


## Accepted Manuscript

Title: Computational Neuroscience across the Lifespan:  
Promises and Pitfalls

Authors: Wouter van den Bos, Rasmus Bruckner, Matthew R.  
Nassar, Rui Mata, Ben Eppinger



PII: S1878-9293(17)30106-8  
DOI: <https://doi.org/10.1016/j.dcn.2017.09.008>  
Reference: DCN 495

To appear in:

Received date: 11-5-2017  
Revised date: 19-8-2017  
Accepted date: 29-9-2017

Please cite this article as: van den Bos, Wouter, Bruckner, Rasmus, Nassar, Matthew R., Mata, Rui, Eppinger, Ben, Computational Neuroscience across the Lifespan: Promises and Pitfalls. *Developmental Cognitive Neuroscience* <https://doi.org/10.1016/j.dcn.2017.09.008>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

**Computational Neuroscience across the Lifespan: Promises and Pitfalls**

Wouter van den Bos<sup>1,2</sup>, Rasmus Bruckner<sup>3,4</sup>, Matthew R. Nassar<sup>5</sup>, Rui Mata<sup>6</sup>, & Ben Eppinger<sup>7,8</sup>

<sup>1</sup>Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany

<sup>2</sup>Department of Psychology, University of Amsterdam, Amsterdam, the Netherlands

<sup>3</sup>Department of Education and Psychology, Freie Universität Berlin, Berlin, Germany

<sup>4</sup>International Max Planck Research School LIFE, Berlin, Germany

<sup>5</sup>Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, USA

<sup>6</sup>Center for Cognitive and Decision Sciences, Department of Psychology, University of Basel, Basel, Switzerland

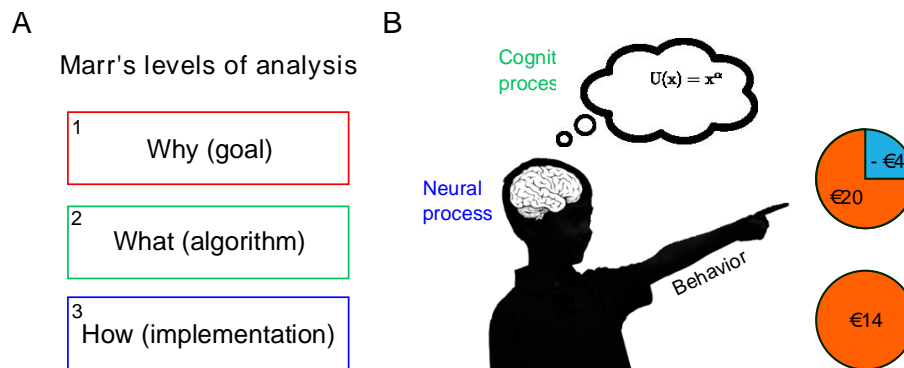
<sup>7</sup>Department of Psychology, Concordia University, Montreal, Canada

<sup>8</sup>Department of Psychology, TU Dresden, Dresden, Germany

Corresponding authors: Wouter van den Bos, [vandenbos@mpib-berlin.mpg.de](mailto:vandenbos@mpib-berlin.mpg.de) & Ben Eppinger, [ben.eppinger@concordia.ca](mailto:ben.eppinger@concordia.ca)

In recent years, the application of computational modeling in studies on age-related changes in decision making and learning has gained in popularity. One advantage of computational models is that they provide access to latent variables that cannot be directly observed from behavior. In combination with experimental manipulations, these latent variables can help to test hypotheses about age-related changes in behavioral and neurobiological measures at a level of *specificity* that is not achievable with descriptive analysis approaches alone. This level of specificity can in turn be beneficial to establish the *identity* of the corresponding behavioral and neurobiological mechanisms. In this paper, we will illustrate applications of computational methods using examples of lifespan research on risk taking, strategy selection and reinforcement learning. We will elaborate on problems that can occur when computational neuroscience methods are applied to data of different age groups. Finally, we will discuss potential targets for future applications and outline general shortcomings of computational neuroscience methods for research on human lifespan development.

Over the past two decades there has been a significant increase in the number of cognitive neuroscience studies of lifespan development (Amso & Scerif, 2015; Li & Rieckmann, 2014; Samanez-Larkin & Knutson, 2015). However, despite this inflation in empirical studies, there has been a serious lag in the development of comprehensive theories linking brain development and behavior (Pfeifer & Allen, 2012, 2016; van den Bos & Eppinger, 2016). The current conundrum of developmental neuroscience consists of two major explanatory problems. First, there is a *specificity problem*: The current verbal theories of neurocognitive development are not specific enough to be translated into precise behavioral and neuroscience predictions. As a result, it is often impossible to tell whether new neuroscientific data confirm *or* falsify existing theories (van den Bos & Eppinger, 2016). Related to this issue is the *identity problem*: We are often unable to precisely identify the processes that underlie developmental differences in behavior.



**Figure 1.** A) Marr's levels of analysis. B) Cartoon of a child choosing between two wheels of fortune. A risky option with a 75% chance of winning 20 Euro and 25% chance of losing 4 Euro versus a safe option with 100% chance of winning 14 Euro. Choosing the option with the highest outcome variance (the risky option) is often considered risk seeking behavior in the context of these tasks (even when total expected value is the same).

David Marr's levels of analysis approach (Figure 1A) is probably the best-known framework devised to formalize our understanding of brain-behavior relationships (Marr & Poggio, 1976). In this approach, the first level is more abstract and is concerned with the "Why" question of behavior (e.g., why would it be beneficial for adolescents take more risks?). The "Why" question may inspire experimental design that can help understand when a particular behavior occurs. For instance, if risk taking is a form of costly signaling, showing off your strength to gain social status in the group, it should occur specifically in presence of relevant peers. Computational analyses of behavior would target Marr's algorithmic level (the "What" question). The algorithmic level focuses on the rules that underlie behavior (e.g., risk preference can be formalized as calculating expected utility (EU) or minimax choice rules, see below for more detail). Finally, the third level ("How" question) refers to the implementation of the algorithm. What is the neurobiological substrate that supports this behavior? Marr famously argued that it is impossible to bridge the gap from the implementational level (neural processes) to behavior (risk seeking) without referring to the algorithmic level, because neural processes per se do not tell us anything about the algorithm they implement (see also (Krakauer, Ghazanfar, Gomez-Marín, Maciver, & Poeppel, 2017)).

Following this line of argument, we argue that computational modeling can foster the understanding of lifespan development by providing testable theories that provide a bridge between Marr's levels of analyses. One key advantage of computational models is that they allow us to capture latent variables that cannot be directly observed from behavior. As a result, by using computational models, we are better able to build specific theories about lifespan

development and to identify processes that underlie developmental changes in behavior. In comparison to purely descriptive theories, the use of computational models can lead to substantially different predictions about behavior and explanations about the underlying processes. We will illustrate our arguments with specific examples in the context of decisions from descriptions (risk-taking and strategy selection) and reinforcement learning. We will also discuss several problems and potential solutions that arise when applying computational modeling to data of different age groups. Finally, we conclude with a discussion of the prospects of neuro-computational approaches in terms of our understanding of developmental changes across the human lifespan.

### **Models of Judgment and Decision-Making**

There has been a considerable increase in the use of paradigms from behavioral economics to study developmental changes underlying judgment and decision making. For instance, several developmental fMRI studies used risky decision making paradigms in which participants have to make repeated choices between safe and risky monetary gambles (for review see (Crone, van Duijvenvoorde, & Peper, 2016)). In most cases the aim of these studies was to investigate the developmental trajectories of risky decision making with the prediction that increased risk taking in teenagers results from an imbalance in the neural systems of motivation and cognitive control. In almost all of these studies, the analytic approach relied on objective measures of risk, e.g. comparing high versus low probability gambles or options of different expected value (EV: probability ( $p$ ) multiplied by outcome ( $x$ )). However, since the development of prospect theory (Kahneman & Tversky, 1979) we know that the subjective evaluation of outcomes in most cases does not match the (objective) expected value of choice options. Furthermore, it is very likely that the subjective judgment of outcomes and probabilities differ between age groups (Eppinger, Mock, & Kray, 2009; Harbaugh, Krause, & Vesterlund, 2002; Paulsen, Platt, Huettel, & Brannon, 2011; van den Bos, Cohen, Kahnt, & Crone, 2012).

One way of addressing these issues is to refer to expected utility (EU) models, such as cumulative prospect theory (CPT; (Tversky & Kahneman, 1992)), that explicitly account for the fact that people do not behave as EV models would predict. One key notion of EU models is that of decreasing marginal sensitivity. That is, outcomes have decreasing marginal effects as more is gained or lost (see Box 1). For instance, this implies that receiving 10 Euro when you have nothing is subjectively experienced as a bigger gain than receiving 10 Euro when you already have 10.000 Euro on your bank account. This principle is very similar to the notion in psychophysics that the objective intensity of a stimulus is distinct from the subjective intensity that guides behavior (G. Fechner, 1966). According to most EU models the transformation of an outcome into its utility ( $U$ ) is described by a power function

$$(1) \quad U(x) = x^\alpha$$

where in a monetary gamble  $x$  denotes the objective amount of money, and where  $0 < \alpha < 1$  denotes the degree that utility is diminished when the amounts of money increase. Thus,  $U(x)$  is a concave value function that describes the degree of risk-averse behavior. Note that because the function is concave, the subjective utility of gaining 20 Euro is only around 10 on the subjective utility scale. As a result, a risky prospect of gaining 20 Euro (at 10% chance) is not so attractive compared to a sure 2 Euro, even if objective EV is equal. CPT extends the more primitive EU value functions by adding three additional assumptions; (1) that there is a

distinction between sensitivity for gains and losses, (2) that gains and losses are defined relative to a reference point, and (3) that there is asymmetry in the steepness of the value functions for losses and gains (losses loom larger than gains, see Box 1).

