/10.1016/j.jecp.2018.07.001>
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229. <https://doi.org/10.1111/tops.12142>
- Hämmerer, D., Li, S.-C., Müller, V., & Lindenberger, U. (2011). Life span differences in electrophysiological correlates of monitoring gains and losses during probabilistic reinforcement learning. *Journal of Cognitive Neuroscience*, 23(3), 579–592. <https://doi.org/10.1162/jocn.2010.21475>
- Hämmerer, D., Schwartenbeck, P., Gallagher, M., FitzGerald, T. H. B., Düzel, E., & Dolan, R. J. (2019). Older adults fail to form stable task representations during model-based reversal inference. *Neurobiology of Aging*, 74, 90–100. <https://doi.org/10.1016/j.neurobiolaging.2018.10.009>
- Hangya, B., Ranade, S. P., Lorenc, M., & Kepecs, A. (2015). Central cholinergic neurons are rapidly recruited by reinforcement feedback. *Cell*, 162(5), 1155–1168. <https://doi.org/10.1016/j.cell.2015.07.057>
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95(2), 245–258. <https://doi.org/10.1016/j.neuron.2017.06.011>

- Hauser, T. U., Will, G.-J., Dubois, M., & Dolan, R. J. (2019). Annual Research Review: Developmental computational psychiatry. *Journal of Child Psychology and Psychiatry*, *60*(4), 412–426. <https://doi.org/10.1111/jcpp.12964>
- Head, D., Kennedy, K. M., Rodrigue, K. M., & Raz, N. (2009). Age differences in perseveration: Cognitive and neuroanatomical mediators of performance on the Wisconsin Card Sorting Test. *Neuropsychologia*, *47*(4), 1200–1203. <https://doi.org/10.1016/j.neuropsychologia.2009.01.003>
- Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, *9*(6), 467–479. <https://doi.org/10.1038/nrn2374>
- Iglesias, S., Mathys, C. D., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E. M., & Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, *80*(2), 519–530. <https://doi.org/10.1016/j.neuron.2013.09.009>
- Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2016). Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron*, *89*(1), 221–234. <https://doi.org/10.1016/j.neuron.2015.11.028>
- Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, *324*(5928), 759–764. <https://doi.org/10.1126/science.1169405>
- Krishnamurthy, K., Nassar, M. R., Sarode, S., & Gold, J. I. (2017). Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nature Human Behaviour*, *1*(6), 0107. <https://doi.org/10.1038/s41562-017-0107>
- Krugel, L. K., Biele, G., Mohr, P. N., Li, S.-C., & Heekeren, H. R. (2009). Genetic variation in dopaminergic neuromodulation influences the ability to rapidly and flexibly adapt decisions. *Proceedings of the National Academy of Sciences*, *106*(42), 17951–17956. <https://doi.org/10.1073/pnas.0905191106>
- Lak, A., Nomoto, K., Keramati, M., Sakagami, M., & Kepecs, A. (2017). Midbrain dopamine neurons signal belief in choice accuracy during a perceptual decision. *Current Biology*, *27*(6), 821–832. <https://doi.org/10.1016/j.cub.2017.02.026>
- Lak, A., Okun, M., Moss, M. M., Gurnani, H., Farrell, K., Wells, M. J., Reddy, C. B., Kepecs, A., Harris, K. D., & Carandini, M. (2020). Dopaminergic and prefrontal basis of learning from sensory confidence and reward value. *Neuron*, *105*(4), 700–711. <https://doi.org/10.1016/j.neuron.2019.11.018>
- Law, C.-T., & Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience*, *12*(5), 655–663. <https://doi.org/10.1038/nn.2304>

