Choice-confirmation bias and gradual perseveration in human reinforcement learning

Stefano Palminteri^{1,2,3*}

(1) Laboratoire de Neurosciences Cognitives et Computationnelles, INSERM, Paris, France

(2) Département d'Etudes Cognitives, ENS, PSL Research University, Paris, France

(3) Institute of Cognitive Neuroscience, HSE, Moscow, Federation of Russia

*Corresponding author: stefano.palminteri@ens.fr

Summary

Do we preferentially learn from outcomes that confirm our choices? This is one of the most basic, and yet consequence-bearing, questions concerning reinforcement learning. In recent years, we investigated this question in a series of studies implementing increasingly complex behavioral protocols. The learning rates fitted in experiments featuring partial or complete feedback, as well as free and forced choices, were systematically found to be consistent with a choice-confirmation bias. This result is robust across a broad range of outcome contingencies and response modalities. One of the prominent behavioral consequences of the confirmatory learning rate pattern is choice hysteresis: that is the tendency of repeating previous choices, despite contradictory evidence. As robust and replicable as they have proven to be, these findings were (legitimately) challenged by a couple of studies pointing out that a choice-confirmatory pattern of learning rates may spuriously arise from not taking into consideration an explicit choice autocorrelation term in the model. In the present study, we re-analyze data from four previously published papers (in total nine experiments;

27 N=363), originally included in the studies demonstrating (or criticizing) the choice-
28 confirmation bias in human participants. We fitted two models: one featured valence-
29 specific updates (i.e., different learning rates for confirmatory and disconfirmatory
30 outcomes) and one additionally including an explicit choice autocorrelation process
31 (gradual perseveration). Our analysis confirms that the inclusion of the gradual
32 perseveration process in the model significantly reduces the estimated choice-
33 confirmation bias. However, in all considered experiments, the choice-confirmation
34 bias remains present at the meta-analytical level, and significantly different from zero
35 in most experiments. Our results demonstrate that the choice-confirmation bias resists
36 the inclusion of an explicit choice autocorrelation term, thus proving to be a robust
37 feature of human reinforcement learning. We conclude by discussing the
38 psychological plausibility of the gradual perseveration process in the context of these
39 behavioral paradigms and by pointing to additional computational processes that may
40 play an important role in estimating and interpreting the computational biases under
41 scrutiny.

42 Introduction

43 A few years ago, evidence in favor of optimistic bias in updating probabilistic beliefs
 44 motivated us to address if a similar bias applied to simple reinforcement learning from
 45 positive and negative prediction errors [1, 2]. Robust observations of dissociable
 46 neural systems for reward and punishment errors further supported the hypothesis of
 47 dissociable networks for positive and negative reward prediction errors [3]. Testing this
 48 hypothesis within the standard reinforcement learning formalism (where value update
 49 is driven by reward prediction errors) translates into postulating that different learning
 50 rates apply to learning from positive and negative reward prediction errors (*RPE*;
 51 valence-specific update). Specifically, in a given trial:

52

$$53 \quad RPE(c) = R(c) - Q(s, c)$$

$$54 \quad \text{if } RPE(c) > 0$$

$$55 \quad \quad Q(s, c) \leftarrow Q(s, c) + \alpha_{c+} \times RPE(c)$$

$$56 \quad \text{elseif } RPE(c) < 0$$

$$57 \quad \quad Q(s, c) \leftarrow Q(s, c) + \alpha_{c-} \times RPE(c)$$

58 end

59

60 Where s denotes a learning context (or state; generally, a fixed pair of options
 61 presented together in two-armed bandit tasks), c denotes the chosen option, α_{c+} and
 62 α_{c-} are respectively the learning rates for positive and negative chosen option
 63 prediction errors, and $R(c)$ is the reward associated with the chosen option. At a very
 64 minimum, testing the presence of a valence-specific update *bias* in reinforcement
 65 learning, would therefore consist in two steps: first, showing that a model including two
 66 learning parameters is justified in terms of parsimony (i.e., it provides a better fit,

despite its increased complexity); second, showing that, at the group level, there is a systematic difference between the two learning rates in a specific direction. Before our studies, several previous studies fitted the above-described models, however no clear answer had emerged concerning the direction of the effect, with some studies reporting a positivity (or optimistic) bias and some other studies a negative (or pessimistic) bias [4–9].