### Box 1. Cumulative Prospect Theory

Cumulative Prospect Theory (CPT) suggest the following value function:

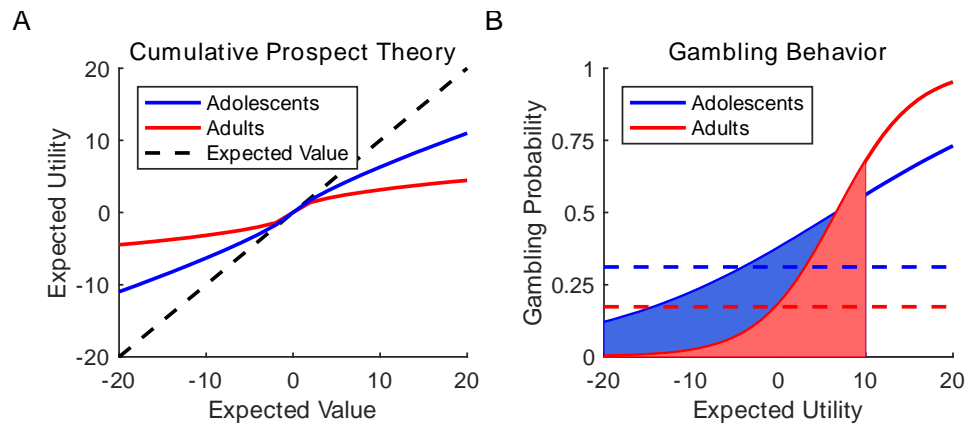
$$U(x) = \begin{cases} x^\alpha, & x \geq 0 \\ -\lambda(-x)^\beta, & x < 0 \end{cases}$$

where in a monetary gamble  $x$  denotes the objective amount of money, and where  $0 < \alpha < 1$  denotes the degree that utility is diminished when positive amounts of money increase, and where  $0 < \beta < 1$  denotes the degree that utility is diminished when the amounts of monetary losses increase. In addition, the  $\lambda$  parameter captures the psychological difference between losses and gains (losses tend to loom larger than gains).

Besides a subjective value function, CPT also assumes a psychological transformation of objective probabilities,  $p$ , into subjective probabilities,  $\pi(p)$ . For simple gambles with just two possible outcomes subjective probabilities are expressed as:

$$\pi(p) = \frac{p^\gamma}{p^\gamma + (1-p)^\gamma}$$

where the parameter  $\gamma$  indicates the strength of the transformation of probabilities. The typical shape of the weighting function suggests that people overestimate small probabilities and underestimate high probabilities. Finally, like most other utility models the expected utility EU of a gamble is the product of  $U(x)$  and  $\pi(p)$ .



**A)** Fit of CPT for a set of choices by hypothetical adolescents and adults. **B)** Two response probability curves for two fictitious groups of adolescents and adults that only differ in their level of choice sensitivity (see eq. 2 in text), but have the exact same EU estimations. Note that in this example the mean level of risky choices (as indicated with the dashed lines) will be higher for the adolescents than the adults.

In heuristic models of adolescent risk taking it is often assumed that an increased proportion of risky choices is the result of increased reward sensitivity. However, in many cases this relationship is only loosely defined and therefore it is difficult to make specific predictions about choice behavior or neural measures (van den Bos & Eppinger, 2016). In contrast, CPT suggests that differences in risk preferences between age groups can be attributed to different mechanisms. Let us consider a possible gamble such as the one displayed in Figure 1B: Option A is associated with a 75% chance of winning 20 Euro and a 25% chance of losing 4 Euro. In contrast, option B represents a guaranteed win of 14 Euro.

The risky option may be more attractive to an adolescent (AD) compared to a young adult (YA) because of:

- 1) adolescents show higher gain sensitivity ( $\alpha_{AD} > \alpha_{YA}$ ). That is, 20 Euro are subjectively more valuable to the adolescent; one possible interpretation of “reward sensitivity”,
- 2) adolescents show decreased loss sensitivity ( $\beta_{AD} < \beta_{YA}$ ). Thus, losing 4 Euro is subjectively less aversive to the adolescent,
- 3) reduced loss aversion ( $\lambda_{AD} > \lambda_{YA}$ ). That is, the psychological difference between gains and losses in adolescents is smaller compared to adults,
- 4) a combination of the above.

What this illustrates is that using EU models can help identifying the specific psychological mechanisms underlying developmental differences in risky choice (see also (van Duijvenvoorde et al., 2015) for a similar approach). As a result, they can also help to better understand how contextual modulations, such as peer presence (Albert, Chein, & Steinberg, 2013), impacts risky choice behavior (e.g., peer presence may specifically alter loss aversion). Indeed, several behavioral studies have already identified that there are significant age differences in both probability weighting (Harbaugh et al., 2002) and value functions (Blankenstein, Crone, van den Bos, & van Duijvenvoorde, 2016; Tymula et al., 2012; van den Bos & Hertwig, 2017) between adults and adolescents.

In most cases utility models are combined with the so-called softmax choice rule that assumes that expected utilities of the options under consideration are probabilistically translated into choices. For instance, the probability of choosing a risky option, when confronted with a safe and risky alternative, would be formalized as follows:

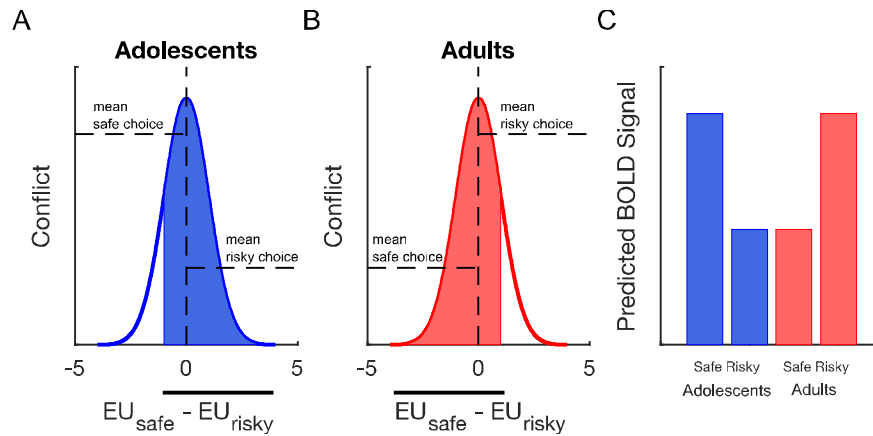
$$(2) \quad P_{risky} = \frac{1}{1 + e^{-\theta(EU_{risky} - EU_{safe})}}$$

where the single free parameter  $\theta$  in this function governs the sensitivity to differences in EU. When participants are less sensitive to EU differences, their preferences become less consistent. For instance, let us compare the hypothetical choice functions in Box 1. The blue curve for adolescents is clearly “flatter” compared to the red curve describing the adults. This flatness indicates a decreased sensitivity to differences in EU. Both groups will prefer the option with the highest EU but adolescent will be less consistent in their choice behavior.

As a result, even when two groups may have the same ‘risk aversion’ parameter ( $\alpha$ ) but different noise parameters ( $\theta$ , like the hypothetical adolescents) it is possible that one group will show a higher proportion of risky choices. This is strongly dependent on the choice set given to the participants. In Box 1, we have indicated a hypothetical choice set of gambles for which  $EU_{safe} - EU_{risk} = [-20, 10]$ . The shaded areas indicate the proportion of risky choices for



each group, given that choice set, the dotted lines indicate mean levels of risk taking. If we would just compare mean levels of risk taking in this hypothetical experiment, we would conclude that adolescents are more risk-seeking than young adults. Yet, in fact, their behavioral differences are due a diminished sensitivity to outcome differences. Importantly, note that with a different choice set, where  $EU_{\text{safe}} - EU_{\text{risk}} = [-10, 20]$ , the adolescents will still be closer to 50% risky choices but then will appear as more risk averse. Thus, this suggests that relying solely on choice proportions, or even just on EV models of risky choice can lead to substantial misinterpretations of age differences in risky decision making. Computational approaches using EU models can help to avoid these misconceptions, and can be used to identify the specific mechanism that contributed to changes in risk behavior.

**Box 2: Expected Utility and Risk Taking**

**A) & B)** Based on subjective utility differences the same choice set may result in mostly risky choices for adolescents and mostly safe choices for adults. Studies have shown that choices that are close to the indifference point (the point at which each option is equally preferred) are associated with increased reaction times (Krajchich, Bartling, Hare, & Fehr, 2015), which is thought to reflect increased choice conflict (Botvinick, 2007). As a result, the most infrequent choices for each age group are associated with increased levels of choice conflict. **C)** In this example one may predict that adolescents show more BOLD activation in conflict monitoring regions (e.g., dorsal anterior cingulate cortex) when choosing the safe option compared to adults. This can falsely be interpreted that in general adolescents need more “control” to avoid risk, although it is a specific effect due to the choice set that is used. This illustrates the usefulness of thinking about our choice experiments in terms of EU. A second implication is that what on an objective level seems like the same set of questions may be very different on the subjective level. This could result in rather imbalanced choice patterns introduce various unexpected and undesirable confounds in the experimental design and the post-hoc exclusion of subjects (e.g., those who chose the safe option over 90% of the time). A simple solution is to let participants perform a pre-test to estimate parameters of the utility model, and subsequently generate unique choice sets for each individual that are equally distributed around the EU indifference point (e.g., van den Bos et al., 2015).

*Implications for imaging*



The notion of EU has far reaching implications for neuroimaging studies. There is ample evidence that there are brain regions, like the ventral striatum and ventromedial prefrontal cortex, that track EU in choice experiments (Kable & Glimcher, 2007; Peters & Büchel, 2010). As we have pointed out above, EU may differ significantly from EV. Thus, if researchers use EV instead of EU as a parametric regressor in their imaging analyses of risky gambles, this will result in a better fit for those subjects for which EV and EU are most closely related (e.g., an  $\alpha$  close to 1). Using EV will therefore be potentially misleading if age groups differ in how close their EV function is to the EU function.

