

RESEARCH ARTICLE

A unifying Bayesian account of contextual effects in value-based choice

Francesco Rigoli^{1,2*}, Christoph Mathys^{3,4,5}, Karl J. Friston¹, Raymond J. Dolan^{1,4}

1 The Wellcome Trust Centre for Neuroimaging, UCL, 12 Queen Square, London, United Kingdom, **2** City, University of London, Northampton Square, London, United Kingdom, **3** Scuola Internazionale Superiore di Studi Avanzati (SISSA), Trieste, Italy, **4** Max Planck UCL Centre for Computational Psychiatry and Ageing Research, London, United Kingdom, **5** Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland

* f.rigoli@ucl.ac.uk



Abstract

Empirical evidence suggests the incentive value of an option is affected by other options available during choice and by options presented in the past. These contextual effects are hard to reconcile with classical theories and have inspired accounts where contextual influences play a crucial role. However, each account only addresses one or the other of the empirical findings and a unifying perspective has been elusive. Here, we offer a unifying theory of context effects on incentive value attribution and choice based on normative Bayesian principles. This formulation assumes that incentive value corresponds to a precision-weighted prediction error, where predictions are based upon expectations about reward. We show that this scheme explains a wide range of contextual effects, such as those elicited by other options available during choice (or within-choice context effects). These include both conditions in which choice requires an integration of multiple attributes and conditions where a multi-attribute integration is not necessary. Moreover, the same scheme explains context effects elicited by options presented in the past or between-choice context effects. Our formulation encompasses a wide range of contextual influences (comprising both within- and between-choice effects) by calling on Bayesian principles, without invoking *ad-hoc* assumptions. This helps clarify the contextual nature of incentive value and choice behaviour and may offer insights into psychopathologies characterized by dysfunctional decision-making, such as addiction and pathological gambling.

OPEN ACCESS

Citation: Rigoli F, Mathys C, Friston KJ, Dolan RJ (2017) A unifying Bayesian account of contextual effects in value-based choice. *PLoS Comput Biol* 13(10): e1005769. <https://doi.org/10.1371/journal.pcbi.1005769>

Editor: Laurence T. Maloney, New York University, UNITED STATES

Received: June 2, 2017

Accepted: September 11, 2017

Published: October 5, 2017

Copyright: © 2017 Rigoli et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its Supporting Information files.

Funding: This work was supported by the Wellcome Trust (wellcome.ac.uk): Karl J Friston is funded by a Wellcome Trust Principal Research Fellowship (088130/Z/09/Z). Raymond J Dolan is funded by a Senior Investigator Award (098362/Z/12/Z) and the Max Planck Society. The Wellcome Trust Centre for Neuroimaging is supported by core funding from the Wellcome Trust 091593/Z/10/Z. The funders had no role in study design, data

Author summary

Research has shown that decision-making is dramatically influenced by context. Two types of influence have been identified, one dependent on options presented in the past (between-choice effects) and the other dependent on options currently available (within-choice effects). Whether these two types of effects arise from similar mechanisms remain unclear. Here we offer a theory based on Bayesian inference which provides a unifying explanation of both between and within-choice context effect. The core idea of the theory is that the value of an option corresponds to a precision-weighted prediction error, where

collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

predictions are based upon expectations about reward. An important feature of the theory is that it is based on minimal assumptions derived from Bayesian principles. This helps clarify the contextual nature of incentive value and choice behaviour and may offer insights into psychopathologies characterized by dysfunctional decision-making, such as addiction and pathological gambling.

Introduction

Standard theories of decision-making assume that the incentive value of an option should be independent of options presented in the past and options available during choice [1–4]. These theories are fundamentally challenged by empirical evidence showing that expectations (derived from past experience) about upcoming options change value attribution and choice behaviour [5–14]. For example, in a series of recent experiments from our lab [8–10], participants made choices in blocks (i.e. contexts) associated with one of two distinct, but partially overlapping, reward distributions. Participants' choices were consistent with attributing a larger incentive value to rewards (common to both contexts) in blocks associated with low compared to high average reward. In other words, the incentive value of a reward increased when the average was lower. In addition to the average reward of a context, evidence from a similar task indicated that reward variance within a given context also exerts an influence on incentive value [11]. These findings highlight contextual effects exerted by expectations about options (induced, for example, by options available during previous choices); namely, *between-choice* contextual effects.

In addition, the empirical literature has highlighted contextual influences elicited by options available during choice; namely, *within-choice* context effects [6, 15–20]. Standard theories of decision-making assume that the incentive value of an option should be independent of other options available during choice [1–4]. This implies that the choice proportion between two options, comprising a more valuable and a less valuable option, should be unaffected by the introduction of a third [2]. However, a recent study [6] has shown that this choice proportion follows a U-shape function, which diminishes as the value of a third option approaches the value of the target options—and starts increasing thereafter (Fig 1A). This is hard to reconcile with standard theories and represents a form of within-choice context effect, whereby the value of an option is affected by other options available during choice.

In this task, it is unnecessary to compare options across different attributes (single-attribute decisions; [6]). However, other forms of within-choice context effect have been observed when options are defined by the same set of attributes that have to be traded off against each other (multiattribute decisions; [15–20]). For example, consider a binary choice between a high-quality and expensive car A versus a low-quality and cheap car B (Fig 1B). Imagine the values of the attributes are such that an agent is indifferent about the two options (i.e., the higher price of car A is exactly compensated by its quality), resulting in the same probability of choosing options A and B. What happens if a third option is also available? Standard models (based on the assumption that values are independent of other options) predict that the choice probability difference will remain zero, independent of a third option. However, empirical data highlight a so-called *similarity effect* [20–23], whereby preference for an option over a second option—which is equally preferable during binary decisions—increases if a third option is available that is similar to the second option (Fig 1B). In our example, the choice probability difference between car A and car B will be positive when a third low-quality and cheap (similar to car B) car C is also available. A form of influence called the *attraction effect* [15, 24–26] has also




Fig 1. A: Empirical evidence concerning contextual effects elicited by multiple options available during choice (within-choice context effects). **A:** Single-attribute choice, where there is no need to integrate multiple attributes to make a decision. Here a better target option associated with reward R_H and a worse target option associated with reward R_L are available together with a third option associated with reward R_3 . The graph plots empirical findings [6] in terms of the ratio between the probability of choosing the better target option ($P[R_H|R_H, R_L, R_3]$) and the probability of choosing the worse target option ($P[R_L|R_H, R_L, R_3]$) as a function of the (normalized) reward of a third options R_3 (see Fig 5C in [6]). **B:** Multiattribute choice, where multiple attributes need to be integrated to make a decision. Here, we consider the difference in choice probability between a high-quality and high-price car A and a low-quality and low-price car B. Although during binary choice this difference is zero, empirical evidence has shown this difference can be non-zero when a third option is also available. A similarity effect favours car A over car B when a low-quality and low-price car C (similar to car B) is available. An attraction effect favours car A over car B when a medium-quality and high-price car D is available. A compromise effect consists in favouring a medium-quality and medium-price car E over both car A and car B during choices in which all three cars are available, despite the fact these cars are equally chosen when they are available in pairs during binary choices.

<https://doi.org/10.1371/journal.pcbi.1005769.g001>

been found with the availability of a third option that is characterized by a low score for one attribute and an intermediate score for the other (Fig 1B). The presence of such a third option favours the option with a high score for the attribute for which the third option has an intermediate score. In our example, the choice probability difference between car A and car B will be positive when a third medium-quality and expensive car D is also available. Finally, empirical data are consistent with a so-called *compromise effect* [17, 25, 27]. This applies when the choice set includes two options scoring high in one attribute and low in another plus a third option characterized by intermediate scores for both attributes. While the three options are equally preferred (i.e., are chosen an equal amount of times) if presented in pairs during binary choices, when they are all available, a preference for the option characterized by intermediate scores is observed (Fig 1B). For instance, although during binary choices an average-price and average-quality car E is not preferred over car A or over car B, car E will be favoured when presented together with both car A and car B.

Several explanations have been proposed to account for contextual effects on incentive value and choice, with most models focusing on within-choice context effects during multiattribute decisions [16–20, 27, 28]. Other theories have been proposed to explain between-choice context effects [29–31], and disregard within-choice effects. We are aware of a single attempt to encompass both between-choice and within-choice effects, though restricted to non-multiattribute

decisions for the latter type of effects [6]. Whether models developed to explain a certain class of context effects generalise to other effects remains unclear—and a unifying account encompassing all known context effects is lacking. Developing a parsimonious account would represent an important theoretical advance, as it would explain diverse empirical phenomena with the same underlying principles.

The goal of the present paper is to describe a unifying theory, referred to as Bayesian model of context sensitive value (BCV) that explains between-context and within-context effects, in single and multiattribute decisions. This theory represents a generalization of a recent model developed to explain between-choice contextual effects [11]. The key idea is that agents build a generative model of reward within a context and, every time a new reward or option is presented, use Bayesian inference to invert this model to form a posterior belief about the underlying reward distribution. Incentive value is computed during this belief update and corresponds to a precision-weighted reward prediction error. The advantage of this theory relies on its grounding upon simple normative principles of Bayesian statistics. In addition, the model can explain between-choice context effects [8–11] and makes specific predictions that have been confirmed empirically.

In brief, BCV proposes that the incentive value of a stimulus (or option) corresponds to the change in reward expected (in any given context) when the stimulus is presented. This makes precise predictions about choice under ideal (Bayesian) observer assumptions (with a minimal number of free parameters). Crucially, predictions include specific forms of context effects, and raise a question of whether these predicted effects are consistent with empirical findings.

In this paper, we applied BCV to multi-alternative choice (considering both single and multiattribute decisions) and ask whether the model predicts the context-effects found empirically. We first present a theoretical extension of BCV applicable to decisions in which multiple options are available and can be characterized by multiple attributes. We show that predictions derived from the model are remarkably similar to empirical findings on within-choice contextual effects, both during non-multiattribute and multiattribute decisions. We next review BCV in relation to between-choice context effects and describe how the model can also explain these empirical findings. On this basis, we offer the model as a principled description of between and within-choice context effects.

Results

Within-choice context effects

The idea behind BCV is to establish a link between theories of value and normative accounts of brain functioning based on Bayesian statistics [32–37]. The Bayesian brain framework rests on the idea that an agent builds a model of the processes generating sensory cues. The generative model comprises a set of random variables (i.e., hidden states or causes of sensory outcomes) and their causal links (i.e., probabilistic contingencies). The variables can be separated into hidden and observable variables, the former representing the latent causes of observations, and the latter representing sensory evidence or cues. Sensory evidence conveyed by observable variables is combined with prior beliefs about hidden causes to produce a posterior belief about the causes of observations. The application of this logic has proved effective in explaining several empirical phenomena in perception [32–37]. For instance, psychophysical data indicate that human perception depends on integrating different perceptual modalities (e.g., visual and tactile) in a manner consistent with Bayesian principles [38], where evidence is weighted by the precision of sensory information. Furthermore, process theories that mediate Bayesian inference (e.g., predictive coding) have a large explanatory scope in terms of neuroanatomy and physiology [39].

Inspired by a recent framework that conceptualises planning and choice as active inference [40–45], our core proposal is that Bayesian inference drives the attribution of incentive value to reward, and this in turn determines choice.

In a previous work, we have developed a version of BCV applicable to conditions where past options elicit context effects by shaping expectancies before a reward is presented ([11]; see below). However, our previous formulation did not consider conditions where multiple options (potentially characterized by multiple attributes) are available. Here we generalize BCV to encompass conditions in which multiple options are available and options can be characterized by multiple attributes. We define a multi-attribute option u_n (e.g., car A or car B) as a contract that yields reward amount $R_{i,n}$ relative to each attribute i (e.g., price or quality):

$$u_n = \{R_{i,n}\}_{i=1,\dots,I_n} \quad \text{with} \quad R_{i,n} \in \mathbb{R} \quad (1)$$

An option set u is the set of options currently available:

$$\text{Set } u = \{u_n\}_{n=1,\dots,N} \quad \text{with} \quad u_n \text{ an option } \forall n \quad (2)$$

The expected value (EV) of an option u_n corresponds to:

$$R_{u_n} = \sum_i R_{i,n} \quad (3)$$

For example, the total reward for car A is equal to the reward associated with price plus the reward associated with quality. BCV assumes that an agent builds a generative model of the reward amounts $R_{i,n}$ (Fig 2). Specifically, an agent believes that, for each attribute i , reward amounts $R_{i,n}$ across options are sampled from the same population. To distinguish among attributes, we assume that an agent believes that an independent population of reward amounts is associated to each attribute. For example, if two attributes characterize options, two independent populations of reward amounts are considered by the agent (Fig 2).

Formally, for each attribute i , the average of the population of the reward amounts $R_{i,n}$ is represented by a random variable C_i , which is assumed to be sampled from a Gaussian distribution with prior mean μ_{C_i} and uncertainty (variance) $\sigma_{C_i}^2$:

$$C_i \sim \mathcal{N}(\mu_{C_i}, \sigma_{C_i}^2) \quad (4)$$

The agent assumes that μ_{C_i} and $\sigma_{C_i}^2$ are known but that C_i is not directly observable and therefore needs to be inferred from observing the different instances of reward amounts $R_{i,n}$ of options for the attribute i . This is realized in the generative model by treating C_i as a hidden cause of Gaussian variables $R_{i,n}$ with mean C_i and uncertainty $\sigma_{R_i}^2$:

$$R_{i,n} \sim \mathcal{N}(C_i, \sigma_{R_i}^2) \quad (5)$$

On the basis of the generative model, for each attribute i , the agent can estimate $\hat{C}_i = \mathcal{N}(\hat{\mu}_{C_i|R_i}, \hat{\sigma}_{C_i|R_i}^2)$, namely the posterior belief about the variable C_i (i.e., the average reward amount relative to the attribute i ; the hat symbol indicates estimates of unknown quantities), given the observation of all reward amounts of all options available for the attribute i , represented by the set R_i . In other words, an agent assumes that there is an average reward for each attribute which is unknown but can be estimated based on the reward amounts.