We therefore believed that the time was ripe for a more systematic investigation of the question. We started by re-analyzing a previously acquired dataset [10] that presented several advantage: I) the task featured only binary rewards, thus avoiding any ambiguity concerning what represents a good or a bad outcome; II) the task featured some critical conditions (detailed later) allowing testing some predictions of the models; III) fMRI data was available, thus allowing testing for neural correlates of putative learning rate biases. In this study, the results were clear-cut: the valence-specific update model outperformed the simple model with only one learning rate and the bias was in the direction of a positivity (or optimistic bias) [11]. Beyond parsimony-driven model comparison (BIC in the first study), we also pointed to a couple of behavioral signatures falsifying the unbiased, standard model [12]. First, in learning contexts where there was no correct response (i.e., both options were associated with the same probability of getting a reward: ‘symmetric’ conditions), subjects tended to develop over time a preferred response. Second, in the more standard learning contexts (i.e., one option is more rewarding than the other; ‘asymmetric’ condition), the distribution of the correct (i.e., reward maximizing) choice rate across subjects showed signs of convexity with most subjects performing very well (close to 100%)

and few of them performing very badly (close to 0%)¹. Both these results were captured by the model with a positivity bias, which induces an overestimation of the value of the most frequently chosen option [13]. The positivity bias was relatively variable across subjects, and we found neural correlates of this variability at the outcome encoding stage: the higher the subject-level positivity bias the higher the activation of the reward system (specifically the ventral striatum²). The two possible outcomes of this first experiment were 0.5€ and 0.0€, thus leaving the possibility open that the smaller learning rate for negative prediction errors was due to the relatively low saliency of the neutral outcome. To rule out this possibility, we ran a fully incentivized experiment³, featuring exactly the same conditions, but replacing the neutral outcome with a monetary loss (-0.5€). The positivity bias was found also in this experiment (if anything, bigger), thus ruling out the possibility that it was due to lower saliency and weak motivation.

While wrapping up the first paper on this topic, we tested an extended model allowing for a learning asymmetry also for the unchosen option outcome a dataset - featuring a slightly more complex design including separate reward seeking and punishment avoidance contexts [14]. In this augmented model the update of the unchosen option was implemented as follows:

$$RPE(u) = R(u) - Q(s, u)$$

$$\text{If } RPE(u) > 0$$

¹ This result is presented in the supplementary materials of the original paper [11].

² The t-map of a between subject contrast between biased and unbiased subjects can be observed here <https://neurovault.org/images/42831/>

³ This experiment was run in an experimental economics laboratory. Experimental economics laboratories have strict rules against deception and in favor of true incentivization.

113 $Q(s, u) \leftarrow Q(s, u) + \alpha_{u+} \times RPE(u)$

114 *elseif* $RPE(u) < 0$

115 $Q(s, u) \leftarrow Q(s, u) + \alpha_{u-} \times RPE(u)$

116 *end*

117

118 Where u denotes the unchosen (or forgone) option (the notation is otherwise the same
 119 as before). This exploratory analysis delivered a clear-cut result: the bias was in the
 120 opposite direction ($\alpha_{u-} > \alpha_{u+}$). We reasoned that the the overall pattern was therefore
 121 consistent with a choice-confirmation⁴ bias. Motivated by this semi-serendipitous⁵
 122 observation, we ran a replication experiment. This time with a dedicated experimental
 123 design involving complete feedback, binary outcomes and (in addition to symmetrical
 124 and asymmetrical reward conditions as before), also a reversal learning condition [15].
 125 The study also included a partial feedback experiment to provide a third replication of
 126 the main finding. The results of this new experiment corroborated our first observation
 127 and showed that in the unchosen option outcome the negative learning rate was higher
 128 than the positive one ($\alpha_{u-} > \alpha_{u+}$). Suspecting that the approach of four learning rates
 129 was not parsimonious enough, we compared it with reduced versions, including a
 130 ‘Confirmation’ model⁶ where we supposed that:

131

132 $\alpha_{c+} = \alpha_{u-} = \alpha_{CON}$

133

⁴ Positive obtained outcome and negative forgone outcome confirms our current choice. Note that our choice-confirmation bias is different and in principle dissociable from a belief-confirmation bias [28].

⁵ We were looking for learning biases, but we were agnostic concerning the directionality of the bias. This stance is made explicit in the introduction of Palminteri et al. [15].

⁶ We label this model ‘Confirmation’ for simplicity, because the overall results go in the direction of the confirmation bias. However, please note that in principle the results could go in the direction of a disconfirmation bias (as it is the case for a small subset of participants in each experiment)

and

$$\alpha_{c-} = \alpha_{u+} = \alpha_{DIS}$$

The ‘Confirmation’ model clearly outperformed the model with four learning rates⁷ (as well as the other reduced models we tested). The second study also allowed us to identify another behavioral signature of this confirmation bias: a reluctance to switch option in the reversal condition.

In the latest published chapter of this line of research, we manipulated subjects’ possibility to choose freely [16]. More precisely, we interleaved the standard ‘free choice’ trials, with ‘forced choice’ trials where the computer randomly picked one of the two options. Learning rate analysis indicated that the learning biases were specific to outcomes following freely chosen options. This study included a fourth experiment where we manipulated motor requirements. In each trial, one of the two options had to be chosen by a ‘no go’ response (i.e., not doing anything). The confirmation bias was not affected by the motoric requirement of the choice (i.e., the bias is in the ‘decision space’, rather than the ‘action space’).