For example, let's assume that the utility curve in adolescents more closely resembles the EV of the presented decision options and, let's assume that activity in the ventral striatum tracks the EU (rather than the EV) of options. In an analysis that only considers the EV of options the result might be greater EV-related activity in the ventral striatum in teenagers compared to adults and the conclusion would be that adolescents show greater reward-sensitivity during risky decision-making. However, if we would run the same analysis with EU instead of EV we might find a very different result. Namely that there are no age differences in EU-related activity in the ventral striatum.

Of course, there are more complex risk tasks in which the probabilities and outcomes have to be learned by experience, such as the Balloon Analog Risk Task (BART) or the Iowa gambling task (IGT). It is possible that these paradigms have more predictive validity in terms of real world risk taking (Schonberg et al., 2012; van den Bos & Hertwig, 2017). However, these tasks partly depend on learning from experience and come with additional challenges in terms of computational modeling, which will be discussed below. Before turning to those issues, we will consider an alternative approach to decisions from description that focuses on the different strategies that individuals of different age might engage in when making choices

**The Strategy View**

The expected utility models discussed above represent an integrative approach that captures individual and developmental differences using specific parameterizations within a single model. The strategy view provides an alternative framework in that differences between individuals or across development are conceptualized as the use of fundamentally different strategies. Strategies can be thought of as sequences of operations or building blocks that can be combined to solve a particular task (Mata, Josef, & Lemaire, 2015). The strategy view has been used widely and successfully to understand the lifespan development of memory, arithmetic, as well as judgment and decision making (Mata, Pachur, et al., 2012; Siegler, 1999).

Gamble	Option A 	Option B 
<b>Hardmax</b> Choose A	$EV_A = 20 \cdot 0.5 - 2 \cdot 0.5 = 9$	$EV_B = 5$
<b>Maximax</b> Choose A	Max gain A: 20	Max gain B: 5
<b>Minimax</b> Choose B	Max loss A: 2	Max loss B: 0

**Figure 2.** When presented with a simple binary choice gamble there are different strategies. The hardmax choice rule deterministically chooses the option with the higher expected value. Other strategies, such as minimax and maximax, only use part of the information that is presented. Maximax tries to maximize to maximum possible gain and minimax tries to minimize the maximum possible loss. Although they may yield seemingly similar behavior, they rely on different mental processes. Furthermore, even though these strategies may sometimes be captured by a parametric model such as cumulative prospect theory parameterizations (Pachur, Suter, & Hertwig, 2017), they make different predictions about what happens on algorithmic and neural level (e.g., no representation of expected utility, no integration of possible outcomes).

There is a long-standing tradition in decision making research to distinguish between different strategies or heuristics (Shah & Oppenheimer, 2008). Two broad classes of strategies include compensatory strategies, that process all relevant information and consider possible trade-offs between attributes of decision options, and non-compensatory strategies, that avoid such trade-offs and typically reduce information processing demands by ignoring potentially relevant information (Gigerenzer & Gaissmaier, 2011; Payne, Bettman, & Johnson, 1988). One example of a non-compensatory strategy is a lexicographic strategy that simply selects the alternative that is best on the most important attribute (e.g. the best possible outcome regardless probability, see Figure 2). On the other hand, the previously described CPT is an example of a compensatory strategy that uses all possible information to come to a decision.

Formally, non-compensatory and compensatory strategies can be distinguished by how they assign weights to different cues in a decision problem, say deciding between two alternatives characterized by two or more cues (e.g., probability and outcome magnitude). Consider an ordered set of cues,  $C_1$  to  $C_M$ . A non-compensatory strategy is defined by assigned weights to each cue,  $W = \{w_1, w_2, w_3, \dots, w_M\}$ , in a manner that each weight is larger than the sum of the subsequent weights. For example, in a five cue example, the set  $\{1, 1/2, 1/4, 1/8, 1/16\}$ , fulfills the requirement that for every cue weight  $1 \leq j \leq M$  we have  $W_j > \sum_{k>j} W_k$ . These weights amount to making a decision whenever the first cue distinguishes between the two alternatives and ignoring all other remaining cues, because the subsequent weights cannot “overrule” the first.

Importantly, the strategy view typically emphasizes not only the specific cue weighting but also the importance of the sequences of processes. Therefore it will typically make predictions about other outcomes beyond choice, including search (Scholz, von Helversen, & Rieskamp, 2015), reaction times (Bröder & Gaissmaier, 2007) or neural processes (von Helversen, Karlsson, Rasch, & Rieskamp, 2014), that can further constrain developmental theories. These additional predictions can help to identify the underlying choice strategy that cannot be distinguished based on choice proportions or patterns alone (see also table 1).

The empirical work that has adopted the strategy approach suggests that compensatory and non-compensatory strategies may tap into different abilities, including the ability to inhibit irrelevant information or integrate many pieces of information that show important individual differences and developmental trends across the lifespan (Betsch, Lehmann, Lindow, & Lang, 2016; Huizenga et al., 2007; Mata, Helversen, & Rieskamp, 2011; Mata, Helversen, & Rieskamp, 2010; Mata, Schooler, & Rieskamp, 2007; Pachur, Mata, & Schooler, 2009). For example, young children seem to have difficulties using “simple” non-compensatory decision strategies due to deficits in cognitive control abilities that develop relatively late during childhood and that are necessary to inhibit accessible but irrelevant information (Mata et al., 2011). All in all, such results emphasize the importance of understanding the various strategies available to decision makers as well as the abilities that such strategies exploit which undergo substantial change during childhood and aging.

#### *Implications for imaging*

These age differences in learning and decision-strategies are important to take into account when analyzing neuroscience data. Knowing that different age groups may apply different strategies, and thus rely on different cognitive processes, will be extremely insightful in understanding different patterns of neural activity. This point is nicely illustrated by a recent fMRI study by van Duijvenvoorde and colleagues (van Duijvenvoorde et al., 2016). In this study participants were presented with a set of gambles designed to identify compensatory and non-compensatory strategies in risky choice. The results revealed that participants who applied a compensatory strategy showed a pattern of activity in the parietal cortex that reflected differences in EV between choice options. In contrast, for those who seemed to apply the non-compensatory strategy, their activation in the dorsomedial prefrontal cortex associated with greater conflict on the attribute level. If the authors would have used a single model to interpret the patterns of activity, it would have most likely led to a misguided interpretation.

The examples discussed above suggest that it may be helpful to distinguish different types of strategies when understanding the lifespan development of judgment and decision processes. The strategy approach, however, is not without empirical and conceptual challenges. For example, empirically, acknowledging heterogeneity in strategy use forces researchers to collect appropriately large samples to capture the cognitive and neural processes of each of the hypothesized strategy types. Conceptually, researchers face the problem of determining the space of hypothesized strategies a priori so as to avoid the problem of dealing with an intractable large number of strategies - the strategy sprawl problem (Scheibehenne, Rieskamp, & Wagenmakers, 2013). Fortunately, there are several possible approaches to deal with these challenges. We present an overview of the main problems, possible solutions, and some references to past exemplary work in Table 1.

Problem	Possible Solution(s)
<p><b>Strategy sprawl:</b> There is a potentially large number of combination of core processes or building blocks that can be combined which may lead to hypothesizing an intractably large number of strategies</p>	<ul style="list-style-type: none"> <li>• Use a pre-defined and constrained set of strategies that have been validated in past research (Mata, Schooler, &amp; Rieskamp, 2007)</li> <li>• Validate the use of several strategies using additional process data, such as search (Scholz et al., 2015), reaction time (Bröder &amp; Gaissmaier, 2007), or neural data (H. B. Fechner et al., 2016; Khader et al., 2011) (Fechner et al., 2016; Helversen et al., 2014; Khader, Pachur, Meier, Bien, &amp; Jost, 2011)</li> <li>• Adopt Bayesian modeling to help quantify the trade-off between flexibility in the number of hypothesized strategies and descriptive adequacy (Scheibehenne et al., 2013)</li> <li>• Explicitly model the mixture of strategies, and use parameter estimation to establish the contributions of individual strategies (Bays, Catalao, &amp; Husain, 2009; Collins &amp; Frank, 2012; van den Berg, Awh, &amp; Ma, 2014).</li> </ul>
<p><b>Identification problem:</b> Different strategies may often provide the same judgment or choice</p>	<ul style="list-style-type: none"> <li>• Design experimental paradigms so as to include critical tests that allow maximizing differences between hypothesized strategies (Mata et al., 2007; Mata, von Helversen, Karlsson, &amp; Cüpper, 2012; von Helversen, Mata, &amp; Olsson, 2010)</li> </ul>
<p><b>Individual heterogeneity:</b> Heterogeneity in strategy selection across individuals may lead to few individuals using the same strategy in a given situation</p>	<ul style="list-style-type: none"> <li>• Calculate power a priori based on hypothesized strategies and collect appropriately large samples</li> <li>• Collapse across individuals in a meaningful way, for example, across similar strategy types, such as compensatory and non-compensatory strategies (Mata, von Helversen, &amp; Rieskamp, 2010) or rule-based and similarity-based processes</li> <li>• Design the task to elicit specific strategies (Juslin, Jones, Olsson, &amp; Winman, 2003)</li> <li>• Instruct or train individuals to execute specific strategies (Mata, Wilke, &amp; Czienskowski, 2009; Siegler &amp; Lemaire, 1997)</li> </ul>
<p><b>Ecological rationality:</b> Different statistical environments (e.g., cue-criterion correlations) favor different strategies and individuals adapt their strategy use accordingly, potentially masking developmental effects</p>	<ul style="list-style-type: none"> <li>• Design the statistical structure of the task to elicit specific strategies (Mata et al., 2007)</li> <li>• Compare different statistical structures and assess individual or developmental differences in adaptivity (Horn, Ruggeri, &amp; Pachur, 2016)</li> </ul>
<p><b>Strategy execution:</b> Differences in judgment or choice outcomes may be due to differences in both strategy selection and strategy execution</p>	<ul style="list-style-type: none"> <li>• Adopt computational models that estimate strategy execution errors (Mata et al., 2010; Mata, von Helversen, &amp; Rieskamp, 2011)</li> <li>• Estimate developmental effects on strategy selection and execution directly using choice vs. no-choice method, that is, comparing experimental conditions in which individuals can select (choice) or simply</li> </ul>