According to Bayes' rule, the posterior belief of C_i can be calculated by considering the associated $R_{i,n}$ sequentially in any order. We propose such sequential belief updating for BCV, even if options (and the associated reward amounts) are presented simultaneously, and we assume that the order of options considered is random (with potentially different orders for

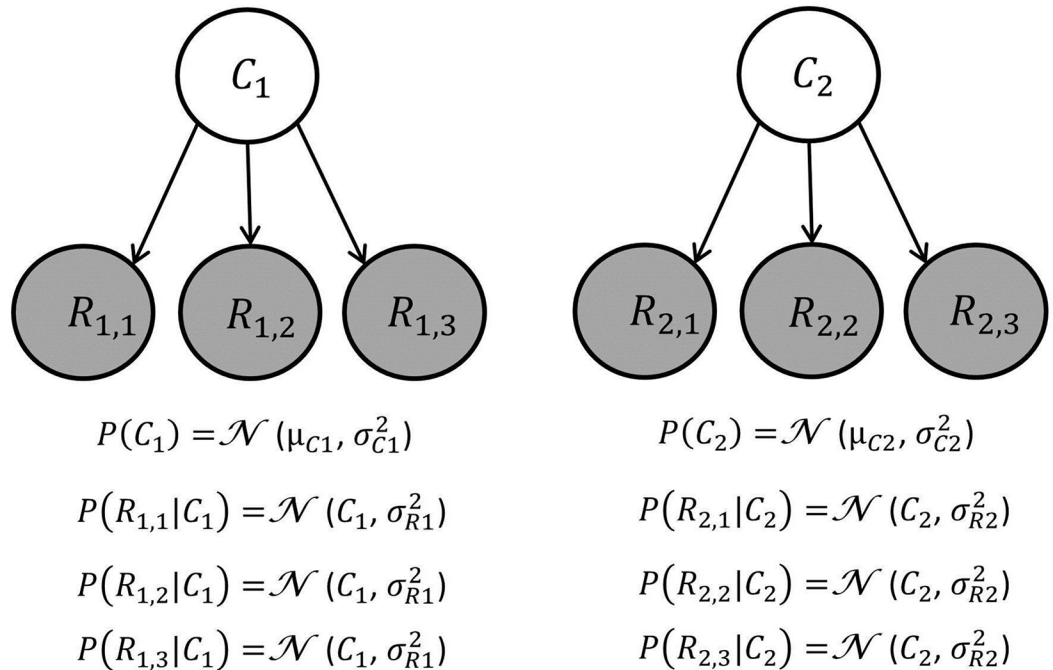


Fig 2. Example of a basic generative model underlying BCV. This is a directed acyclic graph or Bayesian network. Circles represent random variables (shaded and white circles refer to observed and non-observed variables respectively). An arrow denotes a conditional dependence—in which one random variable supplies the mean of the probability distribution of its children. For each attribute i , a hidden variable C_i represents the belief about the average reward amount across options for the attribute i . This generates the mean for Gaussian observable variables $R_{i,n}$, corresponding to reward amounts associated with options available during choice. In this example, three options are available and options are characterized by two attributes. Note that attributes are independent in the generative model, as there is no arrow connecting variables associated with different attributes. Inverting this model, given observations, furnishes posterior beliefs over the mean reward amount across options for each attribute i . This inference is performed sequentially integrating one reward amount observation at each inference step. When a reward observation is considered, its incentive value is conceived as (precision-weighted) reward prediction error.

<https://doi.org/10.1371/journal.pcbi.1005769.g002>

different attributes). For example, when three options characterized by two attributes are available (represented by $R_{1,1}$, $R_{1,2}$ and $R_{1,3}$ for attribute one, and $R_{2,1}$, $R_{2,2}$ and $R_{2,3}$ for attribute two), inference can involve computing, in order, $P(C_1|R_{1,1})$, $P(C_1|R_{1,1}, R_{1,3})$ and $P(C_1|R_{1,1}, R_{1,2}, R_{1,3})$ for attribute one, and $P(C_2|R_{2,3})$, $P(C_2|R_{2,1}, R_{2,3})$ and $P(C_2|R_{2,1}, R_{2,2}, R_{2,3})$ for attribute two. In the example above, an agent may consider first car A and next car B when estimating the average reward for price, and first car B and next car A when estimating the average reward for quality.

The rationale behind sequential belief updating is that the brain is equipped with a limited computational capacity, which precludes the instantaneous (and parallel) evidence accumulation, and hence requires the processing of one option after another. A similar evidence accumulation process is implicit in some theories of perceptual and value-based decision-making (e.g., [16, 46, 47]). Below, we will show that this evidence accumulation, in the form of sequential Bayesian belief updating, endows agents with the right sort of sensitivity to context. Formally, if $R_{i,n}$ is the reward amount considered first during belief updating, in relation to attribute i , the posterior mean $\hat{\mu}_{C_i|R_{i,n}}$ is [48]:

$$\hat{\mu}_{C_i|R_{i,n}} = \mu_{C_i} + \frac{\sigma_{C_i}^2}{\sigma_{C_i}^2 + \sigma_{R_i}^2} (R_{i,n} - \mu_{C_i}) \tag{6}$$

The posterior uncertainty $\hat{\sigma}_{Ci|R_{i,n}}^2$ is:

$$\hat{\sigma}_{Ci|R_{i,n}}^2 = \sigma_{Ci}^2 - \frac{\sigma_{Ci}^2}{\sigma_{Ci}^2 + \sigma_{Ri}^2} \sigma_{Ci}^2 \tag{7}$$

The crucial proposal we advance is that the incentive value $V_i(R_{i,n})$ —attributed to a reward amount $R_{i,n}$ in relation to the attribute i and associated with option u_n —is central to belief updating (see Eq 6) and corresponds to a *precision-weighted prediction error* [49]; namely, to the difference between $R_{i,n}$ and the prior mean μ_{Ci} , multiplied by a gain term which depends on the uncertainty of that attribute σ_{Ri}^2 and the prior uncertainty σ_{Ci}^2 :

$$\begin{aligned} V_i(R_{i,n}) &= \frac{\sigma_{Ci}^2}{\sigma_{Ci}^2 + \sigma_{Ri}^2} (R_{i,n} - \mu_{Ci}) \\ \Rightarrow \\ \hat{\mu}_{Ci|R_{i,n}} &= \mu_{Ci} + V_i(R_{i,n}) \end{aligned} \tag{8}$$

Within BCV, incentive value imbues reward and associated options with behavioural relevance, by favouring either approach to (for positive incentive values) or avoidance of (for negative incentive values) these reward amounts and options.

This implies two fundamental forms of contextual normalization. First, a subtractive normalization is exerted when μ_{Ci} is different from zero. For example, if we assign positive and negative numbers to rewards (i.e., $R_{i,n} > 0$) and punishments (i.e., $R_{i,n} < 0$) respectively, their corresponding incentive values will change in sign, depending on whether punishment (i.e., $\mu_{Ci} < 0$) or reward (i.e., $\mu_{Ci} > 0$) is expected a priori. Small rewards may appear as losses in contexts where large rewards are expected. Second, a divisive normalization depends on considering the gain term $\frac{\sigma_{Ci}^2}{\sigma_{Ci}^2 + \sigma_{Ri}^2}$. This implies that the positive and negative value of profits (i.e., $R_{i,n} > \mu_{Ci}$) and losses (i.e., $R_{i,n} < \mu_{Ci}$) are magnified by a large gain term, when we have precise beliefs about the average reward of the population.

Sequential Bayesian belief updating means that inference proceeds by considering one reward amount at a time. If $R_{i,n}$ is considered at step $t+1$ and $\mathbf{R}_{i,t}$ is a set containing all reward amounts already seen up until step t for attribute i , then a posterior mean $\mu_{Ci|R_{i,t},R_{i,n}}$ is obtained at step $t+1$ equivalent to (Bishop, 2006):

$$\hat{\mu}_{Ci|R_{i,t},R_{i,n}} = \hat{\mu}_{Ci|R_{i,t}} + \frac{\hat{\sigma}_{Ci|R_{i,t}}^2}{\hat{\sigma}_{Ci|R_{i,t}}^2 + \sigma_{Ri}^2} (R_{i,n} - \hat{\mu}_{Ci|R_{i,t}}) \tag{9}$$

Implying a value for the reward amount $R_{i,n}$:

$$V_i(R_{i,n}) = \frac{\hat{\sigma}_{Ci|R_{i,t}}^2}{\hat{\sigma}_{Ci|R_{i,t}}^2 + \sigma_{Ri}^2} (R_{i,n} - \hat{\mu}_{Ci|R_{i,t}}) \tag{10}$$

For each attribute i , incentive values are accumulated in memory until inference is completed (i.e., all reward amounts have been considered). We can assume that inference proceeds in sequence or in parallel across attributes; however, this has no impact on incentive values, as the agent believes that attributes are associated with independent reward populations (formally: $P(C_1, C_2, \dots, C_l | \mathbf{R}) = P(C_1 | \mathbf{R}_1), P(C_2 | \mathbf{R}_2), \dots, P(C_l | \mathbf{R}_l)$).

When all attributes for an option u_n have been considered, we assume that the incentive value of the option corresponds to the sum of the incentive values of associated reward

amounts:

$$V(u_n) = \sum_{i=1} V_i(R_{i,n}) \quad (11)$$

Inference proceeds until, for all attributes i , the posterior expectation about rewards $\hat{\mu}_{C_i R_i}$ is evaluated and, at this point, a choice is realized following a softmax rule based on the incentive values of the available options [2].

In summary, BCV is based on the following assumptions:

1. Each attribute is associated with an average reward (which is a hidden variable).
2. Average rewards for different attributes are independent.
3. For each attribute, the rewards offered (or observed rewards) are treated as samples that depend on the average reward.
4. The observed rewards are used to invert the model and infer the average reward.
5. During inference, observed rewards are considered sequentially.
6. During inference, an incentive value is calculated for each observed reward that corresponds to a precision-weighted prediction error.
7. Incentive values are summed across observed rewards and attributes, and choice follows a softmax rule.

Below these assumptions are discussed in detail. Assumptions (i), (iii), and (iv) are implicit in adopting a Bayesian scheme. Assumption (vii) is based on a standard approach in which incentive values are summed and a softmax choice rule is adopted. Assumption (ii) captures the notion of multiple attributes, in other words it enables an agent to link rewards to their attributes. Assumption (vi) (sequential belief updating or evidence accumulation) reflects the real world constraint that people have to evaluate available options and rewards one by one. In other words, agents cannot magically and instantaneously assimilate all the options on offer—they have to accumulate evidence for the underlying payoffs by evaluating each in turn. This notion plays a central role since we will see that context effects emerge because a reward is contextualized by previous rewards encountered during inference. This underlies assumption (v) that associates incentive value with a precision-weighted prediction error—a central construct in Bayesian inference. Heuristically, this scheme implies that an option is more likely to be selected if it increases expectations of reward, and will be avoided if it decreases expectations. In other words, an option is more likely to be selected if it suggests the situation is better than indicated by options considered previously during belief updating.

Note that a Bayesian perspective may suggest that incentive value corresponds to a posterior belief—rather than a precision-weighted prediction error. As an example, this would imply that the value of the same dish will be perceived as ‘higher’ in a ‘better’ restaurant. However, empirical data are consistent with the opposite notion that (adopting the same example) the value of the same dish is perceived as lower in a better restaurant [6–14]. This evidence motivated our proposal that incentive value corresponds to a precision-weighted prediction error, and not to a posterior belief.

In sum, BCV provides a principled explanation for how Bayesian inference, assigning a key role to prior expectation and uncertainty, might underlie value computation and choice. The key role of uncertainty is reflected in the precision-weighting of prediction errors. The hypothesis we entertain here is that the mechanisms postulated by BCV may be general and explain multiple forms of context effects. We have previously applied BCV to explain between-choice

context effects; namely, those elicited by options presented in the past [11]. Here, we explore the possibility of applying the same model to within-choice effects, which arise when multiple options are available. In what follows, we will consider single and multiattribute choices under this Bayesian formalism.

Single attribute decisions

Here, we apply BCV to explain within-choice contextual influences during non-multiattribute decisions. These comprise choices in which trading-off different attribute is not required, as for instance when options are defined by a single attribute. Consider first how different prior expectations μ_C (i.e., the prior expectation over the average reward of the attribute) and reward uncertainty σ_R^2 affect the choice between two options characterized by a single attribute (Fig 3). We can examine the predicted proportion of choosing a better option (associated with high reward R_H) compared to a worse option (associated with low reward R_L), as a function of prior expectation μ_C and reward uncertainty σ_R^2 . Classical theories predict a flat function because they do not model an influence of the prior mean μ_C and reward uncertainty σ_R^2 [1–4]. In contrast, BCV predicts bell-shape functions over prior expectations, that peak at the prior mean of $\mu_C = (R_H - R_L)/2$ (Fig 3). In this setting, the reward uncertainty σ_R^2 determines the width of the function (larger uncertainties produce narrower functions). Below, we analyse conditions where more than two options are available—and within-choice context effects come into play.

Classical decision-making models predict that, during choice, the choice ratio between two options should not be affected by the reward associated with a third option [1–4]. However, a recent study has challenged this hypothesis, highlighting within-choice context effects [6]. Adopting a choice task in which three options were available during choice, this study showed that the choice proportion between a more valuable and a less valuable target option diminished as a third option value increased towards the value of the target options (Fig 1A). After this point, the choice proportion started increasing (Fig 1A).




Fig 3. A: Generative model involved during choice between two options (characterized by a single attribute), one associated with high reward ($R_H = 10$) and the other with low reward ($R_L = 6$). **B:** Proportion of choices of the better over choices of the worse option predicted by BCV ($P[R_H|R_H, R_L] / P[R_L|R_H, R_L]$), as a function of prior expectation μ_C and reward uncertainty σ_R^2 (100000 trials are simulated for each condition; $\sigma_C^2 = 1$ for simulations). BCV assumes a softmax choice rule (with inverse temperature parameter equal to one for all simulations) and an equal probability for each option of being considered during the first inference step. This shows bell-shape functions where peaks correspond to a prior mean expectation of $(R_H - R_L)/2$ and where the reward uncertainty σ_R^2 determines the width of the function (smaller uncertainties are connected with narrower functions).