Since then, several other studies replicated the positivity, or confirmatory, pattern of learning rates in a broad range of tasks and species, indicating that it is a robust computational result [17–19]. However, as Katahira correctly pointed out in a simulation study, a confirmation bias may spuriously arise, if the true generative model

⁷ Of note, this was not the case in study which replicated the confirmation bias learning rate pattern, but found that a model with four learning rates was preferable to the confirmation bias model (two learning rates) [53].

includes an explicit choice-autocorrelation term, which is not included in the fitted model [20].

To illustrate this point, Katahira, adapted a model of gradual choice perseveration based on previous studies in perceptual decision making to the reinforcement learning scenario [21]. The key idea of the model is that it explicitly tracks for each option a choice trace $C(s, :)$, which is incremented/decremented each time that the option is chosen/unchosen. Choice trace update can be written as follows

$$CPE(c) = 1 - C(s, c)$$

$$CPE(u) = 0 - C(s, u)$$

$$C(s, c) \leftarrow C(s, c) + \tau \times CPE(c)$$

$$C(s, u) \leftarrow C(s, u) + \tau \times CPE(u)$$

Where the $CPE(:)$ can be conceptualized as a choice prediction error for the chosen (c) and unchosen option (u). The τ is an analogue to a learning rate and can be defined as the choice trace accumulation rate. Crucially, in this formulation, choice traces are attached to options in an abstract space without any reference to motor components of the choice (we will return to the plausibility of this assumption in the discussion). The choice traces are then used to bias action selection and accordingly the probability of choosing, say option a over b is given by the following formula:

$$P(a) = 1 / (1 + \exp(-\beta * (Q(s, a) - Q(s, b)) - \phi * (C(s, a) - C(s, b))))$$

Where β is the choice inverse temperature and ϕ can be defined as the choice trace decision bias, because it determines how much choice traces influence choice probabilities. Noteworthy is, unlike β , ϕ can take negative values, thus promoting choice alternation. Katahira also discusses a special case of the model (where $\tau = 1$), that he labels the ‘impulsive’ perseveration model, which is more often used in the literature [22] and only takes into account the last choice. However, the impulsive perseveration process can also generate an artefactual confirmation bias, even though this is to a lesser extent. Notably, in two out of three studies, we explicitly controlled for impulsive perseveration, providing evidence supporting the confirmation bias [15, 16].

Figure 1: box-and-arrow representation of the ‘Full’ computational model. The figure distinguishes between observables (‘public’ facts: choices and outcomes) and latent variables (such as Q-values and prediction errors) that are not publicly observables. The parameters used to compute and transform the latent variables are inferred from the observables via model fitting. When ϕ or τ equal zero the ‘Full’ model reduces to the ‘Confirmation’ model. When $\alpha_{CON}=\alpha_{DIS}$ the ‘Full’ model reduces to the ‘Gradual perseveration’ model.

The fact that an apparent confirmation bias may spuriously arise from an underlying gradual perseveration model (and vice versa) is not surprising. The two models generate largely similar behavioral signatures, such as a preferred response rate in choice contexts with no correct response and a reluctance to reverse the response (whether the perseveration model predicts the concavity of the distribution of correct response rate is unclear). In a recent follow-up study Sugawara and Katahira (thereafter S&K) showed that fitting model that includes both explicit autocorrelation and valence-specific learning rates ('Full' model **Figure 1**) allows to quantify the relative contributions of both confirmation bias and gradual perseveration reasonably well. They also reanalyzed three previously published datasets (including two from one of our own studies; [15]) and ran a new online experiment (using similar learning contexts and contingencies). Here they found that the 'Full' model's estimated learning asymmetries were either reduced or disappeared in three out of five cases.

The goal of the present study is to extend the same analyses to all the datasets at our disposal, while correcting some aspects of S&K's analytical choices that we found debatable. To ensure transparency and replicability, all datasets and code are made available on figshare.com (link below). Finally, in the discussion, we will look at the gradual perseveration processes (and other related computational processes) from a neuro-psychological plausibility perspective.

Results

In total we reanalyzed the 9 experiments originally included in the four papers proposing or challenging the confirmation bias hypothesis [11, 15, 16, 23]. All

experiments featured two-armed bandit tasks, where options are presented in fixed pairs with binary, probabilistic outcomes. Despite these similarities, the experiments greatly differ in many features that are summarized in **Table 1**. Among several methodological differences, the most notable are: length of the experiment (96-640 trials), sample size (20-143 subjects), the presence of complete feedback (3 experiments) or observational trials (3 experiments), response modality and testing set-up (scanner, experimental economics laboratory, online testing).

Table 1: features of the included experiments. Exp: Experiment are labeled by the first letter of the paper's first author, followed by a number that indicates the number of the experiment according to order in which it is presented in the paper. N: sample size of the experiment. T: number of trials in the experiment. States: how many choice contexts (or pair of options). Structure: whether the choices of a given contexts were all presented in a row (blocked) or randomized (interleaved). Complete: the percentage of complete feedback trials. Forced: the percentage of observational trials. Resp.: motor response modality. S-R mapping: the mapping between the stimuli and the responses. Outcomes: values of the highest and lowest possible outcomes. Setting: the setting of the experiments.