---

execute (no-choice) particular strategies (Siegler & Lemaire, 1997)

---

## Reinforcement Learning Models

In the previous two paragraphs, we primarily focused on decisions from description, that is, tasks in which all the information that is necessary for making the decision is available (e.g., information about the value and the probability of an outcome, see Figure 1B and 2). In the world outside the laboratory such scenarios exist, for example, when making decisions between two different medical treatments with known risks and benefits. However, they are relatively rare; in most cases, we have to learn the expected value (EV) of choice options from experience (Hertwig & Erev, 2009). Moreover, we live in dynamically changing environments, which means that we have to constantly adjust our expectations. To do so, we often rely on trial and error learning processes, which undergo substantial developmental changes over the life course. In recent years, several research groups became interested in these processes and applied a range of experimental tasks (e.g., IGT, multi-armed bandit or reversal-learning tasks) to study different aspects of experiential learning. In most of these studies, researchers relied on descriptive summary statistics in their analyses (e.g., percentage of correct choices; (Cauffman et al., 2010; Denburg, Tranel, & Bechara, 2005; Eppinger, Kray, Mock, & Mecklinger, 2008; Eppinger et al., 2009; Hämmerer, Li, Müller, & Lindenberger, 2011; Hooper, Luciana, Conklin, & Yarger, 2004). However, the use of these performance measures may result in imprecise or even misleading interpretations of the underlying computational and neural mechanisms. For example, children and older adults often show similar limitations in learning under uncertainty. However, the computational mechanisms that lead to these learning impairments may differ substantially between groups (Hämmerer & Eppinger, 2012), which may not be evident from descriptive analyses alone. To study learning under uncertainty researchers often use probabilistic choice tasks in which participants have to learn the EV of different options based on probabilistic reward (e.g., reward in 80% versus 60% of the cases). In such tasks, similar performance levels in children and older adults may either emerge from estimating the EV of the choice option too rapidly, that is, by ignoring the recent outcome history or by learning too slowly, i.e., by considering too much of the outcome history (Nassar, Wilson, Heasly, & Gold, 2010). From a psychological point of view, learning too rapidly can be described as a tendency to change expectations about choice options too quickly. In the above described choice task example, this means that a participant is heavily influenced by each outcome and too eager to shift to a different choice option when the outcome does not match the expectation. In contrast, learning too slowly means that outcomes of decisions are not sufficiently considered. Thus, the learner tends to ignore the feedback. Alternatively, similar performance levels in children and older adults could be due to opposing exploration strategies. For example, consider yourself grocery shopping in a super market in a new (unknown) country. You will have to decide between various different types of e.g. cereals. Without any prior knowledge, the best thing that you can do is to sample (explore) the different options that you have. However, exploring too much results in choosing too many low-reward options, whereas exploring too little can lead to missing out on preferable options (Wilson, Geana, White, Ludvig, & Cohen, 2014). What these examples mean to illustrate is that being able to assess these different updating and exploration strategies may allow us to better understand and disambiguate lifespan age differences in learning and decision-making.

As outlined in (Hämmerer & Eppinger, 2012) it may also be that different underlying neurobiological mechanisms contribute to the seemingly similar performance profiles in children and older adults. As such, the observation that different age groups show the same performance impairments may not be sufficient to conclude that these are due to changes in similar neurobiological processes. As we argued above, computational models may help



researchers to develop and test more specific theories about the mechanisms underlying developmental differences in learning and decision making and to identify the cognitive and neurobiological processes from which these differences emerge. In the following, we describe a basic computational implementation of reinforcement learning (RL), the so-called Rescorla-Wagner model and how it can be applied to study learning processes across development and aging (Daw, 2011, 2014; Niv & Schoenbaum, 2008; Sutton & Barto, 1998). RL theory offers formal models for learning from interaction with the environment when an individual has no direct instructions as to what actions to take. Accordingly, previous experiences of reward are used to form expectations about outcomes of future choices.

The Rescorla-Wagner model (and its many derivatives) relies on a simple principle of updating expectations based on prediction errors (Daw, 2011; Sutton & Barto, 1998). This model is often applied in simple tasks as in the example above (Figure 3A), where participants are required to repeatedly choose one of two options (e.g., blue (B) versus red (R)) that provide probabilistic feedback, with the goal of maximizing rewards.

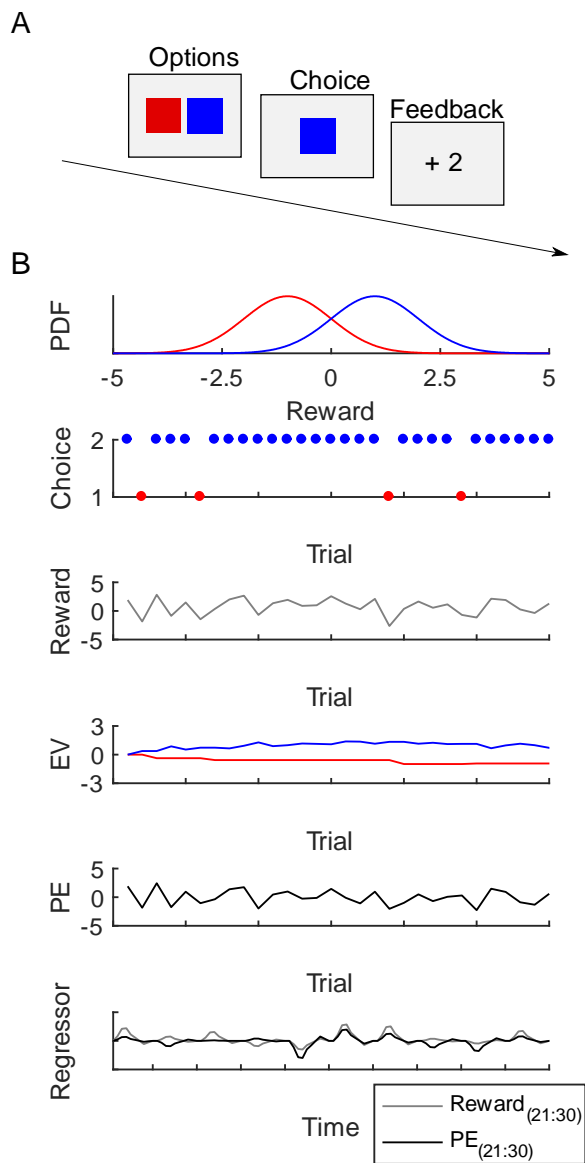
We assume that participants start with a certain expectation about the two options (e.g., the EV is 0 for both B and R (Figure 3B)) and subsequently use trial and error information to approximate the EV of each option. The basic idea of the Rescorla-Wagner model is that value expectations are sequentially updated based on the difference between the EV and the received reward: the prediction error  $\delta$

$$(3) \quad \delta_t = r_t - Q_t(c_t)$$

where  $r_t$  is the received reward at time  $t$  and  $Q_t(c_t)$  is the EV of choice  $c_t$  (B or R) at time  $t$  (see Figure 3B). This prediction error will subsequently be used to update the EV associated with choosing option B or R

$$(4) \quad Q_{t+1}(c_t) = Q_t(c_t) + \alpha \cdot \delta_t.$$

Thus, in this algorithm, the EV is updated in the direction of the prediction error to improve the accuracy of expectations. The prediction error is multiplied by a learning rate  $\alpha$  to scale the influence of the prediction error on the updated stimulus value. A high learning rate ( $\sim 1$ ) will lead to an update in favor of the most recent outcomes whereas with a low learning rate ( $\sim 0$ ) the stimulus value is less affected by the outcome. Thus, the learning rate parameter specifies to what extent new outcomes affect reward expectations and can dissociate participants who learn rapidly from participants who learn slowly (see example above). In addition, in most of the learning scenarios you have two or more options that you can choose between (as in the example above). Thus, in order to make a decision you have to compare the value of the options and translate these values into choice probabilities. In many RL applications, this is achieved with the softmax function as described above (eq. 2). This parameter is often called exploration term and thus determines the degree to which the choice options are explored (see grocery shopping example above).

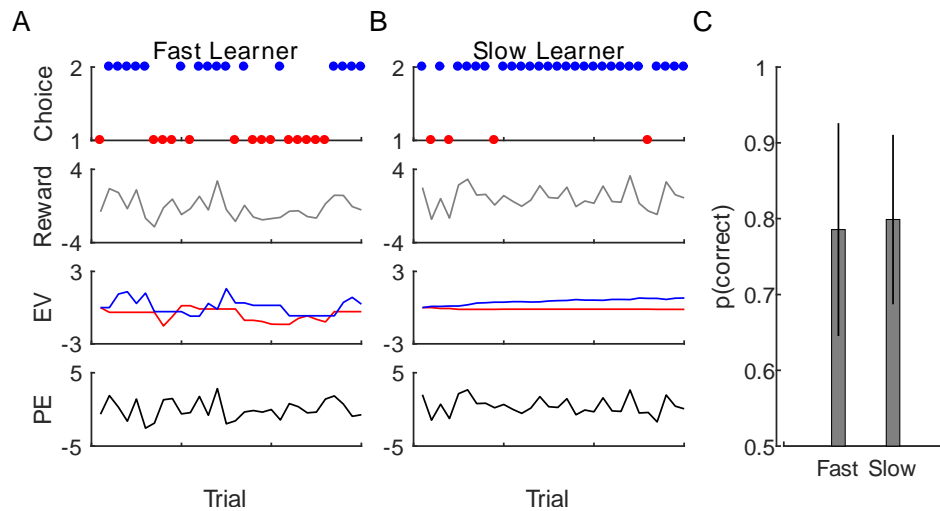


**Figure 3.** Reinforcement learning (RL) and model-based fMRI analyses. **A)** Example of a two-armed bandit task. Participants are required to choose between the blue and red option, which is followed by a reward or a punishment. **B)** Reward distribution of the blue and the red option and application of an RL model. On average, the blue option is associated with reward, the red option is associated with punishment. Choices indicate that the model is able to learn that the blue option is associated with higher reward. As a consequence of a preference for the blue option, the model receives rewards on most trials. Across trials, the model approximately learns the expected value (EV) of both options. The prediction error (PE) indicates the difference between received rewards and EVs and can be utilized to adjust EVs. Finally, model-parameters can be used as parametric regressors in neuroimaging analyses. Note that the predicted BOLD signal of rewards and PEs can go in opposite directions.