<https://doi.org/10.1371/journal.pcbi.1005769.g003>

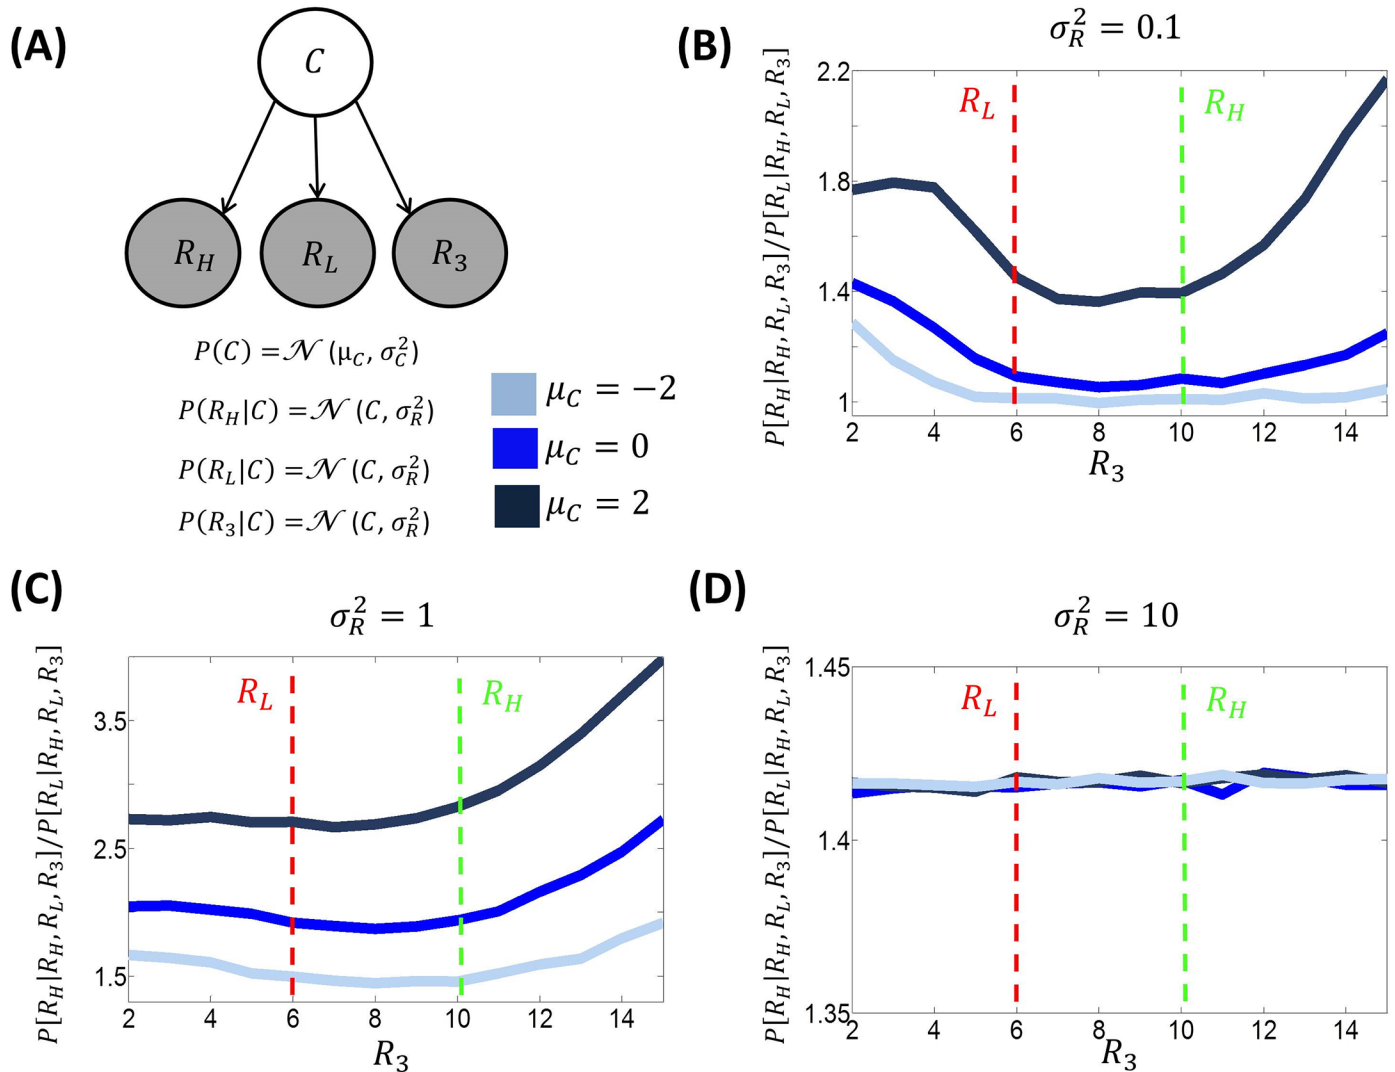


Fig 4. A: Generative model involved during choice between two options, one associated with high reward ($R_H = 10$) and the other with low reward ($R_L = 6$) when a third option (associated with reward R_3) is also available [6]. **B:** Proportion of choices of the better over choices of the worse option predicted by BCV ($P[R_H|R_H, R_L, R_3]/P[R_L|R_H, R_L, R_3]$), as a function of the third option reward R_3 and prior expectation μ_C (100000 trials are simulated for each condition; $\sigma_C^2 = 1$ for simulations). Here, the reward uncertainty σ_R^2 was set to 0.1; **C:** The same simulation is reported except that the reward uncertainty σ_R^2 was set to one. **D:** The same simulation is reported except that the reward uncertainty σ_R^2 was set to ten.

<https://doi.org/10.1371/journal.pcbi.1005769.g004>

Here, we examine the implications of applying BCV in this scenario. Fig 4 illustrates the predictions of BCV of the ratio of choices of the two target options (a better target option R_H and a worse target option R_L) as a function of the reward of a third option R_3 and as a function of the agent's prior belief about the average option reward μ_C and about the reward uncertainty σ_R^2 . This figure shows that all these variables exert an influence. First, for certain values of reward uncertainty σ_R^2 and prior mean μ_C , the reward of a third option R_3 influences the choice proportion between the two target options according to a U-shape function, in a way that is consistent with empirical findings (Fig 1A). Second, the impact exerted by the reward of a third option R_3 decreases as the reward uncertainty σ_R^2 increases. In other words, within-choice context effects emerge only with small reward uncertainty σ_R^2 . This can be explained by the fact that a small σ_R^2 magnifies reward prediction error (RPE), enabling contextual effects to emerge.

Third, when the reward uncertainty σ_R^2 is sufficiently small, the prior mean μ_C comes into play. Overall, a larger prior mean μ_C increases the choice proportion between the two target options (independently of the reward of a third option R_3). Furthermore, the prior mean μ_C exerts a modulatory influence on the effect of the reward of a third option R_3 , as the effect exerted by R_3 is enhanced with a larger prior mean μ_C . Note that context effects exerted by R_3 are obtained with $\mu_C = 0$, which can be considered a default value for this parameter.

Collectively, these simulations provide proof of principle that BCV can explain within-choice contextual effects in single-attribute decisions that are remarkably similar to those seen in empirical studies [6]. In what follows, we now extend the explanatory scope of BCV to multiattribute problems.

Multiattribute decisions

Empirical studies of multi-attribute decisions have highlighted three forms of effects, including the *similarity* [20–23], *attraction* [15, 24–27], and *compromise* effect [17, 25, 27]. Here, we apply BCV to multi-attribute decisions and ask whether the predictions that emerge from the model reproduce the context effects found empirically. To this aim, we consider two options (e.g., the two cars A and B described above) defined by two attributes (e.g., price p and quality q). Considering the reward amounts of car A, we assign $R_{p,A} = 1$ to price (low scores indicate high price) and $R_{q,A} = 10$ to quality. Conversely, when considering the reward amounts of car B, we assign $R_{p,B} = 10$ to price and $R_{q,B} = 1$ to quality. We now consider the choice probability difference between option A and option B as a function of the reward amounts $R_{p,K}$ and $R_{q,K}$ of a third option K.

Empirical evidence is hard to reconcile with standard models of choice, which predict that the choice probability difference between option A and option B should not depend on the value of a third option K. Fig 5A summarises the empirical findings by plotting the probability of choosing A minus the probability of choosing B as a function of the attributes of a third option K. This graph shows conditions in which the choice probability difference is bigger or smaller than zero, illustrating both a similarity and an attraction effect. Specifically, a similarity effect favours option A when option K is good in price and bad in quality (top-left of the graph), and favours option B when option K is bad in price and good in quality (bottom-right of the graph). An attraction effect favours option A when option K is bad in price and has an average quality (bottom-middle of the graph), and favours option B when option K has an average price and is bad in quality (middle-left of the graph).

We can now apply BCV to model choices in this scenario by analysing the influence on the choice probability (difference between option A and option B) of the prior mean μ_C (we use an equal prior mean for both attributes price and quality; formally: $\mu_{Cp} = \mu_{Cq}$), the reward uncertainty σ_R^2 (we use an equal reward uncertainty for both attributes price and quality; formally: $\sigma_{Rp}^2 = \sigma_{Rq}^2$), and the reward amounts $R_{p,K}$ and $R_{q,K}$, associated with price and quality respectively, of option K.

Fig 5B illustrates the choice probability difference (between option A and option B) with prior mean $\mu_C = 0$ and reward uncertainty $\sigma_R^2 = 0.1$. Focusing on areas of the graph where a similarity effect can be tested (i.e., top-left and bottom-right), we see that the similarity effect is reproduced by BCV. Moreover, focusing on areas of the graphs where an attraction effect can be tested (i.e., bottom-middle and middle-left), we can see that this effect can also be explained by BCV. Collectively, these simulations provide proof of principle that, for some sets of values of the prior mean μ_C and of the reward uncertainty σ_R^2 , BCV explains both a similarity and an attraction effect. Note that these effects are obtained with $\mu_C = 0$, which can be considered a default value for this parameter.