Exp.	N	T	States	Structure	Complete	Forced	Resp.	S-R mapping	Outcomes	Site
L1 [11]	50	96	4	Interleaved	0%	0%	two keys	Fixed	0.5€ / 0.0€	Scanner
L2 [11]	35	96	4	Interleaved	0%	0%	two keys	Fixed	0.5€ / -0.5€	Lab(multi)
P1 [15]	20	192	8	Interleaved	0%	0%	two keys	Variable	1pt / -1pt	Lab(single)
P2 [15]	20	192	8	Interleaved	100%	0%	two keys	Variable	1pt / -1pt	Lab(single)
C1 [16]	24	720	12	Blocked	0%	33%	two keys	Variable	1pt / -1pt	Lab(single)
C2 [16]	24	640	16	Blocked	50%	50%	two keys	Variable	1pt / -1pt	Lab(single)
C3 [16]	30	360	12	Blocked	0%	33%	two keys	Variable	1pt / -1pt	Lab(single)

C4 [16]	20	600	6	Blocked	0%	0%	one key (go-nogo)	Variable	1pt / -1pt	Lab(single)
S1 [23]	143	384	16	Interleaved	50%	0%	two keys	Variable	10pt / -10pt	Online

We fitted the ‘Full’ model to all datasets, assuming a different set of free parameters (α_{CON} , α_{DIS} , β , τ and ϕ) per subject, as generally accepted [24]. We also fitted the ‘Confirmation’ model (α_{CON} , α_{DIS} and β) to compare the estimations of the confirmation bias across models. The datasets featuring observational trials were fitted with a model featuring an additional learning rate for the outcomes resulting from forced choices (α_{OBS}). The models were fitted by maximization of the model evidence (see **Methods** for details). Concerning data analysis, the main differences with S&K are the following: 1) Regarding the complete feedback cases, we focused on the ‘Confirmation’ model (and not the model with four learning rates that we never claimed being the winning one); 2) S&K analyzed the new online experiment by splitting each subject’s dataset into two smaller datasets with either partial or complete feedback; here we are fitting a unique set of parameter per subject as it usually done and has been done in previous studies of the confirmation bias.

We will now turn to the analysis of the parameters (**Figure 2**). Unless otherwise specified, we ran statistical tests assuming one point per experiment. For the learning rates, if we focus on the average $(\alpha_{CON} + \alpha_{DIS})/2$, we note the value is significantly lower than the value of the choice trace accumulation rate τ ($t(8)=4.1093$, $p=0.0034$). This observation bears consequences about the plausibility of some interpretation of the gradual perseverance process, which we will be discussed later. Replicating S&K, we found the choice trace decision bias parameter ϕ positive and significantly different

from zero ($t(8)=6.8316$, $p=1.3346e-04$), which confirms that the bias goes in the direction of choice perseverance or hysteresis (rather than an alternation)⁸.

We then turn to the analysis of the confirmation bias metrics, which are the main focus of this study. We focus on two metrics. The first metric we considered is simply difference between the learning rate ($\alpha_{CON} - \alpha_{DIS}$), which is the same metric considered by S&K (**Figure 3A**). Replicating S&K, we found that inclusion of the gradual perseveration process significantly reduced the confirmation bias ($t(8)=-3.8728$, $p=0.0047$). However, we found the bias remains positive in all experiments (and significant at $p<0.05$ in 6 out of 9 experiments) and overall significantly higher than zero ($t(8)=5.5105$, $p=0.0005667$).

We also considered an additional metric, which we proposed to account for a possible difference in the average learning rates and choice temperature⁹: the normalized confirmation bias $(\alpha_{CON} - \alpha_{DIS})/(\alpha_{CON} + \alpha_{DIS})$. The normalized confirmation bias was also reduced by the inclusion of a gradual perseveration process (**Figure 3B**; $t(8)=-5.7798$, $p=0.0004$). However, we found that the bias remains positive in all experiments (and significant at $p<0.05$ in 8 out of 9 experiments) and overall significantly higher than zero ($t(8)=8.3386$, $p=0.0000324$).

⁸ Of note, across experiments the inverse temperature β and the average learning $(\alpha_{CON} + \alpha_{DIS})/2$ rates were negatively correlated ($R=-0.9178$, $p=0.0005$). Similarly, the choice trace accumulation rate τ and decision bias ϕ were also negatively correlated ($R=-0.8507$, $p=0.0036$). No other correlation was significant at $R>0.7$ and $p<0.01$.

⁹ If the choice inverse temperature is very high and the learning rate very small, a small difference between the confirmatory and the disconfirmatory learning rate may have a big effect on choices. The validity of this metric is corroborated by the observation that the inverse temperature and the learning rates are negatively correlated across experiments (of note this is also generally true within the individual experiments).