To illustrate how RL models can be used to dissociate behavioral performance profiles in different groups we can simulate behavior of a fast and a slow learner (a learner with a high or low learning rate): Figure 4 shows behavior of the two RL models (Figure 4A,B) and the average performance (Figure 4C) across simulations in the above described two-armed bandits task. The comparison is based on 500 simulations each consisting of 30 trials. Both models have an exploration parameter  $\theta = 5$ . However, behavior of model 1 (Figure

4A) was generated using a high learning rate ( $\alpha = 0.6$ ) with the effect that the EV of model 1 fluctuates wildly. Model 2 (Figure 4B), in contrast, uses a low learning rate ( $\alpha = 0.05$ ), which consequently leads to a slow change in the EV. As can be nicely seen, although the underlying parameters are clearly different, the average performance of both models is similar.

This simple example demonstrates both promises and pitfalls of computational modeling. It suggests that observed behavioral effects can potentially be generated by different combinations of parameters that may be captured using computational models. However, it also demonstrates the need for careful experimental design: To reliably identify parameter differences between groups or individuals, one needs to carefully think about experimental task conditions that allow to dissociate these parameters (see also Box 3).



**Figure 4.** Simulations using a reinforcement learning (RL) model with different learning rate but equal exploration parameters. **A)** Rapidly learning RL model with a learning rate  $\alpha = 0.6$  and exploration term  $\theta = 5$ . **B)** Slowly learning model with a learning rate  $\alpha = 0.05$ . **C)** Although models have different learning rate parameters, mean performance across 500 simulations with 30 trials each is similar. Error bars represent the standard deviation of the mean between the simulations.

Taken together, one of the advantages of computational RL approaches for developmental science is that they allow us to get access to latent variables (such as learning rates or exploration parameters) that cannot be accessed with descriptive approaches alone. These variables may help us to disambiguate behavioral profiles of different age groups and may provide us with a better mechanistic understanding of developmental differences.

#### *Implications for imaging*

Beyond global parameters such as learning rate or exploration parameter, which are estimated on the individual subject level, RL models can also be used to derive trial-by-trial estimates of two other latent variables: the reward prediction error  $\delta_t$  and the EV  $Q_{t+1}$  of the choice options (actions). When performance differences in learning occur between two different age groups, an obvious question of computational interest is whether these differences are associated with changes in reward prediction error signaling or the representation of EV in different cortical or subcortical areas. In order to answer this question, it is necessary to construct an fMRI design matrix that models the task events (e.g., decision, outcome) and includes a term that reflects the extent to which an event response is modulated by the parametric variable (the prediction error or EV, see Figure 3B). Such a design matrix can be fit to the fMRI data using standard analysis

software and the resulting coefficients for these modulator terms provide a quantitative measure of the relationship between the BOLD signal and the latent variable of interest (e.g., reward prediction error). Thus, in principle this measure can be used to test for differences in neural computations underlying learning (Figure 3B).

Several recent studies have taken advantage of these approaches and show evidence for differences in the correlation between model-derived prediction errors and BOLD activity in the ventral striatum in children and adolescents (Christakou et al., 2013; Hauser, Iannaccone, Walitza, Brandeis, & Brem, 2014; Javadi, Schmidt, & Michael, 2014; van den Bos et al., 2012) as well as older adults (Chowdhury et al., 2013; Eppinger, Schuck, Nystrom, & Cohen, 2013; Samanez-Larkin, Worthy, Mata, McClure, & Knutson, 2014). Most of this work has focused on model-free RL, using for example the Rescorla-Wagner model described above. Recently, however, there is an increasing interest in more complex types of RL (“model-based” learning), which involve learning of a forward model of the environment that can be used for planning (Daw, Niv, & Dayan, 2005; Dolan & Dayan, 2013). Model-based learning may be advantageous, especially in complex environments, because it allows more sophisticated behavioral strategies than model-free learning. However, this advantage comes at the cost of higher demands on, for example, working memory and attention, which makes it an interesting target for research in lifespan cognitive neuroscience (Decker, Otto, Daw, & Hartley, 2016; Eppinger, Walter, Heekeren, & Li, 2013).

In a recent age-comparative study, Nassar and colleagues (Nassar et al., 2016) used RL principles in combination with Bayesian methods to model learning dynamics in uncertain and changing environments. Using a predictive inference task, the authors examined age-related changes in the factors that affect trial-to-trial adjustments of learning rates. The results suggested that age-related learning deficits in older adults are due to a specific deficit in representing uncertainty. This deficit may not directly affect prediction error signaling but rather the computation of the learning rate, which, as described above, determines the degree to which prediction errors are considered during learning.

To conclude, RL offers a theoretical framework to study learning and decision making across the lifespan. The key advantage of these computational models is that they allow us to estimate latent variables such as prediction errors or the learning rate that cannot be assessed using descriptive models. Another important advantage is that RL models can be used to simulate the impact of developmental differences on behavior (see (Nassar et al., 2016; Palminteri, Wyart, & Koehlin, 2016)). This leads to a considerable increase in the specificity of predictions regarding developmental differences in learning and decision processes. Finally, model parameters can be used to inform fMRI or EEG analyses which may allow to identify the neurobiological mechanisms underlying developmental changes in learning. However, despite all these potential advantages, in practice, there are also many technical hurdles that have to be tackled. In Box 3 we focus on a few of the issues that are particularly relevant for developmental comparisons.

### Box 3: Computational Modeling of Lifespan Differences

As outlined above there are several advantages of the application of computational models in developmental cognitive neuroscience studies. However, in practice, there are also many technical hurdles that can impede these lofty research goals. In the following section, we will focus on a few of these issues that are particularly relevant for developmental comparisons.

#### *Model comparison*

Identifying a suitable model of behavior generally requires comparing the fit of several candidate models. For example, several researchers compared RL models with a single learning rate and models with a separate learning rate for positive and negative outcomes (van den Bos et al., 2012). Intuitively, model comparison generally involves some assessment of how likely candidate models would be to generate the observed behavior; however, another important consideration is model complexity. Models that have more freedom to explain nuances in behavior (and nuances in measurement noise) should be penalized for any performance gains achieved through overfitting (capturing variability in the noise). This is often done using penalized likelihood-based metrics such as the Akaike information criterion (AIC) and Bayesian information criterion (BIC), which penalize complexity based on parameter number and can be appropriate under a circumscribed set of assumptions. A more robust approach is to compute out-of-sample likelihoods through k-fold cross-validation or leave-one-out procedures (Friedman, Hastie, & Tibshirani, 2001). Out-of-sample techniques are robust to non-independence of data across trials and subjects, making them ideally suited for model comparison in hierarchical models where dependencies between subjects are modeled directly (Vehtari, Gelman, & Gabry, 2016). In general, out-of-sample methods are justified under a wider range of conditions and can capture more nuanced forms of complexity. In either case, best practice is to always validate the model using *posterior-predictive checking* to verify that the “best fitting” model is likely to generate behavioral data that look like those produced by human subjects (see below).

Note also that model comparison is another way of explicitly testing different theories against each other given the data, which is, as we argued before (van den Bos & Eppinger, 2016) very hard to do with descriptive models. As pointed out in Table 1 (identification) simulations can also be used a priori to check in advance whether an experimental setting is even capable of distinguishing between selected models (Mata et al., 2007; Mata, von Helversen, et al., 2012; von Helversen et al., 2010).

#### *Posterior-predictive checking*

One of the first steps in computational analyses is to find out whether the model that we have used fits the patterns of behavior shown by the participants. Just looking at the parameter values or goodness of fit will not reveal this; an insufficient model can provide highly misleading parameterizations of behavior (Nassar & Gold, 2013; Palminteri et al., 2016). During posterior-predictive checking (Box, 1980; Gelman, Meng, & Stern, 1996), we use the best fitting computational models to simulate data. These simulated data are then submitted to the same descriptive analyses that have been used to analyze the acquired (participant) data (Nassar & Frank, 2016). In some cases, this procedure might reveal that the effects observed in the simulated data closely resemble those observed in subject’s data. In other cases, the comparison may reveal a shortcoming of the model that suggests that the model architecture needs to be changed in order capture the data.

*Group differences in model fit*

A relevant concern for lifespan research is the possibility that two groups of subjects (e.g., younger and older adults) might be best described by different computational models. First, this may reflect true group differences in strategy use (which may be the result you were expecting). In this case, it will not be very meaningful to select a single model, for instance the one that fits the combined data of the groups the best, and pretend that it is the generative model for all participants. However, in some cases it could be that there exists a meaningful hybrid model that includes the features necessary to describe each group, mixed with a single parameter that governs the relative contributions of each model (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Nassar et al., 2010). For example, recent age-comparative studies using such hybrid models showed evidence for qualitative changes in model-free and model-based learning mechanisms across the human lifespan (Decker et al., 2016; Eppinger, Walter, et al., 2013).