- Leat, S. J., Yadav, N. K., & Irving, E. L. (2009). Development of visual acuity and contrast sensitivity in children. *Journal of Optometry*, *2*(1), 19–26. <https://doi.org/10.3921/joptom.2009.19>
- Lieder, F., Griffiths, T. L., Huys, Q. J. M., & Goodman, N. D. (2018). The anchoring bias reflects rational use of cognitive resources. *Psychonomic Bulletin & Review*, *25*(1), 322–349. <https://doi.org/10.3758/s13423-017-1286-8>
- Lindenberger, U., Li, S.-C., & Bäckman, L. (2006). Delineating brain-behavior mappings across the lifespan: Substantive and methodological advances in developmental neuroscience. Editorial. *Neuroscience & Biobehavioral Reviews*, *30*(6), 713–717. <https://doi.org/10.1016/j.neubiorev.2006.06.006>
- Lindenberger, U., & Mayr, U. (2014). Cognitive aging: Is there a dark side to environmental support? *Trends in Cognitive Sciences*, *18*(1), 7–15. <https://doi.org/10.1016/j.tics.2013.10.006>
- Luu, L., & Stocker, A. A. (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *eLife*, *7*, e33334. <https://doi.org/10.7554/eLife.33334>
- Ma, W. J. (2019). Bayesian decision models: A primer. *Neuron*, *104*(1), 164–175. <https://doi.org/10.1016/j.neuron.2019.09.037>
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–1438. <https://doi.org/10.1038/nm1790>
- Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, *37*, 205–220. <https://doi.org/10.1146/annurev-neuro-071013-014017>
- MacKay, D. J. C. (2003). *Information theory, inference and learning algorithms*. Cambridge University Press.
- Maia, T. V., & Frank, M. J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nature Neuroscience*, *14*(2), 154–162. <https://doi.org/10.1038/nn.2723>
- Manning, C., Wagenmakers, E.-J., Norcia, A. M., Scerif, G., & Boehm, U. (2020). Perceptual decision-making in children: Age-related differences and EEG correlates. *Computational Brain & Behavior*. <https://doi.org/10.1007/s42113-020-00087-7>
- Marković, D., & Kiebel, S. J. (2016). Comparative analysis of behavioral models for adaptive learning in changing environments. *Frontiers in Computational Neuroscience*, *10*, 33. <https://doi.org/10.3389/fncom.2016.00033>
- Marshall, L., Mathys, C. D., Ruge, D., de Berker, A. O., Dayan, P., Stephan, K. E., & Bestmann, S. (2016). Pharmacological fingerprints of contextual uncertainty.

- PLoS Biology*, 14(11), e1002575. <https://doi.org/10.1371/journal.pbio.1002575>
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory (Vol. 1)*. Oxford University Press.
- Mather, M., & Harley, C. W. (2016). The locus coeruleus: Essential for maintaining cognitive function and the aging brain. *Trends in Cognitive Sciences*, 20(3), 214–226. <https://doi.org/10.1016/j.tics.2016.01.001>
- Mathys, C. D., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience*, 5, 39. <https://doi.org/10.3389/fnhum.2011.00039>
- Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in Human Neuroscience*, 8, 825. <https://doi.org/10.3389/fnhum.2014.00825>
- McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, 84(4), 870–881. <https://doi.org/10.1016/j.neuron.2014.10.013>
- Mell, T., Heekeren, H. R., Marschner, A., Wartenburger, I., Villringer, A., & Reischies, F. M. (2005). Effect of aging on stimulus-reward association learning. *Neuropsychologia*, 43(4), 554–563. <https://doi.org/10.1016/j.neuropsychologia.2004.07.010>
- Mell, T., Wartenburger, I., Marschner, A., Villringer, A., Reischies, F. M., & Heekeren, H. R. (2009). Altered function of ventral striatum during reward-based decision making in old age. *Frontiers in Human Neuroscience*, 3, 34. <https://doi.org/10.3389/neuro.09.034.2009>
- Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences*, 114(19), E3859–E3868. <https://doi.org/10.1073/pnas.1615773114>
- Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The sense of confidence during probabilistic learning: A normative account. *PLoS Computational Biology*, 11(6), e1004305. <https://doi.org/10.1371/journal.pcbi.1004305>
- Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian probability: From neural origins to behavior. *Neuron*, 88(1), 78–92. <https://doi.org/10.1016/j.neuron.2015.09.039>
- Miller, K. J., Shenhav, A., & Ludvig, E. A. (2019). Habits without values. *Psychological Review*, 126(2), 292–311. <https://doi.org/10.1037/rev0000120>

- Mohr, P. N., Biele, G., & Heekeren, H. R. (2010). Neural processing of risk. *Journal of Neuroscience*, *30*(19), 6613–6619. <https://doi.org/10.1523/JNEUROSCI.0003-10.2010>
- Moran, R. J., Campo, P., Symmonds, M., Stephan, K. E., Dolan, R. J., & Friston, K. J. (2013). Free energy, precision and learning: The role of cholinergic neuromodulation. *Journal of Neuroscience*, *33*(19), 8227–8236. <https://doi.org/10.1523/JNEUROSCI.4255-12.2013>
- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience*, *32*(7), 2335–2343. <https://doi.org/10.1523/JNEUROSCI.4156-11.2012>
- Muller, T. H., Mars, R. B., Behrens, T. E. J., & O'Reilly, J. X. (2019). Control of entropy in neural models of environmental state. *eLife*, *8*, e39404. <https://doi.org/10.7554/eLife.39404.001>
- Munakata, Y., Snyder, H. R., & Chatham, C. H. (2012). Developing cognitive control: Three key transitions. *Current Directions in Psychological Science*, *21*(2), 71–77. <https://doi.org/10.1177/0963721412436807>
- Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related EEG signals and learning. *eLife*, *8*, e46975. <https://doi.org/10.7554/eLife.46975>
- Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, *7*(1), 11609. <https://doi.org/10.1038/ncomms11609>
- Nassar, M. R., & Frank, M. J. (2016). Taming the beast: Extracting generalizable knowledge from computational models of cognition. *Current Opinion in Behavioral Sciences*, *11*, 49–54. <https://doi.org/10.1016/j.cobeha.2016.04.003>
- Nassar, M. R., McGuire, T., J., Ritz, H., & Kable, J. W. (2019). Dissociable forms of uncertainty-driven representational change across the human brain. *Journal of Neuroscience*, *39*(9), 1688–1698. <https://doi.org/10.1523/JNEUROSCI.1713-18.2018>
- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasley, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, *15*(7), 1040–1046. <https://doi.org/10.1038/nn.3130>
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a