Figure 2: Model parameters. The figure displays the individual (grey dots) and average $\pm$ s.e.m. (blue dots) per experiment. The last position in the x-axis ('ALL') of each panel reports the averages of each experiment (blue dots), the across experiment average $\pm$ sem (pink dots).

Finally, we also considered the comparison between the choice inverse temperature β and the choice trace decision bias. Together these determine choice probability ϕ in the 'Full' model. Since ϕ can take both positive (choice repetition) and negative (alternance) values, we considered its absolute values. As for the relation between the learning rates, we considered two metrics: the difference and the log ratio (**Figure 3C** and **Figure 4D**). Both metrics indicate that the β was bigger than the ϕ (difference: $t(8)=3.4212$, $p=0.0064$; log-ratio: $t(8)=3.5939$, $p=0.0070$), thus confirming that, overall, subjects were paying more attention to the history of reward compared to the history

of choices. However, this was not true for each experiment taken individually: in a couple of experiments ϕ was either numerically (C4) or significantly (S1) bigger than β , regardless of the metric (difference or log ratio) under consideration.

Figure 3: Learning and decision biases. The figure displays the individual (grey dots) and average $\pm$ s.e.m. (blue dots) per experiment. The last position ('ALL') of each panel reports, the averages per each experiment (blue dots), the across experiment average $\pm$ sem (pink dots).

Discussion

The nine datasets were fitted with a computational model including both valence-specific updates and explicit autocorrelation. We assessed whether the choice-confirmation bias (i.e., higher learning rate for positive obtained and negative forgone

outcomes) was still detectable, once allowing for gradual perseveration. In line with S&K¹⁰ [23], our results indicate that the estimated choice-confirmation bias is consistently and significantly smaller when gradual perseveration is included in the model. However, our results also indicate that the estimated choice-confirmation bias is robust to the inclusion of gradual perseveration. Here supported by the difference between α_{CON} and α_{DIS} always being positive, and significant in most experiments (a similar result was obtained when considering the normalized confirmation bias). This last observation, at least superficially, is in contrast with S&K's study [23], which reported the bias to be qualitatively persists only in the P2 dataset, not in P1 and S1. Concerning P1, we note that in our case, it was also one of the datasets where the choice-confirmation bias was the most reduced by the inclusion of gradual perseveration. The residual difference could be explained by slightly different choices in parameter optimization between our and S&K's study (see **Methods** below). Likewise in our analyses of S1 the choice-confirmation bias is strongly reduced, even if still very significant (unsurprisingly, given the extremely large sample size). In addition to differences in parameter optimization mentioned above for P1, we believe that other important differences in model and dataset specification may explain this apparent discrepancy between ours and S&K's results. First, S&K split their dataset in two, by separately analyzing the partial and the complete feedback contexts. Second, concerning the complete feedback case, S&K fitted a model with four learning rates and not the 'Confirmation' model, which we showed best fitting the P2 and C2 datasets. Since our analysis is based on a more parsimonious model and an increased

¹⁰ S&K have verified that both processes are reasonably well recovered in simulations, at least using protocol specifications similar to Palminteri et al. [15] - and the four learning rate models concerning complete feedback case. We have not formally verified parameter recoverability in all the protocol specifications used in the dataset analyzed here. Of note most of the experiments involve more trials and the 'Full' model less learning rates, which suggest that parameters should be at least equally recoverable also in the other protocols.

number of trials per subjects, we believe that our parameters estimations are closer to the true value (of course, assuming the model is correct). Therefore, after careful consideration, the discrepancies with S&K's results are small and can easily be explained by different methodological choices.

Our results indicate the coexistence of a choice-confirmation bias and a gradual preservation bias: when both processes are allowed to compete to explain instrumental choice, they both seem to contribute to behavioral variance. The hypothesis of different learning rates for positive and negative prediction errors have been largely justified on neuroanatomical grounds [25–27]. The choice-confirmation observed here in bandit tasks, could represent the reinforcement learning counterpart of update and inferential biases that are ubiquitous in learning and decision-making [28, 29]. In simulation studies, the choice-confirmation bias has been justified on normative grounds [13, 30, 31]. We now turn to discussing the possible (neurophysiological and cognitive) interpretations of the gradual perseveration process and their plausibility in the context of the behavioral paradigms under consideration [32].