*Individual versus group based parameter estimation*

After model selection, the free parameters of the model of interest (such as the learning rate and exploration term) are generally estimated in order to obtain trial-wise latent variables (e.g., prediction error) and to identify parameter differences between groups. On one extreme, models could be fit individually to each subject, and trial-wise latent variables could be extracted according to the model that best fits that particular subject. On the other extreme, the model could be run across data for all subjects using a set of “average” parameters, such that trial-wise latent variables are computed in the same way for each subject. Each of these approaches has its shortcomings; fitting to individual subjects leads to imprecision from limited trial data, whereas group average parameters may not provide a particularly good characterization of any particular subject and thus result in the selection of the wrong generative model (Wulff & van den Bos, 2017). A compromise between these methods is to use a hierarchical approach, such that estimates for individual subjects are constrained by their own data, as well as a group level prior that incorporates data from other subjects. In this way, parameter estimates for all subjects, but particularly those for subjects with limited or highly variable behavioral measurements, will be pulled toward the mean parameter estimate across the entire group “shrinking” the range of parameter estimates. In principle, this can be repeated at multiple levels of analysis, for example, in a nested model that includes a global prior distribution from which parameters for individual age groups are sampled. Individual parameters are then sampled from the lowest level such that they are pulled toward the mean of their own parameter distribution as well as the global mean. Note that the hierarchical approach does require making explicit modeling decisions about the prior distribution and form of hierarchy, and thus care should be taken to ensure that results are not overly conditioned on these choices (Gelman et al., 2014). That said, non-hierarchical models also make implicit assumptions about prior distributions (improper uniform) and relationships between subjects (either zero or full pooling) that could affect results.

*Linking latent computational variables to indirect neural measurements*

Once parameters are estimated, trial-wise parameters can be extracted from the properly parameterized model and linked to indirect measures of neural function such as fMRI or EEG data. This can be done by including trial-wise latent variables from the computational model as an explanatory variable in a GLM to explain trial-wise physiological measurements. One important consideration in this process is how to normalize trial-wise variables; if using either individual subject or hierarchical fitting approaches, it is important to normalize the variance (z-score) of the trial-wise latent variables for each individual subject, to ensure that detected

differences in regression coefficients are not simply picking up on differences in scale emerging from the fitting process.

#### *Parameter recovery*

In order to interpret differences in the best fitting parameters across groups, it is important to show that the model fitting procedures employed are capable of identifying meaningful and stable differences in any of the parameters. For instance, when there are interactions between parameters, it is possible that there are many different parameter settings that generate the same pattern of behavior. Simulating data from models that take a range of different combinations of parameter values and attempting to recover those “true” parameters by fitting the model directly to the data it generated will be informative. Only if these analyses reveal that the simulated and recovered parameter values are related in a systematic and continuous fashion they can be used for meaningful age comparisons.

#### *Linking computational changes to differences in behavior across developmental groups*

Once a computational model has been identified, a key question is whether the constructs captured by its free parameters (e.g. learning rate) differ systematically across groups. When data of individual subjects are fit separately (e.g., by choosing the model parameters that maximize the likelihood of the individual data), this can simply be done by comparing the group median parameter fits (e.g., using a Wilcoxon signed rank test). When interpreting a parameter difference, it is often useful to know whether behavioral differences are selective to a single computational factor (e.g., learning rate for gains but not for losses). To do so, it is important to remember that the existence of a statistical difference in one parameter and the lack of a difference in another parameter does not necessarily imply selectivity; claims of specificity should rely on explicit comparisons of age differences across the different parameters (Nieuwenhuis, Forstmann, & Wagenmakers, 2011). In a hierarchical model, parameter differences between groups can be estimated by examining the posterior distribution over group differences directly, or by computing credible intervals over likely parameter differences.

## General Discussion

In this paper, we have provided an overview about how computational models of cognition can be used to study age-related changes in psychological processes and the underlying neurobiological mechanisms across the human lifespan. The core idea of this neuro-computational approach is that formalized models can provide us with mechanistic links between verbal descriptions of behavior and its cognitive and neurobiological implementations.

Specifically, we propose that computational models allow us to address two current problems in developmental cognitive neuroscience: 1) The problem of making predictions about behavior and the underlying neurobiology that are specific enough to falsify verbal theories (*specificity problem*) and 2) The problem of capturing the identity of developmental processes (*identity problem*). We have shown applications of computational neuroscience approaches in two major domains of decision making: decisions from description and decisions from experience. Here we have focused on the use of models to describe that behavior, note that there is already a wealth of developmental neuroscience studies on these topics (Hartley & Somerville, 2015).

In both of these domains computational approaches in combination with neuroimaging can significantly advance our mechanistic understanding of the underlying processes. This has led to the development of new fields such as computational psychiatry which aims to provide a mechanistic understanding of psychiatric disorders that can guide theory-based clinical interventions (Ahn & Busemeyer, 2016; Huys, Maia, & Frank, 2016). Many of the disorders under study (such as schizophrenia, anxiety, addiction or ADHD) develop during late childhood and adolescence. Therefore, any type of complete theory utilized to explain these disorders has to incorporate a normative perspective of human development. The same is true for aging-related diseases such as Parkinson's disease or dementia. Yet, we are far away from such neuro-computational theories of development and aging. The aim of this work is to provide a starting point for the development of such theories and to encourage researchers to adopt neuro-computational approaches.

It is obvious that there are important questions and research strategies that we have not covered. For example, there is an emerging literature on developmental differences in perceptual decision making and the use of drift diffusion models to discern different perceptual decision making processes (Ratcliff, Thapar, & McKoon, 2006; Schuch & Konrad, 2017; Spaniol, Madden, & Voss, 2006; Thompson, Ratcliff, & McKoon, 2016). Furthermore, there is an increasing interest in studying how the effort that is involved in making a decision affects choice behavior. So far there are only a few studies on age differences in effortful decision making (Benozio & Diesendruck, 2015; Westbrook, Kester, & Braver, 2013) and the existing computational approaches (Kool, Cushman, & Gershman, 2016; Shenhav, Botvinick, & Cohen, 2013) have not (yet) been implemented. We have not addressed a core dilemma in decision making, the question how we solve the trade-off between exploration and exploitation and how this changes across development (Somerville et al., 2017; Wilson et al., 2014). Finally, we also ignored the extensive literature on neural network modeling of development (Mareschal & Shultz, 1996; Plunkett, Karmiloff-Smith, Bates, & Johnson, 1997). These models provide (even) more complex perspectives on cognitive development and aging and may be well suited to address the identity problem.

With respect to the neuroscience approaches we focused on fMRI, but there are new ways of analyzing EEG data using single trial approaches which are promising (Fischer & Ullsperger, 2013) and there are several other recent examples of extremely fruitful



combinations of computational modeling and psychophysiological measures (Cavanagh, Wiecki, Kochar, & Frank, 2014; Nassar et al., 2012).

As outlined there are several potential pitfalls when applying computational models to behavioral data and in using the outcomes of modeling to inform neuroscience data. Some of these pitfalls are specific to developmental research questions, others are more general and our review of these issues is certainly not exhaustive (for more detailed descriptions please refer to (Daw, 2011; Nassar & Frank, 2016; Palminteri et al., 2016) and to (Redish & Gordon, 2016) for an overview of uses in psychiatry). Given the increasing interest in computational neuroscience methods it is important to note that there are several non-trivial inferential problems regarding 1) whether and to which degree a model actually fits the data and 2) what correlations between computational parameters and neurobiological signals actually reflect. The latter point refers to the fact that with the current neuroimaging approaches we cannot make causal inferences about model parameters and neurobiological signals. That is, even though it may be tempting to assume that, for example, learning deficits in older adults are due to diminished striatal prediction error signals, the relationship still remains a correlative one and typically activity in several other areas also correlates with prediction errors (Hayden, Heilbronner, Pearson, & Platt, 2011; Schultz & Dickinson, 2000). To tackle the question of causality, we will have to rely on non-invasive brain stimulation methods (e.g., rTMS) or pharmacological manipulations (Chowdhury et al., 2013), and engage in cross-species comparisons (e.g., involving optogenetic manipulations in rodents). In addition, the correlation (or lack thereof) between computational parameters and neuroscience data can also be used in model selection itself. Often the computational model is fit on the behavioral data alone, but to the extent that the model makes predictions about neural processes, model selection can also be (simultaneously) constrained by neural data. This is an exciting future direction in computational neuroscience that can further our confidence in the identification of cognitive processes underlying life span changes in behavior (Turner et al. 2013).

We are aware of the valid concern that computational neuroscience approaches may lead to a segmentation of behavior into “molecular” psychological and neurobiological processes that are, at some point, far removed from the behavior that was originally set out to be explained. That is, the jargon associated with increasingly complex computational models may become meaningless if the models cannot speak to existing psychological theories of behavior. It is therefore crucial to use our increasing understanding of the neuro-computational mechanisms to answer the question of *why* humans of different ages behave in a certain way, by finding ways of linking the algorithmic level to back the level of verbal theories.

Finally, it is important to point out that computational modeling cannot replace good experimental design, but rather can inform it. That is, the value of a computational model is constrained by the value of the experimental design that it is associated with. Even if the model could perfectly explain the behavior on a specific task, and the accompanying neural processes, it will be of limited value if this task has no external validity. For instance, adolescents probably only rarely encounter “risky” choices where they are presented with full information about probabilities and outcomes such as is often done in monetary gambles (see Figure 1B). Thus, one may wonder how informative age differences in parameters of pure risky choice models will be (cf. (van den Bos & Hertwig, 2017)). In addition, experimental tasks are often simplified in order to capture one specific feature of the real-world environment or one isolated psychological process. This strategy may limit external validity because it may miss out on crucial complexities that explain real world behavior. Having a good computational model will provide a framework for understanding how multiple variables interact (and change over time) and therefore may allow for the design of more complex, externally valid experiments.

To conclude, we think that neuro-computational approaches have a tremendous potential for studying human development across the lifespan. Computational methods can provide access to latent processes that are not accessible with descriptive methods and may thus foster the development of mechanistic theories of normative development. This allows specific predictions about brain behavior relationships in different developmental groups (specificity problem) and may enable us to identify the nature of developmental processes (identity problem).

## Acknowledgments

We would like to thank Joni Shuchat and Robert Lorenz for comments on previous versions of this manuscript. This research was in part funded by grants of the German Research Foundation (Deutsche Forschungsgesellschaft), SFB 940, subproject B7 (B.E.), and the Open Research Area (ORA), ASTA 176 grant (W.B.)