- changing environment. *Journal of Neuroscience*, *30*(37), 12366–12378. <https://doi.org/10.1523/JNEUROSCI.0822-10.2010>
- Nienborg, H., & Cumming, B. G. (2009). Decision-related activity in sensory neurons reflects more than a neuron’s causal effect. *Nature*, *459*(7243), 89–92. <https://doi.org/10.1038/nature07821>
- Nyberg, L., Salami, A., Andersson, M., Eriksson, J., Kalpouzos, G., Kauppi, K., Lind, J., Pudas, S., Persson, J., & Nilsson, L.-G. (2010). Longitudinal evidence for diminished frontal cortex function in aging. *Proceedings of the National Academy of Sciences*, *107*(52), 22682–22686. <https://doi.org/10.1073/pnas.1012651108>
- O’Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, *110*(38), E3660–E3669. <https://doi.org/10.1073/pnas.1305373110>
- Palminteri, S., Wyart, V., & Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences*, *21*(6), 425–433. <https://doi.org/10.1016/j.tics.2017.03.011>
- Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Computational Biology*, *7*(1), e1001048. <https://doi.org/10.1371/journal.pcbi.1001048>
- Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O’Doherty, J. P. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron*, *79*(1), 191–201. <https://doi.org/10.1016/j.neuron.2013.04.037>
- Peters, M. A. K., Thesen, T., Ko, Y. D., Maniscalco, B., Carlson, C., Davidson, M., Doyle, W., Kuzniecky, R., Devinsky, O., Halgren, E., & Lau, H. (2017). Perceptual confidence neglects decision-incongruent evidence in the brain. *Nature Human Behaviour*, *1*(7), 0139. <https://doi.org/10.1038/s41562-017-0139>
- Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. *PLoS Computational Biology*, *16*(7), e1007963. <https://doi.org/10.1371/journal.pcbi.1007963>
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, *19*(3), 366–374. <https://doi.org/10.1038/nm.4240>
- Purcell, A., & Kiani, R. (2016). Hierarchical decision processes that operate over distinct timescales underlie choice and changes in strategy. *Proceedings of the National Academy of Sciences*, *113*(31), E4531–E4540. <https://doi.org/10.1073/pnas.1524685113>

- Qiu, C., Luu, L., & Stocker, A. A. (2020). Benefits of commitment in hierarchical inference. *Psychological Review*, *127*(4), 622–639. <https://doi.org/10.1037/rev0000193>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, *9*(7), 545–556. <https://doi.org/10.1038/nrn2357>
- Raz, N., & Rodrigue, K. M. (2006). Differential aging of the brain: Patterns, cognitive correlates and modifiers. *Neuroscience & Biobehavioral Reviews*, *30*(6), 730–748. <https://doi.org/10.1016/j.neubiorev.2006.07.001>
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement (A. H. Black & W. F. Prokasy, Eds.). In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning ii: Current research and theory*. Appleton-Century-Crofts.
- Ridderinkhof, K. R., Span, M. M., & van der Molen, M. W. (2002). Perseverative behavior and adaptive control in older adults: Performance monitoring, rule induction, and set shifting. *Brain and cognition*, *49*(3), 382–401. <https://doi.org/10.1006/brcg.2001.1506>
- Rorie, A. E., Gao, J., McClelland, J. L., & Newsome, W. T. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE*, *5*(2), e9308. <https://doi.org/10.1371/journal.pone.0009308>
- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd edition). Prentice Hall.
- Rutledge, R. B., Lazzaro, S. C., Lau, B., Myers, C. E., Gluck, M. A., & Glimcher, P. W. (2009). Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *Journal of Neuroscience*, *29*(48), 15104–15114. <https://doi.org/10.1523/JNEUROSCI.3524-09.2009>
- Samanez-Larkin, G. R., Kuhnen, C. M., Yoo, D. J., & Knutson, B. (2010). Variability in nucleus accumbens activity mediates age-related suboptimal financial risk taking. *Journal of Neuroscience*, *30*(4), 1426–1434. <https://doi.org/10.1523/JNEUROSCI.4902-09.2010>
- Samanez-Larkin, G. R., Levens, S. M., Perry, L. M., Dougherty, R. F., & Knutson, B. (2012). Frontostriatal white matter integrity mediates adult age differences in probabilistic reward learning. *Journal of Neuroscience*, *32*(15), 5333–5337. <https://doi.org/10.1523/JNEUROSCI.5756-11.2012>
- Samanez-Larkin, G. R., Worthy, D. A., Mata, R., McClure, S. M., & Knutson, B. (2014). Adult age differences in frontostriatal representation of prediction