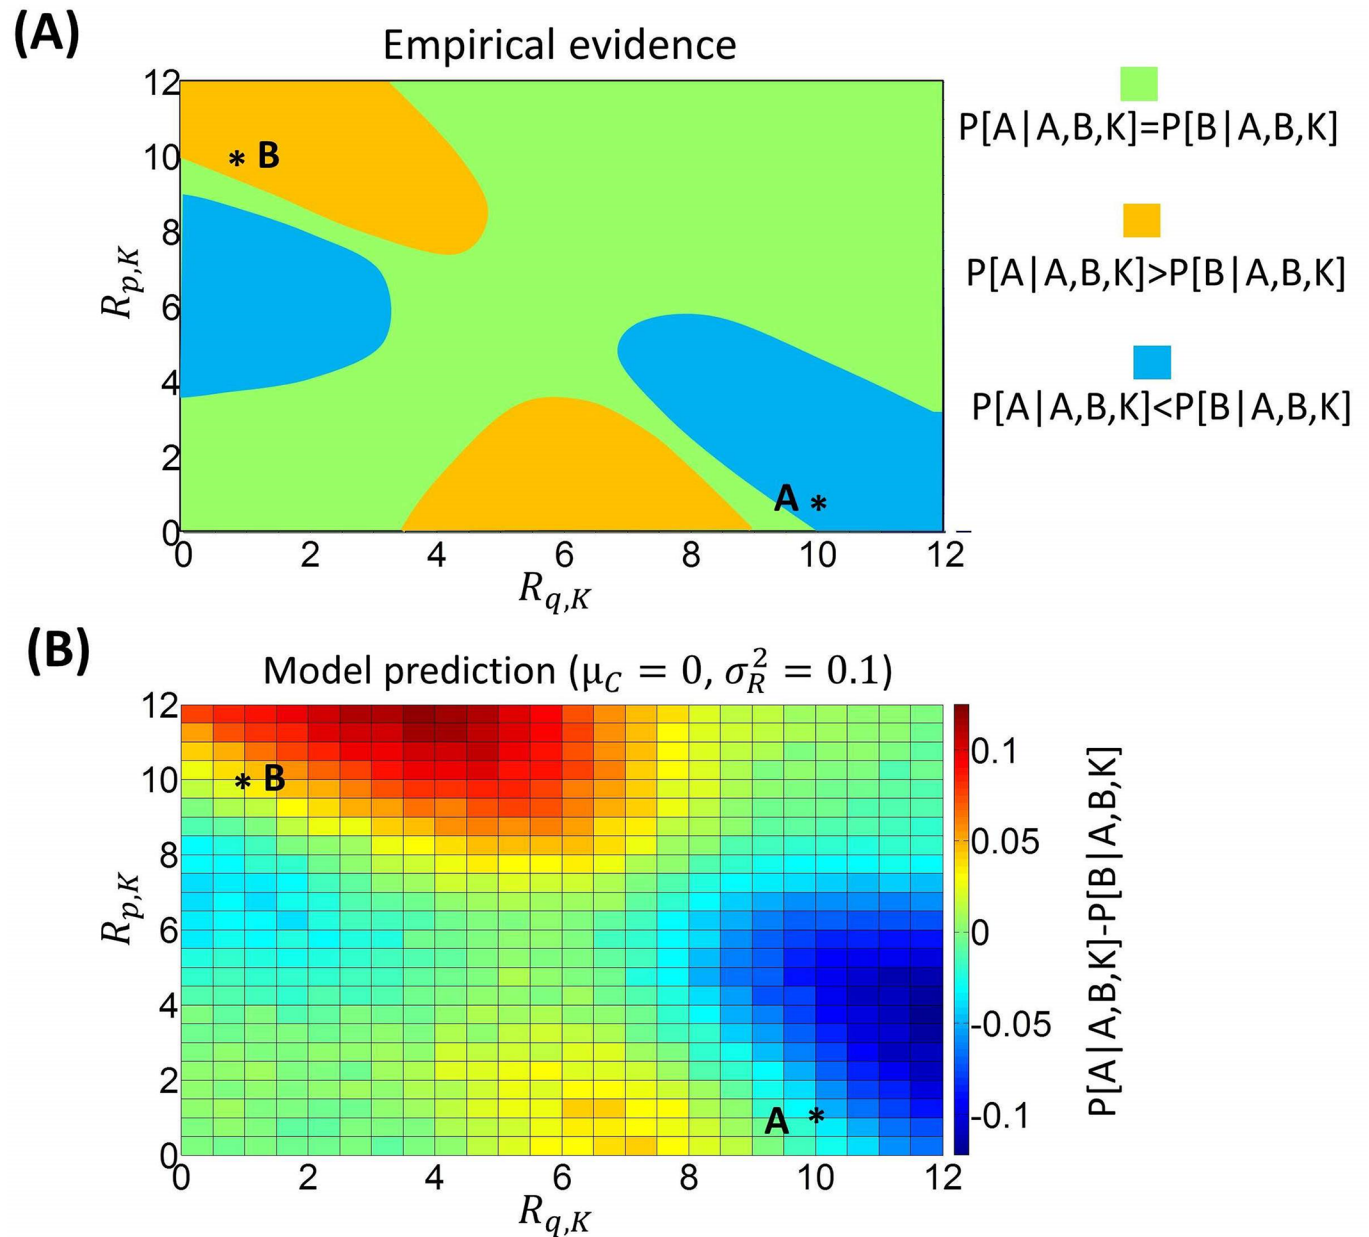


Fig 5. A: Empirical evidence (derived from integrating data from available studies as in [19]) concerning the difference in probability between choosing option A and option B when a third option K is available ($P[A | A, B, K] - P[B | A, B, K]$). Here options are characterized by two attributes (price p and quality q). For car A, we assign $R_{p,A} = 1$ to price (low scores indicate high price) and $R_{q,A} = 10$ to quality. For car B, we assign $R_{p,B} = 10$ to price and $R_{q,B} = 1$ to quality. The graph considers the choice probability difference between option A and option B as a function of the reward amounts $R_{q,K}$ (for quality; x axis) and $R_{p,K}$ (for price; y axis) of a third option K. Green areas indicate values for which no difference is expected based on empirical evidence; orange and blue areas indicates values for which a positive and negative difference is expected, respectively. **B:** The same analysis is performed with data simulated using BCV (100000 trials are simulated for each condition; $\mu_C = 0$; $\sigma_R^2 = 0.1$; $\sigma_C^2 = 1$ for simulations).

<https://doi.org/10.1371/journal.pcbi.1005769.g005>

Fig 6A and 6B examines the effects of adopting other values of the prior mean μ_C (fixing the reward uncertainty σ_R^2 to 0.1) in this scenario. This figure shows that an attraction effect is obtained when the prior mean μ_C is smaller ($\mu_C = -2$ in our simulation), but no similarity effect emerges. Conversely, a similarity effect is evident when the prior mean μ_C is larger ($\mu_C = 2$ in our simulation), but the attraction effect vanishes.

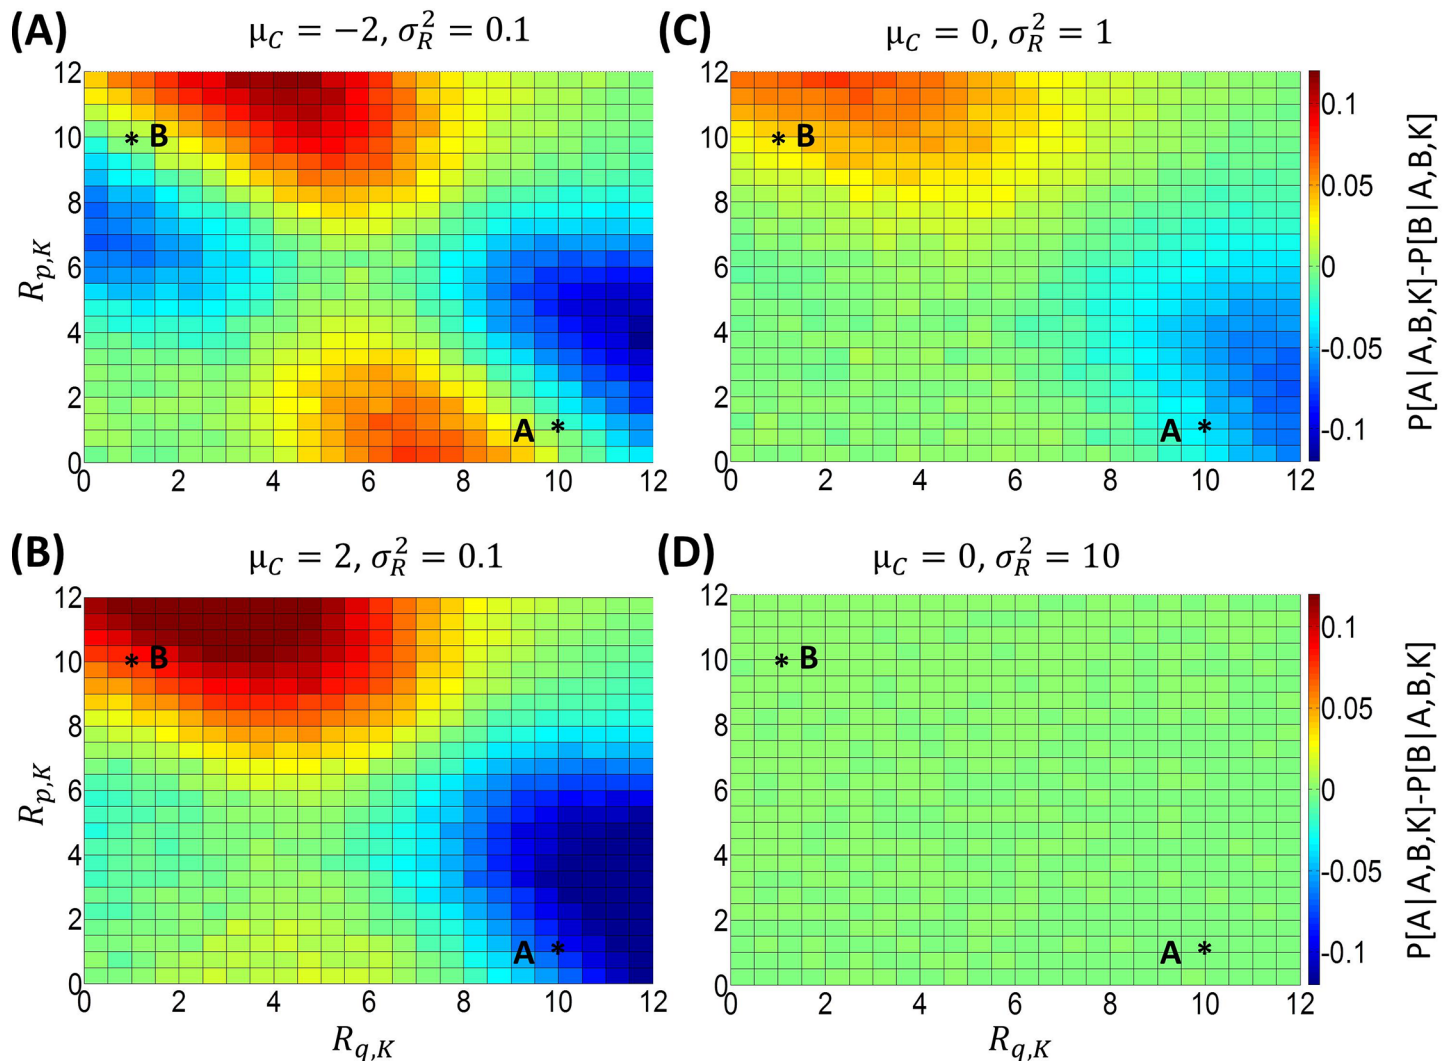


Fig 6. Predictions of BCV about the difference in probability between choosing option A and option B when a third option K is available ($P[A|A,B,K] - P[B|A,B,K]$). Here options are characterized by two attributes (price p and quality q). For car A, we assign $R_{p,A} = 1$ to price (low scores indicate high price) and $R_{q,A} = 10$ to quality. For car B, we assign $R_{p,B} = 10$ to price and $R_{q,B} = 1$ to quality. The graph considers the choice probability difference between option A and option B as a function of the reward amounts $R_{q,K}$ (for quality; x axis) and $R_{p,K}$ (for price; y axis) of a third option K (100000 trials are simulated for each condition; $\sigma_C^2 = 1$ for simulations). Different parameter sets are shown. **A:** Simulation using $\mu_C = -2$ and $\sigma_R^2 = 0.1$. **B:** Simulation using $\mu_C = 2$ and $\sigma_R^2 = 0.1$. **C:** Simulation using $\mu_C = 0$ and $\sigma_R^2 = 1$. **D:** Simulation using $\mu_C = 0$ and $\sigma_R^2 = 10$.

<https://doi.org/10.1371/journal.pcbi.1005769.g006>

Fig 6C and 6D illustrates the choice probability difference (between option A and option B) for different values of the reward uncertainty σ_R^2 (the prior mean μ_C was fixed to zero). We can see that both similarity and attraction effects are not detectable when reward uncertainty σ_R^2 is high. For smaller values of uncertainty, a similarity effect emerges but there is no attraction effect. Both effects can be obtained only when the reward uncertainty σ_R^2 is sufficiently low (Fig 5B). This highlights the role of the reward uncertainty σ_R^2 in determining the degree of contextual effects.

In summary, our analyses show that, when simulating multi-attribute decisions with BCV, similarity and attraction effects emerge for appropriate values of the prior mean μ_C and the reward uncertainty σ_R^2 . The first parameter regulates the balance of the two effects, as an attraction effect (but no similarity effect) is obtained when the prior mean μ_C is small, while a

similarity effect (but no attraction effect) is obtained when the prior mean μ_C is large. Both effects emerge for intermediate values of the prior mean μ_C , including a prior mean $\mu_C = 0$, which is a default value for this parameter. The reward uncertainty σ_R^2 plays a key role too, because context effects vanish when this parameter is high. Decreasing levels of reward uncertainty σ_R^2 reveal a similarity effect first and then an attraction effect. These results indicate that the similarity and attraction effects arise naturally from BCV, without any *ad-hoc* assumptions—and under natural values of model parameters (prior mean μ_C reward uncertainty σ_R^2).

A compromise effect [17, 25, 27] has been observed when the choice set includes two options scoring high in one attribute and low in another, in addition to a third option with intermediate scores for both attributes. Crucially, the three options are equally preferred (i.e., are chosen an equal amount of times) if presented in pairs during binary choices. However, when they are available altogether, a preference for the option characterized by intermediate scores is seen. We model this scenario by manipulating the distance between attributes for two options A and B, namely assigning $R_{p,A} = 5 - d$ and $R_{q,A} = 5 + d$ for option A, and $R_{p,B} = 5 + d$ and $R_{q,B} = 5 - d$ for option B, where the *proximity parameter* d varies (across simulations) from zero to four. To represent the option with intermediate scores for both the two attributes, we assign $R_{p,K} = 5$ and $R_{q,K} = 5$.

Fig 7A, 7B and 7C shows the prediction of BCV using these settings during binary choices between option K and option A, using different parameters for the prior mean μ_C and the reward uncertainty σ_R^2 . The results indicate that the choice probability difference is always zero, irrespective of the values of the proximity parameter d or the parameters of the model (prior mean μ_C and reward uncertainty σ_R^2). Fig 7D, 7E and 7F shows the choice probability difference between option K and option A, when option B is also available. For certain values of the parameters (prior mean μ_C and reward uncertainty σ_R^2), this difference is zero with $d = 0$ and increases with the proximity parameter d . This effect disappears when reward uncertainty σ_R^2 is too large or when the prior mean μ_C is too small. Overall, these results show that the compromise effect emerges naturally from BCV, without any *ad-hoc* assumptions and under default values of the parameters (prior mean μ_C and reward uncertainty σ_R^2). Interestingly, these simulations predict a correlation between the compromise effect and the proximity parameter d , reflecting differences between the intermediate and extreme options. This phenomenon is predicted by another model of the compromise effect [19] but remains to be validated empirically.

In summary, these simulations provide proof of principle that BCV predicts within-choice contextual effects during multiattribute decisions that are remarkably similar to those seen in empirical studies. In other words, the similarity, attraction and compromise effects seen empirically are all emergent properties of BCV. In the next section, we turn from within choice effects and consider between-choice context effects.