**References**

- Ahn, W. Y., & Busemeyer, J. R. (2016). Challenges and promises for translating computational tools into clinical practice. *Current Opinion in Behavioral Sciences*, *11*, 1–7. <http://doi.org/10.1016/j.cobeha.2016.02.001>
- Albert, D., Chein, J., & Steinberg, L. (2013). The teenage brain: Peer influences on adolescent decision making. *Current Directions in Psychological Science*, *22*(2), 114–120. <http://doi.org/10.1177/0963721412471347>
- Amso, D., & Scerif, G. (2015). The attentive brain: Insights from developmental cognitive neuroscience. *Nature Reviews Neuroscience*, *16*(10), 606–619. <http://doi.org/10.1038/nrn4025>
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7.1–11. <http://doi.org/10.1167/9.10.7>
- Benozio, A., & Diesendruck, G. (2015). From effort to value: Preschool children's alternative to effort justification. *Psychological Science*, *26*(9), 1423–1429. <http://doi.org/10.1177/0956797615589585>
- Blankenstein, N. E., Crone, E. A., van den Bos, W., & van Duijvenvoorde, A. C. K. (2016). Dealing with uncertainty: Testing risk- and ambiguity-attitude across adolescence. *Developmental Neuropsychology*, *41*(1–2), 77–92. <http://doi.org/10.1080/87565641.2016.1158265>
- Botvinick, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cognitive, Affective, & Behavioral Neuroscience*, *7*(4), 356–366. <http://doi.org/10.3758/CABN.7.4.356>
- Box, G. E. P. (1980). Sampling and Bayes' inference in scientific modelling and robustness. *Journal of the Royal Statistical Society. Series A (General)*, *143*(4), 383–430. <http://doi.org/10.2307/2982063>
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin & Review*, *14*(5), 895–900. <http://doi.org/10.1093/acprof:oso/9780199744282.003.0020>
- Cauuffman, E., Shulman, E. P., Steinberg, L., Claus, E., Banich, M. T., Graham, S., & Woolard, J. (2010). Age differences in affective decision making as indexed by performance on the Iowa Gambling Task. *Developmental Psychology*, *46*(1), 193–207. <http://doi.org/10.1037/a0016128>

- Cavanagh, J. F., Wiecki, T. V., Kochar, A., & Frank, M. J. (2014). Eye tracking and pupillometry are indicators of dissociable latent decision processes. *Journal of Experimental Psychology: General*, *143*(4), 1476–1488.  
<http://doi.org/10.1037/a0035813>
- 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. <http://doi.org/10.1038/nn.3364>
- Christakou, A., Gershman, S. J., Niv, Y., Simmons, A., Brammer, M., & Rubia, K. (2013). Neural and psychological maturation of decision-making in adolescence and young adulthood. *Journal of Cognitive Neuroscience*, *25*(11), 1807–1823.  
<http://doi.org/10.1162/jocn>
- Collins, A. G. E., & Frank, M. J. (2012). How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *European Journal of Neuroscience*, *35*(7), 1024–1035.  
<http://doi.org/10.1111/j.1460-9568.2011.07980.x>
- Crone, E. A., van Duijvenvoorde, A. C. K., & Peper, J. S. (2016). Annual research review: Neural contributions to risk-taking in adolescence-developmental changes and individual differences. *Journal of Child Psychology and Psychiatry*, *57*(3), 353–368. <http://doi.org/10.1111/jcpp.12502>
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. In M. R. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Decision Making, Affect, and Learning: Attention and Performance XXIII* (pp. 3–38). Oxford: Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199600434.003.0001>
- Daw, N. D. (2014). 16 Advanced Reinforcement Learning. *Neuroeconomics*, 299–320. <http://doi.org/10.1016/B978-0-12-416008-8.00016-4>
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215. <http://doi.org/10.1016/j.neuron.2011.02.027>
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711. <http://doi.org/10.1038/mn1560>
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychological Science*, *27*(6), 848–858.

- <http://doi.org/10.1177/0956797616639301>
- Denburg, N. L., Tranel, D., & Bechara, A. (2005). The ability to decide advantageously declines prematurely in some normal older persons. *Neuropsychologia*, *43*(7), 1099–106.  
<http://doi.org/10.1016/j.neuropsychologia.2004.09.012>
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*(2), 312–325. <http://doi.org/10.1016/j.neuron.2013.09.007>
- Eppinger, B., Kray, J., Mock, B., & Mecklinger, A. (2008). Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia*, *46*(2), 521–39.  
<http://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), 1043–53.  
<http://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. *The Journal of Neuroscience*, *33*(24), 9905–12.  
<http://doi.org/10.1523/JNEUROSCI.2942-12.2013>
- Eppinger, B., Walter, M., Heekeren, H. R., & Li, S.-C. (2013). Of goals and habits: Age-related and individual differences in goal-directed decision-making. *Frontiers in Neuroscience*, *7*(December), 253.  
<http://doi.org/10.3389/fnins.2013.00253>
- Fechner, G. (1966). *Elements of Psychophysics, Vol I (1860)*. Translated by Adler HE.
- Fechner, H. B., Pachur, T., Schooler, L. J., Mehlhorn, K., Battal, C., Volz, K. G., & Borst, J. P. (2016). Strategies for memory-based decision making: Modeling behavioral and neural signatures within a cognitive architecture. *Cognition*, *157*, 77–99. <http://doi.org/10.1016/j.cognition.2016.08.011>
- Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron*, *79*(6), 1243–55. <http://doi.org/10.1016/j.neuron.2013.07.006>
- Friedman, J., Hastie, T., & Tibshirani, R. (2001). The elements of statistical learning. Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian Data Analysis (Vol. 2)*. Boca Raton, FL: CRC press.
- Gelman, A., Meng, X.-L., & Stern, H. (1996). Posterior predictive assessment of

- model fitness via realized discrepancies. Vol.6, No.4. *Statistica Sinica*, 6(4), 733–807. <http://doi.org/10.1.1.142.9951>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–482. <http://doi.org/10.1146/annurev-psych-120709-145346>
- Hämmerer, D., & Eppinger, B. (2012). Dopaminergic and prefrontal contributions to reward-based learning and outcome monitoring during child development and aging. *Developmental Psychology*, 48(3), 862–74. <http://doi.org/10.1037/a0027342>
- 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, 579–592. <http://doi.org/10.1162/jocn.2010.21475>
- Harbaugh, W. T., Krause, K., & Vesterlund, L. (2002). Risk Attitudes of Children and Adults: Choices Over Small and Large Probability Gains and Losses. *Experimental Economics*, 5(1), 53–84. <http://doi.org/10.1023/A:1016316725855>
- Hartley, C. A., & Somerville, L. H. (2015). The neuroscience of adolescent decision-making. *Current Opinion in Behavioral Sciences*, 5, 108–115. <http://doi.org/10.1016/j.cobeha.2015.09.004>
- Hauser, T. U., Iannaccone, R., Walitza, S., Brandeis, D., & Brem, S. (2014). Cognitive flexibility in adolescence: Neural and behavioral mechanisms of reward prediction error processing in adaptive decision making during development. *NeuroImage*. <http://doi.org/10.1016/j.neuroimage.2014.09.018>
- Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., & Platt, M. L. (2011). Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 31(11), 4178–4187. <http://doi.org/10.1523/JNEUROSCI.4652-10.2011>
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517–523. <http://doi.org/10.1016/j.tics.2009.09.004>
- Hooper, C. J., Luciana, M., Conklin, H. M., & Yarger, R. S. (2004). Adolescents' performance on the Iowa Gambling Task: implications for the development of decision making and ventromedial prefrontal cortex. *Developmental Psychology*,

- 40(6), 1148–58. <http://doi.org/10.1037/0012-1649.40.6.1148>
- Horn, S. S., Ruggeri, A., & Pachur, T. (2016). The development of adaptive decision making: Recognition-based inference in children and adolescents. *Developmental Psychology*, *52*(9), 1470–1485. <http://doi.org/10.1037/dev0000181>
- Huys, Q. J. M., Maia, T., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, *19*, 404–413.
- Javadi, A. H., Schmidt, D. H., & Michael, N. S. (2014). Adolescents adapt more slowly than adults to varying reward contingencies. *Journal of Cognitive Neuroscience*, *26*(12), 2670–2681.
- Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(5), 924–941. <http://doi.org/10.1037/0278-7393.29.5.924>
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625–1633. <http://doi.org/10.1038/nn2007>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, *47*, 263–292.
- Khader, P. H., Pachur, T., Meier, S., Bien, S., Jost, K., & Rösler, F. (2011). Memory-based decision-making with heuristics: evidence for a controlled activation of memory representations. *Journal of Cognitive Neuroscience*, *23*(11), 3540–54. [http://doi.org/10.1162/jocn\\_a\\_00059](http://doi.org/10.1162/jocn_a_00059)
- Kool, W., Cushman, F. A., & Gershman, S. J. (2016). When does model-based control pay off? *PLOS Computational Biology*, 1–34.
- Krajbich, I., Bartling, B., Hare, T., & Fehr, E. (2015). Rethinking fast and slow based on a critique of reaction-time reverse inference. *Nature Communications*, *6*, 1–9. <http://doi.org/10.1038/ncomms8455>
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., Maciver, M. A., & Poeppel, D. (2017). Neuroscience needs behavior: Correcting a reductionist bias. *Neuron*, *93*(3), 480–490. <http://doi.org/10.1016/j.neuron.2016.12.041>
- Li, S.-C., & Rieckmann, A. (2014). Neuromodulation and aging: Implications of aging neuronal gain control on cognition. *Current Opinion in Neurobiology*. <http://doi.org/10.1016/j.conb.2014.07.009>



- Mareschal, D., & Shultz, T. R. (1996). Generative connectionist architectures and constructivist cognitive development. *Cognitive Development, 11*, 59–88.
- Mata, R., Josef, A. K., & Lemaire, P. (2015). *Adaptive Decision Making and Aging. Aging and decision making: Empirical and applied perspectives.*  
<http://doi.org/10.1016/B978-0-12-417148-0.00006-6>
- Mata, R., Pachur, T., von Helversen, B., Hertwig, R., Rieskamp, J., & Schooler, L. (2012). Ecological rationality: A framework for understanding and aiding the aging decision maker. *Frontiers in Neuroscience, 6*(FEB).  
<http://doi.org/10.3389/fnins.2012.00019>
- Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology and Aging, 22*(4), 796–810. <http://doi.org/10.1037/0882-7974.22.4.796>
- Mata, R., von Helversen, B., Karlsson, L., & Cüpper, L. (2012). Adult age differences in categorization and multiple-cue judgment. *Developmental Psychology, 48*(4), 1188–1201. <http://doi.org/10.1037/a0026084>
- Mata, R., von Helversen, B., & Rieskamp, J. (2010). Learning to choose: Cognitive aging and strategy selection learning in decision making. *Psychology and Aging, 25*(2), 299–309. <http://doi.org/10.1037/a0018923>
- Mata, R., von Helversen, B., & Rieskamp, J. (2011). When easy comes hard: The development of adaptive strategy selection. *Child Development, 82*(2), 687–700.  
<http://doi.org/10.1111/j.1467-8624.2010.01535.x>
- Mata, R., Wilke, A., & Czienskowski, U. (2009). Cognitive aging and adaptive foraging Behavior. *Journals of Gerontology - Series B Psychological Sciences and Social Sciences, 64*(4), 474–481. <http://doi.org/10.1093/geronb/gbp035>
- 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*(May 2015), 11609.  
<http://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.
- Nassar, M. R., & Gold, J. I. (2013). A healthy fear of the unknown: Perspectives on the interpretation of parameter fits from computational models in neuroscience. *PLoS Computational Biology, 9*(4), 1–6.