- error but not reward outcome. *Cognitive, Affective & Behavioral Neuroscience*, *14*(2), 672–682. <https://doi.org/10.3758/s13415-014-0297-4>
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. *Neuron*, *91*(6), 1402–1412. <https://doi.org/10.1016/j.neuron.2016.08.019>
- Shing, Y. L., Schmiedek, F., Lövdén, M., & Lindenberger, U. (2012). Memory updating practice across 100 days in the COGITO study. *Psychology and Aging*, *27*(2), 451–461. <https://doi.org/10.1037/a0025568>
- Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(6), 1865–1897. <https://doi.org/10.1037/a0016926>
- Soltani, A., & Izquierdo, A. (2019). Adaptive learning under expected and unexpected uncertainty. *Nature Reviews Neuroscience*, *20*(10), 635–644. <https://doi.org/10.1038/s41583-019-0180-y>
- Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, *20*(4), 581–589. <https://doi.org/10.1038/nn.4520>
- Stocker, A. A., & Simoncelli, E. P. (2007). A Bayesian model of conditioned perception (J. C. Platt, D. Koller, Y. Singer, & S. Roweis, Eds.). In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems*. MIT Press.
- Stolyarova, A., Rakhshan, M., Hart, E. E., O’Dell, T. J., Peters, M. A. K., Lau, H., Soltani, A., & Izquierdo, A. (2019). Contributions of anterior cingulate cortex and basolateral amygdala to decision confidence and learning under uncertainty. *Nature Communications*, *10*(1), 4704. <https://doi.org/10.1038/s41467-019-12725-1>
- Summerfield, C., & Koechlin, E. (2010). Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Frontiers in Human Neuroscience*, *4*, 208. <https://doi.org/10.3389/fnhum.2010.00208>
- Summerfield, C., & Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: Neural and computational mechanisms. *Frontiers in Neuroscience*, *6*, 70. <https://doi.org/10.3389/fnins.2012.00070>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.
- Terry, W. S. (2015). *Learning and memory: Basic principles, processes, and procedures*. Psychology Press.

- Toga, A. W., Thompson, P. M., & Sowell, E. R. (2006). Mapping brain maturation. *Trends in Neurosciences*, *29*(3), 148–159. <https://doi.org/10.1016/j.tins.2006.01.007>
- Troller-Renfree, S. V., Buzzell, G. A., & Fox, N. A. (2020). Changes in working memory influence the transition from reactive to proactive cognitive control during childhood. *Developmental Science*, e12959. <https://doi.org/10.1111/desc.12959>
- Ullsperger, M., Danielmeier, C., & Jocham, G. (2014). Neurophysiology of performance monitoring and adaptive behavior. *Physiological Reviews*, *94*(1), 35–79. <https://doi.org/10.1152/physrev.00041.2012>
- Ullsperger, M., Fischer, A. G., Nigbur, R., & Endrass, T. (2014). Neural mechanisms and temporal dynamics of performance monitoring. *Trends in Cognitive Sciences*, *18*(5), 259–267. <https://doi.org/10.1016/j.tics.2014.02.009>
- Urai, A. E., de Gee, J. W., Tsetsos, K., & Donner, T. H. (2019). Choice history biases subsequent evidence accumulation. *eLife*, *8*, e46331. <https://doi.org/10.7554/eLife.46331>
- Van Bergen, R. S., & Jehee, J. F. M. (2017). Modeling correlated noise is necessary to decode uncertainty. *NeuroImage*, *180*, 78–87. <https://doi.org/10.1016/j.neuroimage.2017.08.015>
- Van Bergen, R. S., Ma, W. J., Pratte, M. S., & Jehee, J. F. M. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*, *18*(12), 1728–1730. <https://doi.org/10.1038/nn.4150>
- Van den Bos, W., Bruckner, R., Nassar, M. R., Mata, R., & Eppinger, B. (2018). Computational neuroscience across the lifespan: Promises and pitfalls. *Developmental Cognitive Neuroscience*, *33*, 42–53. <https://doi.org/10.1016/j.dcn.2017.09.008>
- Van den Bos, W., Cohen, M. X., Kahnt, T., & Crone, E. A. (2012). Striatum-medial prefrontal cortex connectivity predicts developmental changes in reinforcement learning. *Cerebral Cortex*, *22*(6), 1247–1255. <https://doi.org/10.1093/cercor/bhr198>
- Van der Schaaf, M. E., Warmerdam, E., Crone, E. A., & Cools, R. (2011). Distinct linear and non-linear trajectories of reward and punishment reversal learning during development: Relevance for dopamine's role in adolescent decision making. *Developmental Cognitive Neuroscience*, *1*(4), 578–590. <https://doi.org/10.1016/j.dcn.2011.06.007>
- Van Duijvenvoorde, A. C. K., Jansen, B. R. J., Griffioen, E. S., van der Molen, M. W., & Huizenga, H. M. (2013). Decomposing developmental differences in probabilistic feedback learning: A combined performance and heart-rate