Between-choice context effects

To characterize between-choice context-effects [11], BCV uses the same generative model as above, characterized by a prior belief μ_C (here we consider only options defined by a single attribute) over reward (with uncertainty σ_C^2) and by an observation of reward amount R (with uncertainty σ_R^2). Here, the generative model is extended to include a Gaussian observation variable O that reflects contextual information provided before an option is presented (Fig 8A). This depends on the hidden cause C and is endowed with uncertainty σ_O^2 (as for the reward amount):

$$O \sim \mathcal{N}(C, \sigma_O^2) \tag{12}$$

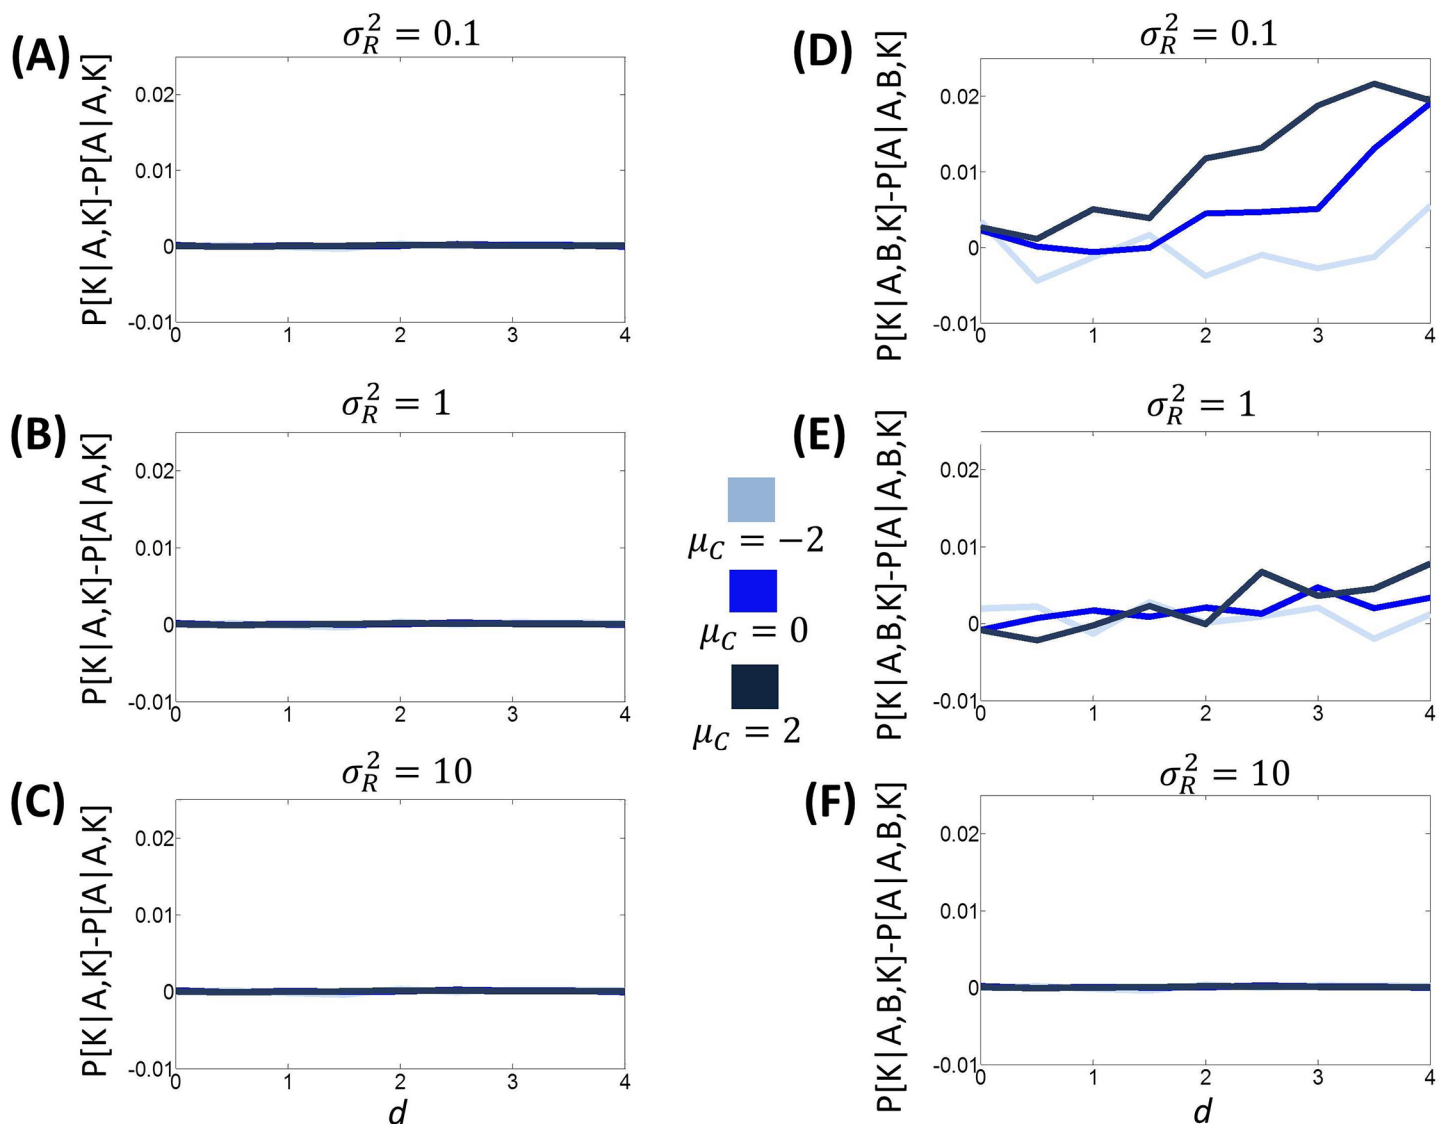


Fig 7. Predictions of BCV related to the compromise effect, involving option A, B and K. Here, in relation with option A, we assign $R_{p,A} = 5 - d$ for price and $R_{q,A} = 5 + d$ for quality; in relation with option B, we assign $R_{p,B} = 5 + d$ for price and $R_{q,B} = 5 - d$ for quality. The proximity parameter d varies across simulations from zero to four. To represent the option K with intermediate scores for both the two attributes, we assign $R_{p,K} = 5$ and $R_{q,K} = 5$ (100000 trials are simulated for each condition; $\sigma_C^2 = 1$ for simulations). **A:** The difference in probability between choosing option K and option A during binary choice ($P[K|A,K] - P[A|A,K]$). The reward uncertainty σ_R^2 is set to 0.1 and different values of the prior mean μ_C are considered. **B:** The same simulation is reported except that the reward uncertainty σ_R^2 was set to one. **C:** The same simulation is reported except that the reward uncertainty σ_R^2 was set to ten. **D:** The difference in probability between choosing option K and option A during choices in which option B is also available ($P[K|A,B,K] - P[A|A,B,K]$). The reward uncertainty σ_R^2 is set to 0.1 and different values of the prior mean μ_C are considered. **E:** The same simulation is reported except that the reward uncertainty σ_R^2 was set to one. **F:** The same simulation is reported except that the reward uncertainty σ_R^2 was set to ten.

<https://doi.org/10.1371/journal.pcbi.1005769.g007>

As above, we assume that an agent infers the posterior expected reward of options afforded by a given context, based on the reward amount but also now on contextual information (i.e., $\hat{\mu}_{C|O,R}$). Since the latter is provided before the option, we assume that the agent infers $\hat{\mu}_{C|0}$ first and then $\hat{\mu}_{C|O,R}$, when the option is presented. Assuming a prior mean equal to zero $\mu_C = 0$, then:

$$\hat{\mu}_{C|0} = \frac{\sigma_C^2}{\sigma_C^2 + \sigma_O^2} O \tag{13}$$

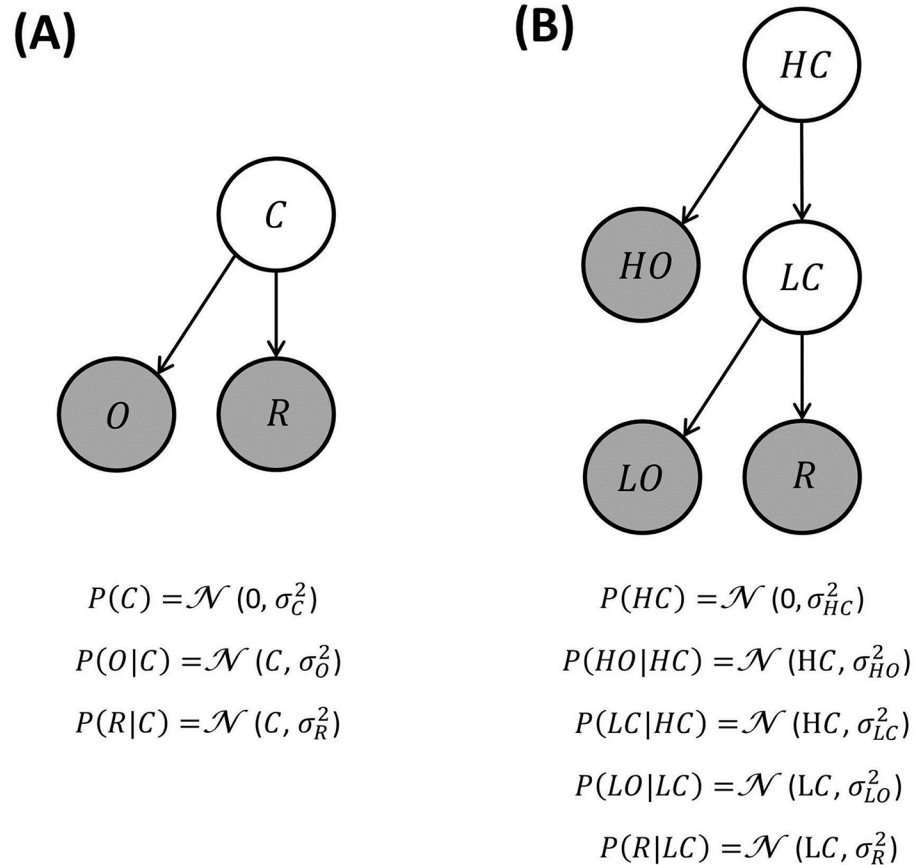


Fig 8. A: Generative model where a contextual variable C reflects a prior expectancy of zero over the reward mean, and a noisy observation O of the context value is provided. **B:** Generative model where context is organized hierarchically and comprises a high level (HC ; e.g., a neighbourhood) and a low level (LC ; e.g., a restaurant), both associated with noisy observations (HO and LO respectively).

<https://doi.org/10.1371/journal.pcbi.1005769.g008>

And the posterior uncertainty:

$$\hat{\sigma}_{C|O}^2 = \sigma_C^2 - \frac{\sigma_C^2}{\sigma_C^2 + \sigma_O^2} \sigma_C^2 \tag{14}$$

The mean of the posterior distribution $P(C|O,R)$ corresponds to:

$$\hat{\mu}_{C|O,R} = \hat{\mu}_{C|O} + \frac{\hat{\sigma}_{C|O}^2}{\hat{\sigma}_{C|O}^2 + \sigma_R^2} (R - \hat{\mu}_{C|O}) \tag{15}$$

Implying the following incentive value for the option:

$$V(R) = \frac{\hat{\sigma}_{C|O}^2}{\hat{\sigma}_{C|O}^2 + \sigma_R^2} (R - \hat{\mu}_{C|O}) \tag{16}$$

This shows that, other things being equal, information about context (reflected in the value of O) induces subtractive value normalization. For instance, when contextual cues O supports a larger reward, $\hat{\mu}_{C|O}$ will be larger and hence the reward prediction error (i.e., $R - \hat{\mu}_{C|O}$) will be smaller.

An extension of this generative model is illustrated in Fig 8B, where contexts are organized hierarchically. Combining the influence of reward expectancies within a hierarchy allows the generative model to explain the impact of context at multiple levels. For instance, the value attributed to a certain dish may depend on the reward distribution associated with a restaurant (a more specific context), integrated with the reward distribution associated with a city (a more general context). In detail, a higher-level prior belief about the average reward amount of options (e.g., at the level of the neighbourhood) is represented by a Gaussian distribution with mean μ_{HC} equal to zero and uncertainty σ_{HC}^2 , from which a value HC is sampled. Contextual information about HC is provided and represented by HO that is sampled from a Gaussian distribution with mean HC and uncertainty σ_{HO}^2 . A lower-level belief about the average reward amount of options (e.g., the restaurant) is represented by a (Gaussian) distribution with mean HC and uncertainty σ_{LC}^2 , from which a value LC is sampled. Contextual information about LC is provided and represented by LO, which is sampled from a Gaussian distribution with mean LC and uncertainty σ_{LO}^2 . A reward is obtained and sampled from a Gaussian distribution with mean LC and uncertainty σ_R^2 .

We propose that agents infer the posterior expectation $\hat{\mu}_{LC|HO,LO,R}$ $P(LC|HO,LO,R)$ sequentially by estimating $\hat{\mu}_{HC|HO}$, $\hat{\mu}_{LC|HO}$, $\hat{\mu}_{HC|HO,LO}$ and finally $\hat{\mu}_{LC|HO,LO,R}$. This produces an equation for incentive value with the following form (see Materials and Methods for derivation):

$$V(R) = K(R - \tau_{LO}LO - \tau_{HO}HO) \tag{17}$$

Three normalization factors are implicit here. The first ($\tau_{LO}LO$) is a subtractive normalization factor proportional to the value LO observed at the low contextual level. The second ($\tau_{HO}HO$) is a subtractive normalization factor proportional to the value HO observed at the high contextual level. The terms τ represent gain-dependent effects and describe the relative precision of information conveyed by the low-level (τ_{LO}) and high-level (τ_{HO}) observations. Finally, a third factor (K) implements divisive normalization and depends on a gain term which includes reward uncertainty (see Materials and Methods for details).

In recent studies [8–11], we have investigated the nature of contextual influence on incentive value that depends on reward expectations established before choice presentation (between-choice effects). In these studies, we have used a simple decision-making task, where participants had to repeatedly choose between a sure monetary reward and a fifty-fifty gamble. These options comprised double the sure monetary reward and a zero outcome, ensuring that the two options had equivalent expected reward or value (EV). Across blocks, we manipulated the distribution of EVs, such that these distributions overlapped. We analysed choice behaviour with EVs common to both contexts to examine whether incentive value attributed to the objective EV changed according to BCV predictions.

In one experiment (Fig 9A and 9B; [8, 9]), in different blocks, the sure monetary gain was drawn from one of two distinct, but partially overlapping, distributions of rewards (low-average and high-average context). Choice behaviour was consistent with attributing a larger incentive value to common EVs in the low average compared to high-average context. This and similar evidence [5–14, 50] suggests that incentive values are, to some extent, rescaled to the average reward expected in a given context, such that they increase (resp. decrease) with smaller (resp. larger) average reward expectations.