Firstly, of note is that gradual perseveration model was explicitly inspired by Akaishi et al.'s paper [21], where it was proposed to explain choice history effects in perceptual decision-making tasks without trial-by-trial feedback. In such scenario, choice traces could be conceived as the building up of a low-level sensorimotor facilitation, where the activation threshold of a motor command is lowered by mere repetition¹¹. We do

¹¹ But note that following this interpretation (i.e., sensori-motor facilitation) the phi parameter should not be allowed to take negative values as negative sensorimotor facilitation is not documented to our knowledge.

not question the possibility that sensorimotor facilitation may play a role in the perceptual tasks that lead to the formulation of the gradual perseveration model [33]. We also conceive the possibility that sensorimotor facilitation processes may play a role in some reinforcement learning tasks, where options can be represented in a fixed perceptual and motor space. In a previous study we administered healthy subjects to a reinforcement learning task, where they had to choose between two options materialized as right/left hand button presses (i.e., fixed sensorimotor mapping) for an unusually long number of trials (1200). In accordance with the sensorimotor interpretation, we found that an (impulsive) perseveration model fitted this data well. Moreover, the task included trials where the outcome was masked, thus mimicking situations where feedback is not provided (as in perceptual decision-making tasks). We found that choice bias parameter ϕ was higher for decisions following masked outcome trials, suggesting that the reliance on choice history was increased when outcome information was degraded. However, while we think that the sensorimotor facilitation interpretation of the gradual perseveration model is plausible in tasks with fixed stimulus-response mapping and blocked design [22], we do not believe it could be extended to the behavioral protocols considered here. In fact, all the considered datasets, involved either randomized stimulus-response associations (a stimulus could appear in both sides of the scree) or interleaved design (the apparition of stimuli pairs were counterbalanced). Additionally, one of the experiments even involved a go / no go response mode, where in order to choose one of the two options the participants had to refrain from any action (**Table 1**). That is to say, in these experiments, options are represented in an abstract stimuli space, rather than a fixed perceptual and motor space, thus ruling out sensorimotor facilitation as a sound interpretation for the gradual perseveration process.

377

378 Another alternative could be to consider the gradual perseveration process as a model
379 of habit formation. Indeed, prominent theories in animal learning literature recognize
380 two different phases in instrumental learning: a goal-directed phase (outcome-based)
381 and a habitual phase, where behavioral control is independent from any outcome [34].
382 Indeed, the gradual perseveration model is computationally very similar to a recent
383 model of habit formation, which suppose that choosing an action increases a 'habit
384 strength' latent variable (in a way very similar to how the choice traces are learnt and
385 represented) [35]. Choices are then based on 'decision variable', which is a weighted
386 average between outcome-dependent action values and choice history-dependent
387 habit strengths (a weighting parameter determines the balance between goal-directed
388 versus habitual control, thus playing a computational role similar to that of the
389 parameter ϕ). Although the structural similarity between the two models make the
390 analogy tempting, we do not believe that habit formation is a plausible interpretation
391 for the gradual perseveration process. The main reason being that habitual control is
392 supposed to kick in only after extensive training, over time scales spanning from days
393 to months [36]. The datasets under consideration here derived from behavioral
394 experiments whose duration reached, at most, a couple of hours. Furthermore, a
395 universally recognized property of habits is that they kick in *after* the goal directed
396 phase, because the habitual traces are extremely slow to build and unbuild (hence
397 several habit-related problems). Our analyses show that the choice trace accumulation
398 rate (i.e., the equivalent of the learning rate for the choice traces) was significantly
399 higher than the outcome-learning rates, which is inconsistent with the 'habitual'
400 interpretation of the gradual perseveration model.

401

A third possible interpretation of the gradual perseveration process could be to consider it a computational instantiation of a (putative) psychological phenomenon traditionally referred to as ‘choice-based preference change’. The choice-based preference change hypothesis is rooted in cognitive dissonance theory [37]. The idea would be that the mere fact of expressing a preference for an item A (over an alternative B) increases the subsequent evaluation of A and decreases the subsequent valuation of B. The effect was traditionally demonstrated using a paradigm (the ‘free-choice’ paradigm) where subjective values for various (not abstract) items are first elicited in isolation [38]. Later subjects are asked to express preference between pairs of randomly selected items. In the final phase of the experiment, the items are re-evaluated individually. Of particular interest are pairs of items that present the same value on the first phase, for which the choice-based preference change theory predicts an increased/decreased valuation of the previously chosen/forgone item to alleviate the cognitive dissonance of having rejected an otherwise good option. Could the gradual perseveration model be an expression of choice-based preference change operating at the level of reinforcement learning? Answering this question requires first establishing whether choice-based preference change is a robust phenomenon. In this respect it is important to mention that the original version of the free-choice paradigm has been shown to provide inflated results [39]. An amended version of the paradigm has later been proposed and a meta-analysis suggests that the choice-based preference change is still detectable, even if the effect size is very small [38, 40, 41]. This is in striking contrast with the values of the gradual perseveration parameters observed here. First, as mentioned before, the choice accumulation rate is quite high (close to 0.6 in average). Second, choice trace decision bias ϕ , even if it is generally smaller (although bigger in two experiments) than the