- analysis. *Biological Psychology*, *93*(1), 175–183. <https://doi.org/10.1016/j.biopsycho.2013.01.006>
- Vossel, S., Bauer, M., Mathys, C. D., Adams, R. A., Dolan, R. J., Stephan, K. E., & Friston, K. J. (2014). Cholinergic stimulation enhances Bayesian belief updating in the deployment of spatial attention. *Journal of Neuroscience*, *34*(47), 15735–15742. <https://doi.org/10.1523/JNEUROSCI.0091-14.2014>
- Whiteley, L., & Sahani, M. (2008). Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of Vision*, *8*(3), 1–15. <https://doi.org/10.1167/8.3.2>
- Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *eLife*, *8*, e49547. <https://doi.org/10.7554/eLife.49547>
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Computation*, *22*(9), 2452–2476. https://doi.org/10.1162/NECO_a_00007
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, *81*(2), 267–279. <https://doi.org/10.1016/j.neuron.2013.11.005>
- Wood, W., & Runger, D. (2016). Psychology of habit. *Annual Review of Psychology*, *67*(1), 289–314. <https://doi.org/10.1146/annurev-psych-122414-033417>
- Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: Confidence and error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1594), 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681–692. <https://doi.org/10.1016/j.neuron.2005.04.026>

9 Appendix

9.1 Deutsche Zusammenfassung

Adaptives Verhalten verlangt eine ständige Verarbeitung von neuen Ereignissen sowie eine Reaktion auf diese. In der Psychologie und den Neurowissenschaften wird dies als Lern- und Entscheidungsprozess bezeichnet. Solche Prozesse finden in der Regel in Situationen statt, in denen Unsicherheit über aktuelle und zukünftige Ereignisse herrscht. Um sich in derartigen Situationen erfolgreich zurechtfinden zu können, muss man aus den Erfahrungen der Vergangenheit Vorhersagen über zukünftige Ereignisse ableiten.

Die Dissertation behandelt folgende Themen:

1. Normative Berechnungen, die dem Lernen unter Unsicherheit zugrunde liegen sollten.
2. Verzerrungen, die bei der Berücksichtigung von Unsicherheit zu Abweichungen vom normativen Lernen führen.
3. Altersrelatierte Unterschiede über die Lebensspanne, die beim Lernen unter Unsicherheit charakteristisch sind.

Der Begriff normative Berechnungen aus dem Forschungsfeld *Computational Neuroscience* bezieht sich in dieser Dissertation auf Berechnungen, die zu einer optimalen Lösung eines Lern- und Entscheidungsproblems führen. Meine Dissertationsstudien basieren auf Computermodellen, die normative Berechnungen implementieren und Unsicherheit formal definieren. Anhand dieser Modelle wird systematisch untersucht, inwieweit Menschen im jüngeren Erwachsenenalter und über die Lebensspanne Unsicherheit berücksichtigen, um aus ihren Erfahrungen zu lernen.

Zu Beginn der Dissertation wird demonstriert, dass adaptives Verhalten aus mehreren Schritten besteht, von der Repräsentation der Umgebung über die Entscheidungsfindung bis hin zu Lernprozessen ([Introduction](#)). Auf dieser Grundlage stelle

ich zunächst ein Modell vor, das Unsicherheit in drei Formen unterteilt: Perzeptuelle Unsicherheit, erwartete Unsicherheit und unerwartete Unsicherheit ([Normative computations](#)). Perzeptuelle Unsicherheit hängt mit der Verarbeitung sensorischer Informationen zusammen, erwartete Unsicherheit ergibt sich aus der Variabilität von Ereignissen und unerwartete Unsicherheit ist die Folge von Veränderungen in der Umgebung. Für jede dieser drei Formen beschreibe ich, wie Unsicherheit beim Lernen aufgrund von normativen Berechnungen berücksichtigt werden sollte. Danach zeige ich, dass Verzerrungen, also Abweichungen von den normativen Berechnungen, durch die man sich an Unsicherheit anpasst, charakteristisch für menschliches Lernen sind ([Uncertainty biases](#)). Abschließend erfolgt eine Darstellung, die verdeutlicht, warum die Erfassung dieser Verzerrungen mit Computermodellen nützlich ist, um altersrelatierte Unterschiede über die Lebensspanne beim Lernen unter Unsicherheit besser verstehen zu können ([Lifespan differences](#)).