These data fit within predictions of BCV. In addition, BCV postulates a between-choice influence of expected reward variance on incentive values (Fig 9C and 9D). In a recent study [11], we used the same gambling task described above and manipulated contextual variance on two levels; one associated with blocks where two target trial EVs were presented (low-variance context), and another with blocks where the same two target trial EVs plus a larger and a

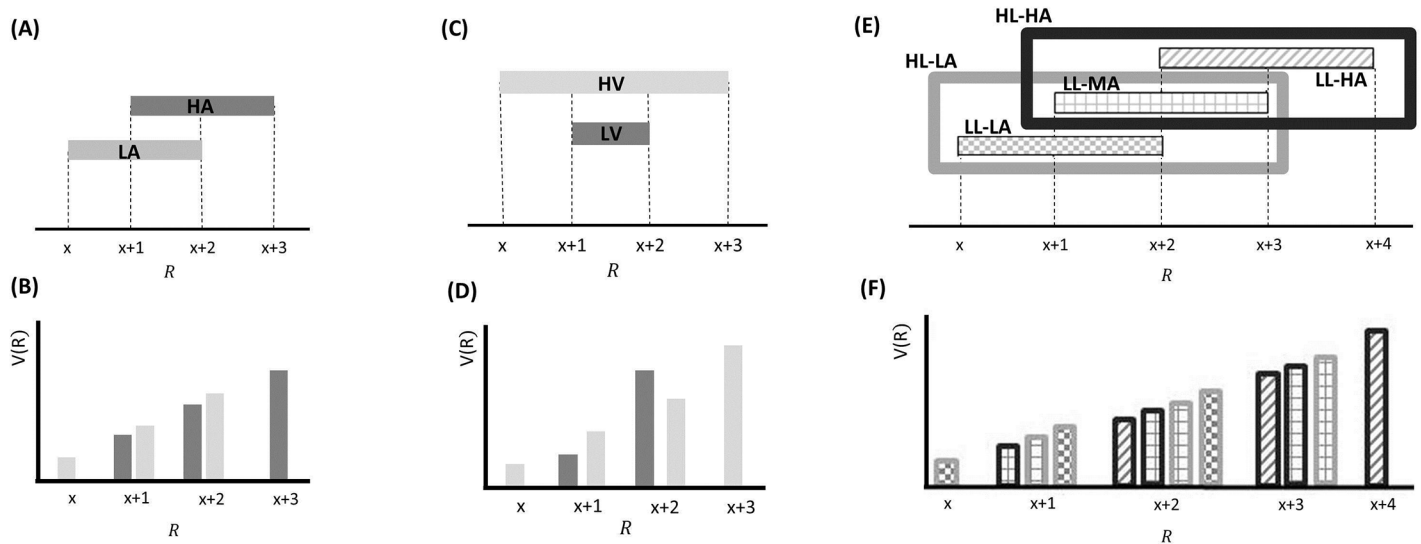


Fig 9. Between-choice effects (predicted by BCV) that depend on contexts with different averages. Contexts are associated with certain distribution of rewards presented sequentially over trials (arranged in blocks for each context). **A:** Example with a single hierarchical level, where two contexts have different average rewards. In blocks associated with a low-average context (LA; in lighter grey), possible rewards are x , $x+1$ and $x+2$; in blocks associated with a high-average context (HA; in darker grey), possible rewards are $x+1$, $x+2$ and $x+3$. **B:** BCV prediction of the incentive value attributed to rewards depending on these contexts. Larger values are predicted in the LA compared to the HA for amounts common to both contexts. **C:** Effects predicted by BCV dependent on contexts with different variance. In blocks associated with a high-variance context (HV; in lighter grey), possible rewards are x , $x+1$ and $x+2$; in blocks associated with a low-variance context (LV; in darker grey), possible rewards are $x+1$ and $x+2$. **D:** BCV prediction of the incentive value attributed to rewards depending on these contexts. Considering rewards common to both contexts, BCV predicts higher incentive value for $x+1$ in the high-variance context and for $x+2$ in the low-variance context. **E:** Example with two hierarchical levels (low-level (LL) contexts, represented by filled rectangles, and high-level (HL) contexts, represented by frames). Blocks associated with HL contexts comprise several sub-blocks associated with LL contexts having specific average reward. In the HL context with low-value (HL-LA; light frame), a LL context with low average (LL-LA, where rewards are x , $x+1$ and $x+2$) and a LL context with medium average (LL-MA, where rewards are $x+1$, $x+2$ and $x+3$) alternate. In the HL context with high-value (HL-HA; dark frame), a LL-MA context and a LL context with high average (LL-HA, where rewards are $x+2$, $x+3$ and $x+4$) alternate. **F:** BCV prediction of the incentive value attributed to rewards depending on these contexts. The fill colour of bars represent the LL context condition, the outline colour represent the HL context condition. BCV predicts that incentive values derive from integrating both hierarchical levels, with larger values emerging when average reward is lower at both context levels.

<https://doi.org/10.1371/journal.pcbi.1005769.g009>

smaller EV were presented (high-variance context). Crucially, this ensured that the two contexts had equivalent average reward but different variance. BCV predicts that the incentive value of the smaller target trial EV will be lower in the low-variance compared to the high-variance context, and the incentive value of the larger target trial EV will be higher in the low-variance compared to the high-variance context. In other words, BCV predicts a larger value difference between the two target trial EV in the low compared to high-variance context. This derives from the gain term, which depends on contextual reward variance. Specifically, low variance magnifies the reward prediction error and hence further reduces the value of rewards that are lower than expected and enhances the value of rewards that are larger than expected. We have previously provided data that are consistent with this prediction [11].

This latter study supports the hypothesis that between-choice reward variance influences incentive value consistent with BCV. In the same study (Fig 9E and 9F; [11]), we also reported that between-choice context effects can be expressed at different hierarchical levels, in line with predictions of BCV. Participants played a computer-based task, where two decks of cards (representing a low-level context) appeared. Each card was associated with a monetary reward, and decks contained cards with different average rewards. A card was drawn from a selected deck and participants had to choose between half of the card reward for sure and a gamble between the full reward and a zero outcome, each with 50% chance. Two sets of decks (representing a high-level context) alternated in a pseudo-random way. The empirical data showed

that the lowest incentive values were attributed when both high-value decks and deck-sets were simultaneously presented, while the highest incentive values were attributed when low-value decks and deck-sets were simultaneously presented. Intermediate incentive values were attributed when decks and deck-sets had one high value and the other low value.

Collectively, these empirical studies provide evidence consistent with between-choice contextual effects on incentive value that depends on beliefs about the average reward and variance expected across choices at multiple hierarchical levels. Furthermore, the empirical findings endorse the predictions derived from BCV.

Discussion

We advance BCV as a unifying theory of contextual effects in value-based choice under the normative principles of Bayesian statistics. BCV assumes that the brain calls on Bayesian inference to invert a generative model and compute (independently for each attribute) the average reward based on observing different reward amounts of options that are available in a given context. Our key proposal is that incentive value emerges during this inferential process, and corresponds to a precision-weighted reward prediction error. Here, we show that these principles are sufficient to explain a wide range of between-choice and within-choice contextual influences; in the latter case encompassing both single and multiattribute effects. To our knowledge, this is the first time a theory has been applied to the full range of context effects.

An important advantage of BCV is its grounding in normative principles of Bayesian statistics [32–37]. Several arguments have been made in support of a Bayesian approach. These are based on a formal and clear definition of the functions that motivate cognitive processes, which are formulated as Bayesian inference and learning. This allows BCV to establish a direct link with Bayesian schemes in other domains—a step towards formulating a unifying theory of brain function. Remarkably, we show that the same basic processes postulated by BCV can be applied to a wide range of conditions in which contextual effects on value and choice are involved. Beyond explaining the available empirical evidence, this scheme can generate new hypotheses (see below). Indeed one of our previous studies [11] was motivated by testing predictions arising out of our initial formulations of BCV.

BCV is associated with planning as inference and active inference [40–45]. The basic idea is that an agent considers the rewards on offer as samples drawn from a population. The latter is not known directly, but can be inferred based on the rewards on offer. Heuristically, agents are interested in inferring how much reward is available on a given trial, which they estimate by combining prior expectations with observations of available rewards. On this view, agents primarily aim to *infer*—and not maximize—the reward; implying that utility-maximization is an emergent process. We argue that an advantage of this perspective is that it offers a normative interpretation of contextual effects, which emerge from the inferential treatment offered here.

Although our theoretical treatment is grounded in Bayesian inference one might argue that the Bayesian gloss is unnecessary to understand the particular inferential mechanisms we have called upon [51]. To a certain extent, there is tautology in Bayesian explanations for behaviour. This follows from the complete class theorem (i.e., for every loss function and behaviour there is a prior belief that renders the behaviour Bayes optimal) [52]. In other words, in principle, everything is Bayes optimal under some priors. This means that the interesting questions reduce to the form of prior beliefs that constitute a subject's generative model. Our focus has been on the form of these models and the particular role of precision weighting in belief updating and choice. The results of our analysis are consistent with empirical data on several forms of context effect, and hence may contribute to a clarification of the computational principles at play.

In BCV incentive value, and in turn choice behaviour, emerges from Bayesian belief updating. Under continuous state space models of the hidden causes of reward values, belief updates and incentive value can be cast as precision-weighted (reward) prediction error. A possibility consistent with BCV is that action is steered by (precision-weighted) prediction errors and is oriented to error cancellation, with approach and avoidance responses elicited by positive and negative prediction errors, respectively. The crucial role of prediction error highlights a perspective in which incentive value is inherently *relative* with respect to reward expectation. Eliciting approach and avoidance behaviour in response to positive and negative prediction errors can be conceived as a basic error-cancellation process (crystallized during evolution of biological organisms), which is a core tenet of active inference schemes.

BCV postulates that the two fundamental determinants of incentive value are prediction error and relative precision. A prediction error is determined by the difference between the observed and expected reward which, in BCV, derives from integrating different expectations under contextual uncertainty. Relative precision depends on the (relative) precision or prior confidence—and ensures that the prediction error is normalised and (Bayes) optimally weighted in relation to uncertainty about both context and reward cues. BCV predicts precision exerts an influence in two ways. First, at high hierarchical levels, precision determines the optimal integration of multiple contextual representations—as it mandates that contexts characterized by a high precision (greater reliability) will exert more influence on reward expectancy. For instance, if we assume that subjects have very precise beliefs about the low-level context (e.g., the card deck in the final experiment on between-choice context effects), then the effect of the high-level (e.g., the deck set) will disappear. Formally, this is because in the hierarchical model the low-level context constitutes a Markov blanket for the posterior expectation about the reward option (Bishop, 2006). In other words, the effect of the high-level context tells us that if subjects are using a hierarchical model, there must be posterior uncertainty about the low-level context. Heuristically, even though they can see which deck they are currently playing with, they still nuance their expectations about this deck based upon the deck-set from which it came. Second, at the lowest hierarchical level, precision determines the gain assigned to the prediction error and hence is a direct determinant of incentive value.

Within BCV, the ratio between reward uncertainty and prior uncertainty determines the gain term (or relative precision) which is used for belief updating (see Eq 6). This means that manipulating the prior uncertainty produces exactly opposite effects compared to manipulating the reward uncertainty, meaning that varying one during simulations is sufficient for testing the predictions of the model (above, we manipulated reward uncertainty and kept the prior uncertainty constant). Thus BCV has only two parameters; namely, the prior mean and reward uncertainty. The role of the latter is straightforward, as context effects are allowed only with small reward uncertainty, and the size of these effects decreases with this reward uncertainty. The role of the prior mean is more complex: for instance, a large prior mean permits a similarity effect but interferes with an attraction effect, while a small prior mean allows an attraction effect but interferes with a similarity effect. Notably, all contextual effects are expressed when setting the prior reward expectation to zero, which can be considered the default value. In short, relying on only two parameters endows BCV with simplicity and constrains the predictions that can be derived, making BCV easy to validate or falsify (see below).

Comparison with other models

We have shown that the principles underlying BCV can explain a wide range of empirical findings on the context sensitivity of value-based choice. Several previous accounts have focused on a single context effect, especially during multiattribute decisions. Some models have been

developed explicitly for explaining the similarity effect [20, 53–55], other models for explaining the attraction effect [56, 57], and other models for the compromise effect [27]. However, a shortcoming of these models is their inability to explain all three effects within a single formal framework. More recently, adopting connectionist architectures, the multi-alternative decision field theory [16, 58, 59] and the leaky competing accumulator [19, 59, 60] have been able to reproduce all three effects (see also [61]). The first model [16, 58] is based on a process modeling attentional switches across attributes and a comparator mechanism which, for the attribute under attention, computes the difference between the reward of each option and the mean reward across options. The second model [19, 59, 60] is similar, except that the comparator applies a non-linear asymmetric (loss-averse) value function to the difference. Although these models fit remarkably with empirical literature and shed light on the neural mechanisms underlying choice, we argue that BCV presents several advantages. First, it is based on normative principles of Bayesian inference. This constrains the model in terms of empirical predictions. In other words, the similarity, attraction and compromise effect are implicit in the way the model works. In fact, these effects arise when default parameters are used. Second, BCV is a more parsimonious model; as the number of free parameters is much lower (essentially, the prior mean and the reward uncertainty). Third, without any further assumptions, BCV applies to a wider range of phenomena including single-attribute decisions and also accounts for between-context effects. Overall, while previous connectionist models are informative especially at the implementation level, BCV helps clarify context sensitivity at the algorithmic and computational level.

The concept of *wealth* in expected utility theory [3] and *status quo* in prospect theory [62] have been recently re-casted in terms of average expected reward [29]. This formulation opens the possibility of context effects dependent on changes in reward expectation. In line with this view, empirical evidence indicates a between-choice context effect that depends on the average contextual reward (as for example inferred from past choices), consisting in attributing larger incentive values in contexts characterized by lower reward. A similar idea has inspired decision by sampling theory [14, 31], which evokes a few basic cognitive processes to explain choice behaviour. According to this model, each choice option elicits retrieval from memory (in the form of random sampling) of stimuli encountered in the past, especially those associated with the current context. A set of binary comparisons follows between the option and the samples, and the number of comparisons in which the option is favoured over each sample is recorded. This number corresponds to the incentive value of the option and is computed for all options available, hence determining their relative preference. Since samples are drawn from memory, they depend on past experience and therefore reflect the distribution of options and outcomes characterizing the environment of an agent. This model can account for an attribution of larger incentive value to the same reward in contexts where lower compared to higher reward is expected before options are provided. This effect is explained by a decreased likelihood, in the former compared to the latter context, of sampling stimuli from memory that are preferred to rewards common to both contexts (assuming a recency effect in memory sampling; [14, 31]). BCV extends these views by appealing explicitly to Bayesian principles (i.e. Bayesian belief updating and evidence accumulation), with implications for empirical predictions. For instance, contrary to BCV and empirical findings, it remains unclear whether these previous models can account for between-choice contextual influence of reward variance or any within-choice contextual effects.