choice temperature β , is not negligible in size. To evaluate the plausibility of the choice-based preference change hypothesis, it is important to remember that it is only documented for options whose initial subjective valuations are very close (virtually the same): a situation that has not been systematically implemented in the reinforcement learning tasks under consideration. In fact, only 5 out of 9 experiments contained 'same value' options (often labelled as 'symmetric' conditions), ranging from 50% (L1, L2) to 25% (P1, P2, and S1) of the trials¹². Another important aspect must be considered. In the free-choice paradigms, participants are expressing *pre-existing* values and preferences for fully described items, while in reinforcement learning tasks they are *forming* values and preferences for otherwise perfectly neutral and unknown stimuli. One can even say that they are testing hypotheses about the option values, which is a very different intentional stance compared to evaluating well known and well described items. Accordingly, while we can conceive that one may want to find a 'reason' for rejecting Pepsi over Coke, or Bordeaux over Chianti, what cognitive dissonance could arise from picking one (never seen before) specific agathodaimon letter (or identicon symbol)? We believe not much. Finally, it is also important to note that the computational similarity between the gradual perseveration process and the cognitive dissonance-minimizing choice-based preference change is only superficial. The latter supposes that, if a subject is forced to make a choice between two options (whose subjective values are initially the same), they will increase/decrease the value of the chosen/forgone option to eliminate the cognitive dissonance of a poorly motivated choice. On the other hand, the gradual perseveration process postulates that the choice trace of a given option will be revised upward and downward as a

¹² Of note, in our studies, we fitted the confirmatory and the disconfirmatory learning rates separately per condition and found the results were largely not affected by the outcome contingencies, contrary to what choice-based preference change hypothesis predicts [15].

function of the actual choice history (which is quite erratic in the paradigms under consideration, whose correct and preferred choice rates rarely exceed 80% at the end of the learning curves). Is it therefore reasonable to assume that the ups and downs of the choice traces reflect cognitive dissonance-minimization, for say, rejecting B at trial=1 and rejecting A at trial=2? We believe not. Based on these considerations, we do not believe it plausible, to consider the gradual perseveration process as a manifestation of cognitive dissonance-minimizing choice-based preference change in the reinforcement learning tasks of question.

To sum up, assuming the parameter recovery analysis run by S&K extended to the other behavioral protocols considered here, and therefore the learning rates and the perseveration parameters can be concomitantly (and correctly) estimated by fitting the 'Full' model, our results suggest that choice hysteresis are not uniquely due to the confirmation bias. However, we struggle to find a plausible cognitive (and physiological) interpretation for the gradual perseveration model. What if the gradual perseveration process was capturing the contribution of other, unspecified, computational processes? If yes, what could be plausible candidates? In the following section of the discussion, we identify several computational processes that, in addition to the confirmation bias, may affect choice autocorrelation and could have been captured by the gradual perseveration process in the analysis implemented here.

First, it could be hypothesized that, when the outcome of the unchosen option is not presented, subjects also update the unchosen option, but in the opposite direction. Assuming that subjects will naturally converge to choosing the more rewarding option, this fictive learning rule will induce a magnification of the difference between the Q-

values between the best and worst option. Consequently, subjects will display more choice hysteresis than predicted by the confirmation bias alone. This fictive learning heuristic is rational when one option is systematically better than other one: a situation that was verified in many of the experiments considered here (but not all conditions, such as the ‘symmetric’ ones). Of note, behavioral studies showed that subjects tend to assume an anticorrelation between obtained and forgone outcomes, even when by both design and instructions are explicitly told not to do so [42, 43]. We believe that this fictive update heuristic is plausible and may play a role, even if it suffers from a strong limitation: it only applies to partial feedback conditions (3 out of the 9 experiments under consideration involved complete feedback).

Second, the models considered here assume static learning rates: the values of α_{CON} and α_{DIS} remain the same during the whole experimental session. Static learning rates are not perfectly suited for tasks involving stationary reward probabilities. Here it is optimal to progressively decrease them, in order to ‘protect’ the learner from non-informative negative feedback (once they have identified the correct response) [44, 45]. Progressively reducing learning rates, will also reliably increase choice hysteresis, additional to the confirmation bias. There is ample evidence that learning rates are adaptively modulated in human (and non-human) reinforcement learning [44, 45]. Static reward contingencies were featured in all experiments and most learning contexts¹³. Even if the learning contexts were quite short (24 trials per condition, which is much shorter than the standard stable/volatile task), it is nonetheless plausible to

¹³ Only three experiments (P1, P2 and S1) featured a reversal in reward contingencies, in a small proportion (25%) of trials, which could make unsuited for dynamic reduction of the learning rates.

assume that a progressive reduction of learning rates may play a role in the experiments under consideration, especially at the end of the learning process [46].

Third, until very recently, the tacit assumption in reinforcement learning literature was that the selection process (i.e., the softmax) was the only source of stochasticity or choice variability. In a recent study, based on compelling results in perceptual decision making [47], we tested whether the update process (delta or Rescorla-Wagner rule) also presented a stochastic component [48]. Over four behavioral experiments we demonstrated that the update process is stochastic as well. Specifically, combining model simulations, fitting and neuroimaging, we showed that a gaussian noise (with mean equals to zero and variance proportional to the size of the prediction error) is systematically added at each update step. In the original study we also demonstrated that noisy update induces an increase in choice hysteresis. The reason being that, at any given trial, chances are the subject is choosing an option whose values have been locally inflated by (positive) noise. Therefore, the chances of repeating a choice are generally higher than what one would estimate from uncorrupted (i.e., noisy-free) action values. The plausibility of the noisy update hardly demands to be justified (if anything, is infinite precision update that is implausible). We therefore strongly believe that noisy updates contribute to increase choice hysteresis additional to the confirmation bias¹⁴.