In der ersten Dissertationsstudie ([Bruckner et al., 2020a](#)) wurde untersucht, welche normativen Berechnungen beim Lernen unter perzeptueller Unsicherheit wichtig sind, in welchem Maße jüngere Erwachsene dementsprechend lernen und wie dieser Prozess durch vorherige perzeptuelle Entscheidungen verzerrt wird. In der zweiten Studie ([Nassar et al., 2016](#)) wurde Lernen unter erwarteter und unerwarteter Unsicherheit bei jüngeren und älteren Erwachsenen untersucht. Insbesondere wurde hier erforscht, inwiefern Verzerrungen bei der Berücksichtigung dieser Unsicherheiten altersrelatierte Lernunterschiede erklären. Die dritte Studie ([Bruckner et al., 2020b](#)) hat darauf aufgebaut und speziell bei Kindern und älteren Erwachsenen untersucht, inwiefern sie auf vereinfachte Lernstrategien zurückgreifen und auf normative Berechnungen verzichten. Die vierte Studie ([Van den Bos et al., 2018](#)) hat schließlich beschrieben, wie Computermodelle die Erforschung altersrelatierter Lernunterschiede über die Lebensspanne unterstützen können.

Die Ergebnisse der oben genannten Studien werden im Folgenden kurz zusammengefasst. In [Bruckner et al. \(2020a\)](#) konnten wir zeigen, dass perzeptuelle Un-

sicherheit beim Lernen zu vorschnellen Schlussfolgerungen auf Basis von Fehlinterpretationen einer Wahrnehmung führen kann. Um vorschnelle Schlussfolgerungen zu vermeiden, sollte man sich beim Lernen unter perzeptueller Unsicherheit vorsichtiger verhalten als in perzeptuell eindeutigen Situationen. Wir fanden in dieser Studie heraus, dass Menschen perzeptuelle Unsicherheit beim Lernen berücksichtigen. Zusätzlich stellten wir allerdings eine Verzerrung bei der Berücksichtigung perzeptueller Unsicherheit aufgrund von früheren perzeptuellen Entscheidungen beim Lernen fest, die wiederum zu einer weniger vorsichtigen Anpassung des Lernverhaltens führt.

In [Nassar et al. \(2016\)](#) fanden wir Hinweise darauf, dass altersrelatierte Lernunterschiede mit Verzerrungen bei der Anpassung an erwartete Unsicherheit zusammenhängen. Insbesondere stellten wir fest, dass ältere Erwachsene (60 bis 80 Jahre) dazu neigen, die Unsicherheit über ihre Erwartungen im Vergleich zu jüngeren Erwachsenen (20 bis 30 Jahre) zu unterschätzen. Diese Form der Unsicherheitsunterschätzung führt zu einem weniger flexiblen Lernverhalten im Vergleich zu jüngeren Erwachsenen.

In [Bruckner et al. \(2020b\)](#) wurde gezeigt, dass altersrelatierte Unterschiede beim Lernen unter Unsicherheit damit zusammenhängen, dass Kinder (7 bis 11 Jahre) und ältere Erwachsene häufig auf vereinfachte Lernstrategien zurückgreifen, was dazu führt, dass Verhalten wiederholt (Perseveration) oder stärker durch die Umgebung beeinflusst wird (externe Kontrolle).

Abschließend wurde in [Van den Bos et al. \(2018\)](#) argumentiert, dass Computermodellierung eine wichtige Methode ist, um altersrelatierte Unterschiede beim Lernen und in der Entscheidungsfindung besser zu verstehen. Hier wurden sowohl die Vorteile der Anwendung von Computermodellen zur Erforschung altersrelatierter Verhaltensunterschiede (in Bezug auf Risikobereitschaft, Strategiewahl und Verstärkungslernen) als auch potenzielle Fallstricke aufgezeigt.