- <http://doi.org/10.1371/journal.pcbi.1003015>
- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*.
- Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of Neuroscience*, *30*(37), 12366–78. <http://doi.org/10.1523/JNEUROSCI.0822-10.2010>
- Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E.-J. (2011). Erroneous analyses of interactions in neuroscience: a problem of significance. *Nature Neuroscience*, *14*(9), 1105–1107. <http://doi.org/10.1038/nn.2886>
- Niv, Y., & Schoenbaum, G. (2008). Dialogues on prediction errors. *Trends in Cognitive Sciences*, *12*(7), 265–72. <http://doi.org/10.1016/j.tics.2008.03.006>
- Palminteri, S., Wyart, V., & Koehlin, E. (2016). Computational cognitive neuroscience: what is a good model of brain processes? *bioRxiv*. <http://doi.org/10.1101/079798>
- Paulsen, D. J., Platt, M. L., Huettel, S. A., & Brannon, E. M. (2011). Decision-making under risk in children, adolescents, and young adults. *Frontiers in Psychology*, *2*(APR), 72. <http://doi.org/10.3389/fpsyg.2011.00072>
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*(3), 534–552. <http://doi.org/10.1037/0278-7393.14.3.534>
- Peters, J., & Büchel, C. (2010). Neural representations of subjective reward value. *Behavioural Brain Research*. <http://doi.org/10.1016/j.bbr.2010.04.031>
- Pfeifer, J. H., & Allen, N. B. (2012). Arrested development? Reconsidering dual-systems models of brain function in adolescence and disorders. *Trends in Cognitive Sciences*, *16*(6), 322–329. <http://doi.org/10.1016/j.tics.2012.04.011>
- Pfeifer, J. H., & Allen, N. B. (2016). The audacity of specificity: Moving adolescent developmental neuroscience towards more powerful scientific paradigms and translatable models. *Developmental Cognitive Neuroscience*. <http://doi.org/10.1016/j.dcn.2015.12.012>
- Plunkett, K., Karmiloff-Smith, A., Bates, E., & Johnson, M. H. (1997). Connectionism and developmental psychology. *Journal of Child Psychology and Psychiatry*, *38*, 53–80.

- Ratcliff, R., Thapar, A., & McKoon, G. (2006). Aging and individual differences in rapid two-choice decisions. *Psychonomic Bulletin and Review*, *13*, 626–635.
- Redish, A. D., & Gordon, J. A. (Eds.). (2016). *Computational psychiatry: New perspectives on mental illness*. MIT Press.
- Samanez-Larkin, G. R., & Knutson, B. (2015). Decision making in the ageing brain: changes in affective and motivational circuits. *Nature Reviews. Neuroscience*, *16*(5), 278–289. <http://doi.org/10.1038/nrn3917>
- 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), 672–82. <http://doi.org/10.3758/s13415-014-0297-4>
- Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: a Bayesian hierarchical approach. *Psychological Review*, *120*(1), 39–64. <http://doi.org/10.1037/a0030777>
- Scholz, A., von Helversen, B., & Rieskamp, J. (2015). Eye movements reveal memory processes during similarity- and rule-based decision making. *Cognition*, *136*, 228–246. <http://doi.org/10.1016/j.cognition.2014.11.019>
- Schonberg, T., Fox, C. R., Mumford, J. A., Congdon, E., Trepel, C., & Poldrack, R. A. (2012). Decreasing ventromedial prefrontal cortex activity during sequential risk-taking: an fMRI investigation of the balloon analog risk task. *Frontiers in Neuroscience*, *6*(June), 80. <http://doi.org/10.3389/fnins.2012.00080>
- Schuch, S., & Konrad, K. (2017). Investigating task inhibition in children versus adults: A diffusion model analysis. *Journal of Experimental Child Psychology*, *156*, 143–167.
- Schultz, W., & Dickinson, A. (2000). Neuronal Coding of Prediction Errors. *Annu. Rev. Neurosci*, *23*, 473–500. <http://doi.org/10.1146/annurev.neuro.23.1.473>
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*(2), 207–222. <http://doi.org/10.1037/0033-2909.134.2.207>
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, *79*, 217–240.
- Siegler, R. (1999). Strategic development. *Trends in Cognitive Sciences*, *3*(11), 430–435. [http://doi.org/10.1016/S1364-6613\(99\)01372-8](http://doi.org/10.1016/S1364-6613(99)01372-8)
- Siegler, R., & Lemaire, P. (1997). Older and younger adults' strategy choices in

- multiplication: testing predictions of ASCM using the choice/no-choice method. *Journal of Experimental Psychology: General*, *126*(1), 71–92.
- Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the Expansion of Strategic Exploratory Behavior During Adolescence. *Journal of Experimental Psychology: General*, *146*, 155–164.
- Spaniol, J., Madden, D. J., & Voss, A. (2006). A diffusion model analysis of adult age differences in episodic and semantic long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *32*, 101–117.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT press.
- Thompson, C. A., Ratcliff, R., & McKoon, G. (2016). Individual differences in the components of children’s and adults’ information processing for simple symbolic and non-symbolic numeric decisions. *Journal of Experimental Child Psychology*, *150*, 48–71.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. <http://doi.org/10.1007/BF00122574>
- Tymula, A., Rosenberg Belmaker, L. A., Roy, A. K., Ruderman, L., Manson, K., Glimcher, P. W., & Levy, I. (2012). Adolescents’ risk-taking behavior is driven by tolerance to ambiguity. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(42), 17135–17140. <http://doi.org/10.1073/pnas.1207144109>
- van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, *121*(1), 124–49. <http://doi.org/10.1037/a0035234>
- 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. <http://doi.org/10.1093/cercor/bhr198>
- van den Bos, W., & Eppinger, B. (2016). Developing developmental cognitive neuroscience: From agenda setting to hypothesis testing. *Developmental Cognitive Neuroscience*, *17*, 138–144. <http://doi.org/10.1016/j.dcn.2015.12.011>
- van den Bos, W., & Hertwig, R. (2017). Adolescents display distinctive tolerance to

- ambiguity and to uncertainty during risky decision making. *Nature Scientific Reports*, 7, 40962. <http://doi.org/10.1038/srep40962>
- van Duijvenvoorde, A. C. K., Figner, B., Weeda, W. D., van der Molen, M. W., Jansen, B. R. J., & Huizenga, H. M. (2016). Neural mechanisms underlying compensatory and noncompensatory strategies in risky choice. *Journal of Cognitive Neuroscience*, 28(9), 1358–1373. [http://doi.org/10.1162/jocn\\_a\\_00975](http://doi.org/10.1162/jocn_a_00975)
- van Duijvenvoorde, A. C. K., Huizenga, H. M., Somerville, L. H., Delgado, M. R., Powers, A., Weeda, W. D., ... Figner, B. (2015). Neural correlates of expected risks and returns in risky choice across development. *Journal of Neuroscience*, 35(4), 1549–60. <http://doi.org/10.1523/JNEUROSCI.1924-14.2015>
- Vehtari, A., Gelman, A., & Gabry, J. (2016). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *arXiv*, (September), 1–20. <http://doi.org/10.1007/s11222-016-9696-4>
- von Helversen, B., Karlsson, L., Rasch, B., & Rieskamp, J. (2014). Neural substrates of similarity and rule-based strategies in judgment. *Frontiers in Human Neuroscience*, 8, 809. <http://doi.org/10.3389/fnhum.2014.00809>
- von Helversen, B., Mata, R., & Olsson, H. (2010). Do children profit from looking beyond looks? From similarity-based to cue abstraction processes in multiple-cue judgment. *Developmental Psychology*, 46(1), 220–229. <http://doi.org/10.1037/a0016690>
- Westbrook, A., Kester, D., & Braver, T. S. (2013). What Is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS 1*, 8, e68210.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6), 2074–2081. <http://doi.org/10.1037/a0038199>
- Wulff, D. U., & van den Bos, W. (2017). Modeling choices in delay discounting. *Psychological Science*, 1–5. <http://doi.org/10.1177/0956797616664342>
- Fechner, H. B., Pachur, T., Schooler, L. J., Mehlhorn, K., Battal, C., Volz, K. G., & Borst, J. P. (2016). Strategies for memory-based decision making: Modeling behavioral and neural signatures within a cognitive architecture. *Cognition*, 157, 77–99. DOI: 10.1016/j.cognition.2016.08.011
- von Helversen B, Karlsson L, Rasch B, Rieskamp J. Neural substrates of similarity

and rule-based strategies in judgment. *Frontiers in Human Neuroscience*. 8: 809.  
PMID 25360099 DOI: 10.3389/fnhum.2014.00809