Divisive normalization theory [6, 63–68] has been proposed recently to explain both between-choice and within-choice contextual effects during single attribute decisions. Divisive normalisation was first proposed in the sensory domain to explain phenomena such as neural adaptation within the retina to stimuli of varying intensity [63]. There is evidence that similar

principles can explain higher-order cognitive processes, such as selective attention and perceptual decision-making [63, 69]. Recently, divisive normalisation has been extended to contextual adaptation effects in value-guided choice [6], and proposes that incentive value corresponds to the reward divided by the average reward of past or current choices. This can explain contextual influences elicited both within-choice effects during non-multiattribute decisions and between-choice effects that depend on the average contextual reward. Though this scheme relies on a normalization scheme similar to BCV, different empirical predictions arise. It remains unclear whether this divisive normalization scheme is able to explain between-choice effects deriving from reward variance, and can explain data on multi-attribute choices. In addition, BCV, but not divisive normalization theory, is based on normative principles of Bayesian statistics. However, an attractive aspect of divisive normalization theory is the explicit connection with mechanisms characterizing biological neural processes [63]. A similar connection can be motivated for BCV, given several proposals showing how Bayesian inference (the framework of BCV) is compatible with neuronal processes [49, 70, 71].

The manner in which BCV conceptualizes incentive value is similar to recent economic models that postulate incentive value is adapted to the statistics of the expected reward distribution [29, 30]. These theories can be broadly classified into those based on subtractive normalization, which assume that incentive value corresponds to the reward minus a reference value [29], and those based on divisive normalization, assuming that incentive value corresponds to the reward divided (or multiplied) by the range of an expected distribution of rewards [30]. An important difference between BCV and these theories is the derivation of the former but not the latter from normative assumptions of Bayesian inference. From Bayesian belief updating, BCV derives the proposition that incentive value corresponds to precision-weighted prediction error, hence implying both a subtractive normalization to the expected reward and a divisive normalization with respect to the reward uncertainty. Importantly, these predictions are not *ad hoc* but derive from Bayesian assumptions, distinguish BCV from other models, and have been recently supported empirically [11]. In addition, while these recent economic models focus on between-choice context effects, BCV is more general as it can reproduce within-choice effects in both single and multiattribute decisions.

Like BCV, a recent proposal has interpreted multi-attribute within-choice effects based on the notion that perception of reward is stochastic [72]. The idea is that, for each attribute, an agent forms noisy observations of reward amounts and of the ordinal positions of the reward amounts. Multi-attribute effects can then be obtained by integrating these two observations [72]. Though there are analogies between BCV and the model of Howes et al. [72], we emphasize several important differences. First, the latter does not employ a Bayesian framework, since it is not based on integrating prior beliefs and observations, nor it is based on optimal weighting of different sources of information (as in multi-sensory integration). Second, the model of Howes et al. [72] has been applied to aspects of multi-attribute effects (such as the impact on reaction times), which remain to be explored with BCV. On the other hand, the model of Howes et al., [72] remains to be explored in relation to within-choice effects involving a single attribute and in relation to between-choice effects.

Predictions and limitations of BCV

Specific empirical predictions can be derived from BCV, and here we highlight some of these. Standard economic theories assume that choice should be independent of whether options are presented simultaneously or sequentially. However, the latter case remains largely to be investigated. BCV may inspire this investigation, as it predicts that a higher value will be attributed to an option after presentation of lower value options. This because BCV proposes a sequential

belief updating in which options considered so far contextualize the option observed now. Other predictions involve interactions regarding between- and within-choice effects. For example, consider the example above in which an agent usually evaluates equally car A (expensive and high quality) and car B (cheap and low quality). One may design an experiment where participants are first exposed to a set of cars having a fixed level of quality and varying on price. BCV predicts that this manipulation would determine a lower reward uncertainty for quality compared to price. In other words, quality would become more salient than price, predicting a preference for car A over car B. In addition, BCV predicts other forms of interactions regarding between- and within-choice effects dependent on manipulations of the reward uncertainty and the prior mean (see above), which also remain to be explored empirically. Finally, BCV may be relevant for research on the neural underpinnings of decision-making. A main aspect of this theory is the idea that incentive value corresponds to a precision-weighted reward prediction error. Interestingly, reward prediction error is reflected in activity of brain regions involved in reward processing [73]. BCV raises the possibility that a stimulus which elicits a stronger prediction error response in the brain will be attributed a higher incentive value.

There are shortcomings to BCV, though we argue that the same framework may be fruitfully used to address some of these shortcomings. A shortcoming of our current formulation assumes that model parameters are given. In reality, these parameters need to be learned in the first place. Questions about the mechanisms that might underpin learning of generative models adopted for Bayesian inference are still largely open, though substantial contributions exist, particularly in the context of structure learning [74–80]. A second shortcoming is that here we have assumed that choices occur after inference has considered all observations. An important extension of BCV is a consideration that action tendencies actually develop during evidence accumulation, and this speaks to models of choice that focus on action dynamics, sequential policy optimisation and reaction times [16, 46, 47]. Another important extension of BCV would be to generalize to domains outside incentive value computation. Context effects similar to those observed in value-based decision-making have been reported in many other conditions during perception and judgement [81–84]. Notably, multi-attribute context effects have been recently shown outside incentive value computation [85, 86], suggesting that they may derive from a general way in which the brain works [61].

Conclusions

We offer BCV as a unifying theory of contextual effects during choice behaviour based on Bayesian normative principles. BCV predictions are in line with available empirical evidence about context sensitivity seen empirically both within and between-choice. These different effects are explained using the same simple set of principles, invoking minimal assumptions. We argue that strengths of this model are its foundation on normative principles, simplicity, the link with other influential models of brain function, and the ability to explain a wide range of empirical data. This theory may help clarify the nature of incentive value attribution and choice behaviour. This is particularly prescient when trying to understand ecological phenomena and psychopathologies characterized by dysfunctional choice, such as addiction.

Materials and methods

Here we derive Eq 17 from the generative model shown in Fig 9B. A higher-level contextual variable (e.g., a neighbourhood containing several restaurants) is represented by a Gaussian distribution with mean μ_{HC} equal to zero and uncertainty σ_{HC}^2 , from which a value HC is sampled. Sensory evidence about HC is provided and represented by HO which is sampled from a Gaussian distribution with mean HC and uncertainty σ_{HO}^2 . A lower-level contextual variable

(e.g., one of the restaurants) is represented by a (Gaussian) distribution with mean HC and uncertainty σ_{LC}^2 , from which a value LC is sampled. Sensory evidence about LC is provided and represented by LO, which is sampled from a Gaussian distribution with mean LC and uncertainty σ_{LO}^2 . A reward is obtained and sampled from a Gaussian distribution with mean LC and uncertainty σ_R^2 . The posterior distribution $P(LC|HO,LO,R)$ can be inferred sequentially in the order $P(HC|HO)$, $P(LC|HO)$, $P(LC|HO,LO)$, and $P(LC|HO,LO,R)$. The posterior mean of $P(HC|HO)$ is:

$$\hat{\mu}_{HC|HO} = \frac{\sigma_{HC}^2}{\sigma_{HC}^2 + \sigma_{HO}^2} HO \tag{18}$$

And the posterior uncertainty:

$$\hat{\sigma}_{HC|HO}^2 = \sigma_{HC}^2 - \frac{\sigma_{HC}^2}{\sigma_{HC}^2 + \sigma_{HO}^2} \sigma_{HC}^2 \tag{19}$$

The posterior mean of $P(LC|HO)$ is equal to $\hat{\mu}_{LC|HO}$ ($\hat{\mu}_{LC|HO} = \hat{\mu}_{HC|HO}$), while the posterior uncertainty is:

$$\hat{\sigma}_{LC|HO}^2 = \hat{\sigma}_{HC|HO}^2 + \sigma_{LC}^2 \tag{20}$$

The posterior mean of $P(LC|HO,LO)$ is:

$$\hat{\mu}_{LC|HO,LO} = \hat{\mu}_{LC|HO} + \frac{\hat{\sigma}_{LC|HO}^2}{\hat{\sigma}_{LC|HO}^2 + \sigma_{LO}^2} (LO - \hat{\mu}_{LC|HO}) \tag{21}$$

And the posterior uncertainty:

$$\hat{\sigma}_{LC|HO,LO}^2 = \hat{\sigma}_{LC|HO}^2 - \frac{\hat{\sigma}_{LC|HO}^2}{\hat{\sigma}_{LC|HO}^2 + \sigma_{LO}^2} \hat{\sigma}_{LC|HO}^2 \tag{22}$$

The posterior mean of $P(LC|HO,LO,R)$ is:

$$\hat{\mu}_{LC|HO,LO,R} = \hat{\mu}_{LC|HO,LO} + \frac{\hat{\sigma}_{LC|HO,LO}^2}{\hat{\sigma}_{LC|HO,LO}^2 + \sigma_R^2} (R - \hat{\mu}_{LC|HO,LO}) \tag{23}$$

Finally, with few rearrangements, we obtain the following incentive value for a reward offer:

$$V(R) = \frac{\hat{\sigma}_{LC|HO,LO}^2}{\hat{\sigma}_{LC|HO,LO}^2 + \sigma_R^2} \left(R - \frac{\hat{\sigma}_{LC|HO}^2}{\hat{\sigma}_{LC|HO}^2 + \sigma_{LO}^2} LO - \frac{\sigma_{LO}^2}{\hat{\sigma}_{LC|HO}^2 + \sigma_{LO}^2} \frac{\sigma_{HC}^2}{\sigma_{HC}^2 + \sigma_{HO}^2} HO \right) \tag{24}$$

This equation implements three normalization factors: (i) a subtractive normalization factor $\left(\frac{\hat{\sigma}_{LC|HO}^2}{\hat{\sigma}_{LC|HO}^2 + \sigma_{LO}^2} LO \right)$ proportional to the value LO observed at the low contextual level, (ii) a subtractive normalization factor $\left(\frac{\sigma_{LO}^2}{\hat{\sigma}_{LC|HO}^2 + \sigma_{LO}^2} \frac{\sigma_{HC}^2}{\sigma_{HC}^2 + \sigma_{HO}^2} HO \right)$ proportional to the value HO observed at the high contextual level, (iii) a divisive normalization factor $\left(\frac{\hat{\sigma}_{LC|HO,LO}^2}{\hat{\sigma}_{LC|HO,LO}^2 + \sigma_R^2} \right)$ that captures the weighting dependent on the (relative) reward uncertainty. If we define the three factors as τ_{LO} and τ_{HO} and K respectively, we obtain Eq 17.

Author Contributions

Conceptualization: Francesco Rigoli, Christoph Mathys, Karl J. Friston, Raymond J. Dolan.

Formal analysis: Francesco Rigoli, Christoph Mathys, Karl J. Friston.

Funding acquisition: Raymond J. Dolan.

Methodology: Francesco Rigoli, Christoph Mathys, Karl J. Friston.

Project administration: Francesco Rigoli.

Software: Francesco Rigoli.

Supervision: Francesco Rigoli, Karl J. Friston, Raymond J. Dolan.

Validation: Francesco Rigoli.

Visualization: Francesco Rigoli.

Writing – original draft: Francesco Rigoli, Christoph Mathys, Karl J. Friston, Raymond J. Dolan.

Writing – review & editing: Francesco Rigoli, Christoph Mathys, Karl J. Friston, Raymond J. Dolan.