Nach der Diskussion der Dissertationsprojekte ([General discussion and future directions](#)) stelle ich ein kognitives Modell zum Lernen unter Unsicherheit vor, das auf den neuen Erkenntnissen meiner Studien und früheren Arbeiten aus der Literatur

9 Appendix

basiert ([Uncertainty in the cycle of adaptive behavior](#)). Zusammenfassend legt meine Dissertation dar, dass Lernen ein dynamischer Prozess ist, der von vielfältigen Formen der Unsicherheit beeinflusst wird. Menschen berücksichtigen ihre Unsicherheit beim Lernen, weisen aber charakteristische Unsicherheitsverzerrungen auf, die sich im Laufe der Lebensspanne erheblich verändern.

9.2 List of publications

1. Bruckner, R., Heekeren, H. R., & Oswald, D. (2020). Belief states and categorical-choice biases determine reward-based learning under perceptual uncertainty. *bioRxiv*. <https://doi.org/10.1101/2020.09.18.303495>
2. Bruckner, R., Nassar, M. R., Li, S.-C., & Eppinger, B. (2020). Default beliefs guide learning under uncertainty in children and older adults. *PsyArXiv*. <https://doi.org/10.31234/osf.io/nh9bq>
3. Frömer, R., Nassar, M. R., Bruckner, R., Stürmer, B., Sommer, W., & Yeung, N. (2020). I knew that! Response-based outcome predictions and confidence regulate feedback processing and learning. *bioRxiv*. <https://doi.org/10.1101/442822>
4. Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related EEG signals and learning. *eLife*, 8, e46975. <https://doi.org/10.7554/eLife.46975>
5. Oswald, D., Schneider, S., Bruckner, R., & Horvarth, L. (2019). Power, positive predictive value, and sample size calculations for random field theory-based fMRI inference. *bioRxiv*. <https://doi.org/10.1101/613331>
6. Oswald, D., Schneider, S., Bruckner, R., & Horvarth, L. (2018). Random field theory-based p-values: A review of the SPM implementation. *arXiv*. <https://arxiv.org/abs/1808.04075>
7. Van den Bos, W., Bruckner, R., Nassar, M. R., Mata, R., & Eppinger, B. (2018). Computational neuroscience across the lifespan: Promises and pitfalls. *Developmental Cognitive Neuroscience*, 33, 42–53. <https://doi.org/10.1016/j.dcn.2017.09.008>

9 Appendix

8. Nassar, M. R., Bruckner, R., & Eppinger, B. (2016). What do we GANE with age? [Invited peer commentary]. *Behavioral and Brain Sciences*, *39*, E218. <https://doi.org/10.1017/S0140525X15001892>
9. Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, *7*(1), 11609. <https://doi.org/10.1038/ncomms11609>
10. Eppinger, B., & Bruckner, R. (2015). Towards a mechanistic understanding of age-related changes in learning and decision making: A neuro-computational approach (T. M. Hess, J. Strough, & C. E. Löckenhoff, Eds.). In T. M. Hess, J. Strough, & C. E. Löckenhoff (Eds.), *Aging and decision making: Empirical and applied perspectives*. Academic Press. <https://doi.org/10.1016/B978-0-12-417148-0.00004-2>

9.3 Talks and colloquium presentations

- Jan 2020** Leipzig, Germany, *MPI for Human Cognitive and Brain Sciences*, Dr. Roland G. Benoit. "Adaptive learning under uncertainty: Computational mechanisms and lifespan differences".
- Dec 2019** Egmond aan Zee, The Netherlands, *Symposium at the Brain and Cognition conference*. "Lifespan age differences in the regulation of learning in changing and uncertain environments".
- Dec 2019** Dublin, Ireland, *Trinity College Dublin*, Prof. Tomás J. Ryan. "Adaptive learning under uncertainty: Computational mechanisms and lifespan differences".
- Oct 2019** Providence, USA, *Brown University*, Prof. Matthew R. Nassar (online). "Lifespan age differences in the regulation of learning during sequential decisions under uncertainty."
- Oct 2019** Berlin, Germany, *Freie Universität Berlin*, Science Slam (in German). "Computermodelle in den Kognitiven Neurowissenschaften".
- June 2019** Dresden, Germany, *Symposium at the Psychologie und Gehirn conference*. "Computational mechanisms of human state-action-reward contingency learning under perceptual uncertainty".
- Dec 2018** Frankfurt, Germany, *Goethe Universität Frankfurt*. Prof. Yee Lee Shing. "Lifespan differences in the regulation of learning rates".
- Sep 2018** Tegernsee, Germany, *MPS-UCL Symposium and advanced course on computational psychiatry and ageing research at Marbach Castle*. "Computational mechanisms of human state-action-reward contingency learning under perceptual uncertainty".

Feb 2018 Berlin, Germany, MPI for Human Development, Dr. Nicolas W. Schuck.

"Learning to make economic decisions under perceptual uncertainty."

June 2017 Providence, USA, Brown University, Prof. Michael J. Frank, Dr.