References

1. Bernoulli D. Specimen theoriae novae de mensura sortis (Exposition of a new theory on the measurement of risk). *Comentarii Acad. Scient. Petropolis* (translated in *Econometrica*). 1738; 5: 23–36.
2. Luce RD. On the possible psychophysical laws. *Psychol Rev.* 1959; 66: 81. PMID: [13645853](#)
3. von Neumann J, Morgenstern O. *Theory of Games and Economic Behavior*. 1944; Princeton: Princeton UP.
4. Vlaev I, Chater N, Stewart N, Brown GD. Does the brain calculate value?. *Trends Cogn Sci*, 2011, 15: 546–554. <https://doi.org/10.1016/j.tics.2011.09.008> PMID: [21983149](#)
5. Ludvig EA, Madan CR, Spetch ML. Extreme outcomes sway risky decisions from experience. *J Behav Decis Mak.* 2013;
6. Louie K, Khaw MW, Glimcher PW. Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci U S.* 2013; 110: 6139–6144.
7. Morgan KV, Hurlly TA, Bateson M, Asher L, Healy SD. Context-dependent decisions among options varying in a single dimension. *Behav Process.* 2012; 89: 115–120.
8. Rigoli F, Rutledge RB, Dayan P, Dolan RJ. The influence of contextual reward statistics on risk preference. *NeuroImage.* 2016; 128: 74–84. <https://doi.org/10.1016/j.neuroimage.2015.12.016> PMID: [26707890](#)
9. Rigoli F, Rutledge RB, Chew B, Ousdal OT, Dayan P, Dolan RJ. Dopamine Increases a Value-Independent Gambling Propensity. *Neuropsychopharmacolog.* 2016; 41: 2658–2667.
10. Rigoli F, Friston KJ, Dolan RJ. Neural processes mediating contextual influences on human choice behaviour. *Nat Commun.* 2016; 7, 12416. <https://doi.org/10.1038/ncomms12416> PMID: [27535770](#)
11. Rigoli F, Friston KJ, Martinelli C, Selaković M, Shergill SS, Dolan RJ. A Bayesian model of context-sensitive value attribution. *eLife.* 2016; 5, e16127. <https://doi.org/10.7554/eLife.16127> PMID: [27328323](#)
12. Simonsohn U. New Yorkers commute more everywhere: contrast effects in the field. *Rev Econ Stat.* 2006; 88: 1–9.
13. Simonsohn U, Loewenstein G. Mistake# 37: The Effect of Previously Encountered Prices on Current Housing Demand*. *Econ J.* 2006; 116: 175–199.
14. Stewart N. Decision by sampling: The role of the decision environment in risky choice. *Q J Exp Psychol.* 2009; 62: 1041–1062.
15. Huber J, Payne JW, Puto C. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *J Cons Res.* 1982; 43: 90–98.
16. Roe RM, Busemeyer JR, Townsend JT. Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychol Rev.* 2001; 108: 370. PMID: [11381834](#)

17. Simonson I, Tversky A. Choice in context: Tradeoff contrast and extremeness aversion. *J Mark Res*. 1992; 29: 281.
18. Soltani A, De Martino B, Camerer C. A range-normalization model of context-dependent choice: a new model and evidence. *PLoS Comput Biol*. 2012; 8, e1002607–e1002607. <https://doi.org/10.1371/journal.pcbi.1002607> PMID: 22829761
19. Tsetsos K, Usher M, Chater N. Preference reversal in multiattribute choice. *Psychol Rev*, 2010; 117: 1275. <https://doi.org/10.1037/a0020580> PMID: 21038979
20. Tversky A. Elimination by aspects: A theory of choice. *Psychol Rev*. 1972; 79: 281.
21. Batsell RR, Polking JC. A new class of market share models. *Mark Sci*. 1985; 4: 177–198.
22. Lehmann DR, Pan Y. Context effects, new brand entry, and consideration sets. 1994; *J Mark Res*, 364–374.
23. Sjöberg L. Choice frequency and similarity. *Scand J Psychol*. 1977; 18: 103–115.
24. Ratneshwar S, Shocker AD, Stewart DW. Toward understanding the attraction effect: The implications of product stimulus meaningfulness and familiarity. *J Cons Res*. 1987; 13: 520–533.
25. Simonson I. Choice based on reasons: The case of attraction and compromise effects. *J Cons Res*. 1989; 16: 158–174.
26. Wedell DH. Distinguishing among models of contextually induced preference reversals. *J Exp Psychol Learn Mem Cogn*. 1991; 17: 767.
27. Tversky A, Simonson I. Context-dependent preferences. *Manag Sci*. 1993; 39: 1179–1189.
28. Pettibone JC, Wedell DH. Testing alternative explanations of phantom decoy effects. *J Behav Dec Mak*. 2007; 20: 323–341.
29. Köszegi B, Rabin M. A model of reference-dependent preferences. *Q J Econ*. 2006; 20: 1133–1165.
30. Köszegi B, Szeidl A. A model of focusing in economic choice. *Q J Econ*. 2013; 128: 53–104.
31. Stewart N, Chater N, Brown GD. Decision by sampling. *Cogn Psychol*. 2006; 53: 1–26. <https://doi.org/10.1016/j.cogpsych.2005.10.003> PMID: 16438947
32. Chater N, Tenenbaum JB, Yuille A. Probabilistic models of cognition: Conceptual foundations. *Trends Cogn Sci*. 2006; 10: 287–291. <https://doi.org/10.1016/j.tics.2006.05.007> PMID: 16807064
33. Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci*. 2013; 36: 181–204. <https://doi.org/10.1017/S0140525X12000477> PMID: 23663408
34. Oaksford M, Chater N. *Bayesian rationality: The probabilistic approach to human reasoning*. 2007; Oxford University Press.
35. Oaksford M, Chater N. *Précis of Bayesian rationality: The probabilistic approach to human reasoning*. *Behav Brain Sci*. 2009; 32: 69–84. <https://doi.org/10.1017/S0140525X09000284> PMID: 19210833
36. Dayan P, Hinton GE, Neal RM, Zemel RS. The Helmholtz machine. *Neural Comput*. 1995; 7: 889–904. PMID: 7584891
37. Friston K. The free-energy principle: a unified brain theory?. *Nature Rev Neurosci*. 2010; 11: 127–138.
38. Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*. 2002; 415: 429–43. <https://doi.org/10.1038/415429a> PMID: 11807554
39. Bastos AM, Vezoli J, Bosman CA, Schoffelen JM, Oostenveld R, Dowdall JR, De Weerd P, Kennedy H, Fries P. Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*. 2015; 85: 390–401. <https://doi.org/10.1016/j.neuron.2014.12.018> PMID: 25556836
40. Botvinick M, Toussaint M. Planning as inference. *Trends Cogn Sci*. 2012; 16: 485–488. <https://doi.org/10.1016/j.tics.2012.08.006> PMID: 22940577
41. Friston KJ, Schwartenbeck P, FitzGerald T, Moutoussis M, Behrens T, Dolan RJ. The anatomy of choice: active inference and agency. *Front Hum Neurosci*. 2013; 7.
42. Friston KJ, Rigoli F, Ognibene D, Mathys C, Fitzgerald T, Pezzulo G. Active inference and epistemic value. *Cogn Neurosci*. 2015; 2: 1–28.
43. Pezzulo G, Rigoli F. The value of foresight: how prospection affects decision-making. *Front Neurosci*. 2011; 5.
44. Pezzulo G, Rigoli F, Friston KJ. Active Inference, homeostatic regulation and adaptive behavioural control. *Prog Neurobiol*. 2015; 134: 17–35. <https://doi.org/10.1016/j.pneurobio.2015.09.001> PMID: 26365173
45. Solway A, Botvinick M. Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. *Psychol Rev*. 2012; 119: 120. <https://doi.org/10.1037/a0026435> PMID: 22229491

46. Johnson JG, Busemeyer JR. A dynamic, stochastic, computational model of preference reversal phenomena. *Psychol Rev.* 2005; 112: 841. <https://doi.org/10.1037/0033-295X.112.4.841> PMID: 16262470
47. Ratcliff R, McKoon G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 2008; 20: 873–922. <https://doi.org/10.1162/neco.2008.12-06-420> PMID: 18085991
48. Bishop CM. *Pattern recognition and machine learning.* 2006; Springer: New York.
49. Friston KJ. A theory of cortical responses. *Phil Trans Royal Soc B.* 2005; 360: 815–836.
50. Stewart N, Chater N, Stott HP, Reimers S. Prospect relativity: how choice options influence decision under risk. *J Exp Psychol Gen.* 2003; 132: 23. PMID: 12656296
51. Bowers JS, Davis CJ. Bayesian just-so stories in psychology and neuroscience. *Psychol Bull.* 2012; 138: 389. <https://doi.org/10.1037/a0026450> PMID: 22545686
52. Brown LD. A complete class theorem for statistical problems with finite sample spaces. *Ann Stat.* 1981; 1289–1300.
53. Candel MJ. A probabilistic feature model for unfolding tested for perfect and imperfect nestings. *J Math Psychol.* 1997; 41: 414–430. PMID: 9473403
54. Edgell SE, Geisler WS. A set-theoretic random utility model of choice behavior. *J Math Psychol.* 1980; 21: 265–278.
55. Mellers BA, Biagini K. Similarity and choice. *Psychol Rev.* 1994; 101: 505.
56. Ariely D, Wallsten TS. Seeking subjective dominance in multidimensional space: An explanation of the asymmetric dominance effect. *Organ Behav Hum Decis Process.* 1995; 63: 223–232.
57. Dhar R, Glazer R. Similarity in context: Cognitive representation and violation of preference and perceptual invariance in consumer choice. *Organ Behav Hum Decis Process.* 1996; 67: 280–293.
58. Hotaling JM, Busemeyer JR, Li J. Theoretical developments in decision field theory: comment on Tsetsos, Usher, and Chater (2010). *Psychol Rev.* 2010; 117: 1294–1298. <https://doi.org/10.1037/a0020401> PMID: 21038981
59. Usher M, Tsetsos K, Chater N. Postscript: Contrasting predictions for preference reversal. *Psychol Rev.* 2010; 117: 1291–1293.
60. Usher M, McClelland JL. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol Rev.* 2004; 111: 757. <https://doi.org/10.1037/0033-295X.111.3.757> PMID: 15250782
61. Trueblood JS, Brown SD, Heathcote A. The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychol Rev.* 2014; 121: 179. <https://doi.org/10.1037/a0036137> PMID: 24730597
62. Kahneman D, Tversky A. Prospect theory: An analysis of decision under risk. *Econometrica.* 1979; 263–291.
63. Carandini M, Heeger DJ. Normalization as a canonical neural computation. *Nat Rev Neurosci.* 2012; 13: 51–62.
64. Louie K, LoFaro T, Webb R, Glimcher PW. Dynamic divisive normalization predicts time-varying value coding in decision-related circuits. *J Neurosci.* 2014; 34: 16046–16057. <https://doi.org/10.1523/JNEUROSCI.2851-14.2014> PMID: 25429145
65. Louie K, Glimcher PW, Webb R. Adaptive neural coding: from biological to behavioral decision-making. *Curr Opin Behav Sci.* 2015; 5: 91–99. <https://doi.org/10.1016/j.cobeha.2015.08.008> PMID: 26722666
66. Rangel A, Clithero JA. Value normalization in decision making: theory and evidence. *Curr Opin Neurobiol.* 2012; 22: 970–981. <https://doi.org/10.1016/j.conb.2012.07.011> PMID: 22939568
67. Summerfield C, Tsetsos K. Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Front Neurosci.* 2012; 6.
68. Summerfield C, Tsetsos K. Do humans make good decisions?. *Trends Cogn Sci.* 2015 19: 27–34. <https://doi.org/10.1016/j.tics.2014.11.005> PMID: 25488076
69. Cheadle S, Wyart V, Tsetsos K, Myers N, De Gardelle V, Castañón SH, Summerfield C. Adaptive gain control during human perceptual choice. *Neuron.* 2014; 81: 1429–1441. <https://doi.org/10.1016/j.neuron.2014.01.020> PMID: 24656259
70. Hennequin G, Aitchison L, Lengyel M. Fast Sampling-Based Inference in Balanced Neuronal Networks. *Adv Neural Inf Process Syst.* 2014; 2240–2248.
71. Knill DC, Pouget A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 2004; 27: 712–719. <https://doi.org/10.1016/j.tins.2004.10.007> PMID: 15541511
72. Howes A, Warren PA, Farmer G, El-Deredy W, Lewis RL. Why contextual preference reversals maximize expected value. *Psychol Rev.* 2016; 123: 368. <https://doi.org/10.1037/a0039996> PMID: 27337391

73. Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275: 1593–1599. PMID: [9054347](https://pubmed.ncbi.nlm.nih.gov/9054347/)
74. Acuña DE, Schrater P. Structure Learning in Human Sequential Decision-Making. *PLOS Comput Biol*. 2010; 6, <https://doi.org/10.1371/journal.pcbi.1001003> PMID: [21151963](https://pubmed.ncbi.nlm.nih.gov/21151963/)
75. Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat Neurosci*. 2007; 10: 1214–1221. <https://doi.org/10.1038/nn1954> PMID: [17676057](https://pubmed.ncbi.nlm.nih.gov/17676057/)
76. Collins AG, Frank MJ. Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychol Rev*. 2013; 120: 190.
77. Courville AC, Daw ND, Touretzky DS. Bayesian theories of conditioning in a changing world. *Trends Cogn Sci*. 2006; 10: 294–300. <https://doi.org/10.1016/j.tics.2006.05.004> PMID: [16793323](https://pubmed.ncbi.nlm.nih.gov/16793323/)
78. FitzGerald T, Dolan RJ, Friston KJ. Model averaging, optimal inference, and habit formation. *Front Hum Neurosci*. 2014; 8.
79. Gershman SJ, Niv Y. Learning latent structure: carving nature at its joints. *Curr Opin Neurobiol*. 2010; 20: 251–256. <https://doi.org/10.1016/j.conb.2010.02.008> PMID: [20227271](https://pubmed.ncbi.nlm.nih.gov/20227271/)
80. Mathys C, Daunizeau J, Friston KJ, Stephan KE. A Bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci*; 2011; 5: 39. <https://doi.org/10.3389/fnhum.2011.00039> PMID: [21629826](https://pubmed.ncbi.nlm.nih.gov/21629826/)
81. Parducci A. Category judgment: a range-frequency model. *Psychol Rev*. 1965; 72: 407. PMID: [5852241](https://pubmed.ncbi.nlm.nih.gov/5852241/)
82. Parducci A. Happiness, pleasure, and judgment: The contextual theory and its applications. 1995; Lawrence Erlbaum Associates, Inc.
83. Maltby J, Wood AM, Vlaev I, Taylor MJ, Brown GD. Contextual effects on the perceived health benefits of exercise: The exercise rank hypothesis. *J Sport Exerc Psychol*. 2012; 34: 828–841. PMID: [23204361](https://pubmed.ncbi.nlm.nih.gov/23204361/)
84. Watkinson P, Wood AM, Lloyd DM, Brown GD. Pain ratings reflect cognitive context: A range frequency model of pain perception. *Pain*. 2013; 154: 743–749. <https://doi.org/10.1016/j.pain.2013.01.016> PMID: [23498366](https://pubmed.ncbi.nlm.nih.gov/23498366/)
85. Trueblood JS. Multialternative context effects obtained using an inference task. *Psychon Bull Rev*. 2012; 19: 962–968. <https://doi.org/10.3758/s13423-012-0288-9> PMID: [22736437](https://pubmed.ncbi.nlm.nih.gov/22736437/)
86. Trueblood JS, Brown SD, Heathcote A, Bussemeyer JR. Not just for consumers context effects are fundamental to decision making. *Psychol Sci*. 2013; 24: 901–908. <https://doi.org/10.1177/0956797612464241> PMID: [23610134](https://pubmed.ncbi.nlm.nih.gov/23610134/)