Matthew R. Nassar. "Learning to make economic decisions under perceptual uncertainty".

Oct 2016 Berlin, Germany, IMPRS LIFE Academy. "Value-based decision mak-

ing under perceptual uncertainty".

May 2016 Virginia, USA, IMPRS LIFE Academy. "Lifespan differences in the

regulation of learning rates".

May 2016 Berlin, Germany, Symposium at the Psychologie und Gehirn conference.

"Individual and age-related differences in the regulation of learning rates".

9.4 Eigenanteil

Erklärung gemäß § 7 Abs. 3 Satz 4 der Promotionsordnung über den Eigenanteil an den veröffentlichten oder zur Veröffentlichung vorgesehenen eingereichten wissenschaftlichen Schriften im Rahmen meiner publikationsbasierten Arbeit

I. Name, Vorname: Bruckner, Rasmus

Institut: Fachbereich Erziehungswissenschaft und Psychologie

Promotionsfach: Psychologie

Titel: Master of Science Psychologie

II. Nummerierte Aufstellung der eingereichten Schriften (Titel, Autoren, wo und wann veröffentlicht bzw. eingereicht):

1. Bruckner, R., Heekeren, H. R., & Ostwald, D. (2020). Belief states and categorical-choice biases determine reward-based learning under perceptual uncertainty. *bioRxiv*. <https://doi.org/10.1101/2020.09.18.303495>
2. Bruckner, R., Nassar, M. R., Li, S.-C., & Eppinger, B. (submitted). Default beliefs guide learning under uncertainty in children and older adults. <https://doi.org/10.31234/osf.io/nh9bq>.
3. Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, *7*(1), 11609. <https://doi.org/10.1038/ncomms11609>.
4. Van den Bos, W., Bruckner, R., Nassar, M. R., Mata, R., & Eppinger, B. (2018). Computational neuroscience across the lifespan: Promises and pitfalls. *Developmental Cognitive Neuroscience*, *33*, 42–53. <https://doi.org/10.1016/j.dcn.2017.09.008>.

III. Darlegung des eigenen Anteils an diesen Schriften:

Die Bewertung des Eigenanteils erfolgt auf der Skala: “vollständig – überwiegend – mehrheitlich – in Teilen”.

Zu II. 1.: Konzeption (mehrheitlich), Versuchsdesign (überwiegend), Datenerhebung (vollständig), Methodenentwicklung (in Teilen), Datenanalyse (vollständig), Software (mehrheitlich), Ergebnisdiskussion (vollständig), Erstellen des Manuskriptes (mehrheitlich).

Zu II. 2.: Konzeption (überwiegend), Versuchsdesign (überwiegend), Datenerhebung (vollständig), Methodenentwicklung (mehrheitlich), Datenanalyse (vollständig), Ergebnisdiskussion (mehrheitlich), Erstellen des Manuskriptes (mehrheitlich).

Zu II. 3.: Datenerhebung (überwiegend), Datenanalyse (in Teilen), Programmierung (in Teilen), Ergebnisdiskussion (in Teilen), Erstellen des Manuskriptes (in Teilen).

Zu II. 4.: Konzeption (in Teilen), Simulationen (vollständig), Literaturrecherche (in Teilen), Abbildungen (mehrheitlich), Ergebnisdiskussion (in Teilen), Erstellen des Manuskriptes (in Teilen).

9.5 Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt,

- dass ich die vorliegende Arbeit selbstständig und ohne unerlaubte Hilfe verfasst habe,
- dass ich mich nicht bereits anderwärts um einen Doktorgrad beworben habe und keinen Doktorgrad in dem Promotionsfach Psychologie besitze und
- dass ich die zugrunde liegende Promotionsordnung vom 08.08.2016 kenne.

Berlin, 01.10.2020

Rasmus Bruckner

9.6 Research articles

Bruckner, R., Heekeren, H. R., & Oswald, D. (2020). Belief states and categorical-choice biases determine reward-based learning under perceptual uncertainty. *bioRxiv*. <https://doi.org/10.1101/2020.09.18.303495>

Bruckner, R., Nassar, M. R., Li, S.-C., & Eppinger, B. (2020). Default beliefs guide learning under uncertainty in children and older adults. *PsyArXiv*. <https://doi.org/10.31234/osf.io/nh9bq>.

Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, *7*(1), 11609. <https://doi.org/10.1038/ncomms11609>.

Van den Bos, W., Bruckner, R., Nassar, M. R., Mata, R., & Eppinger, B. (2018). Computational neuroscience across the lifespan: Promises and pitfalls. *Developmental Cognitive Neuroscience*, *33*, 42–53. <https://doi.org/10.1016/j.dcn.2017.09.008>.