¹⁴ At this stage, one may wonder whether noisy update explains away the confirmation bias. We (with Vasilisa Skvortsova and Valentin Wyart) have tested this hypothesis in several datasets (including those presented in the original publication) and found that it is not the case. Both biases co-exists and explain unique part of behavioral variance. Of note, these analyses also show that the confirmation bias is detected in tasks including continuous rewards and is not restricted to task with binary rewards.

Our results suggest that a confirmation bias and gradual perseveration are both detectable and significant in considered datasets. This suggests that in our data there are more choice hysteresis (or autocorrelation) than predicted by only a choice-confirmation bias in the option value update. We proceeded to challenge several physiological and psychological interpretations of the gradual perseveration model, which is, essentially, a form of choice self-reinforcement: in the form of sensorimotor facilitation, habit formation and, finally, choice-based preference change. We reviewed alternative (learning-based) computational processes that may contribute to increase choice hysteresis additional to of the confirmation bias (fictive update of the unchosen option, decreasing learning rates and noisy update). These we consider to be more plausible than the gradual perseveration process for the paradigms in question. We also believe that these computational processes present at least three other advantages compared to the gradual perseveration process. First, they do not suffer for scaling issues. Second, they have a higher epistemic value. Third they are more representationally parsimonious. Concerning the first point, in all formulations of the choice trace model (including its habit version [20, 21, 35]), the chosen option receives a +1 bonus. In the considered analysis (and in most of the experiments) the outcomes were recoded -1 and +1 (outcome magnitude was never manipulated), but what if the tasks involved rewards in much higher or much smaller reward range? Transposing the current implementation of the model, one would predict that the process would become irrelevant (big outcome ranges) or totally dominate the decision-making process. Even if this issue could be circumvented by assuming some form option value normalization [49, 50], it is worth noting that fictive updates, decreasing learning rates and noisy updates do not suffer from this limitation. Concerning the second point, in the respect to the question, “*what are the causes of choice hysteresis?*”, the gradual

perseveration model provides a rather frustrating answer. It boils down to: “*humans tend to repeat choices because they tend to repeat choices*”. In other words, this does not explain the phenomenon because it is a tautology. On the other hand, fictive updates, decreasing learning rates and noisy updates, all provide an *external* explanation for choice hysteresis. Concerning the third point, it should be remembered that model parsimony should be evaluated beyond the simple number of free parameters¹⁵. The gradual perseveration model maintains two memory or representational systems, while all the other processes under consideration only require one memory system¹⁶.

To conclude, in a series of studies we hypothesized that in human reinforcement learning (at least for simple bandit tasks) choices are consistent with a choice - confirmation bias. The destiny of clearly specified and testable (cognitive) hypotheses is to be challenged. This is how scientific process unfolds to ultimately provide better, i.e., more accurate, theories [23]. Following this golden rule S&K have argued that the confirmation bias results may be artefactual and choices may be better explained by a gradual perseveration process [23]. Here we challenged this claim, by empirically showing that the confirmation bias resists the inclusion of gradual perseveration process and by challenging its psychological plausibility (at least in the paradigms

¹⁵ Whose number can always be reduced by fixing them to the average value or as fixed function of another free parameter: the strong correlations between them would justify this approach.

¹⁶ On a more subjective note, during our career, we personally administered these (or similar) tasks to hundreds of participants. In between experiments, in the ‘piloting’ phase, we have ourselves performed different versions of these tasks’ uncountable times. The fact is, that the gradual perseveration process does not square with our subjective experience of performing (as experimental subject) and administering (as experimenter) the task. If anything, in our experience the gradual perseveration process (i.e., basing choices on past choices rather than past outcomes) would correspond more to the behavior of those subjects not really engaged in the task. With this in mind, it is perhaps unsurprising that the only experiment where the parameter ϕ was significantly bigger than the parameter β is an online experiment, where attention and motivation are usually lower (S1). Of course, we recognize that subjective experience and introspection should be considered with caution (even it apparently played a major role in the development of prospect theory).

under consideration). We hope the results and arguments presented here (together with the associated datasets and codes) will help this collective, dialectic and self-correcting effort that is the Scientific process. Ultimately, we hope this will contribute to elucidating one the most fundamental, and yet apparently elusive, questions in human reinforcement learning research.

Methods

Datasets

We re-analyzed nine datasets. The inclusion criterion for dataset inclusion is being part of a study demonstrating [11, 15, 16] or challenging the confirmation bias hypothesis [23]. We refer to the original publications for the details concerning the demographic data. All experiments featured two-armed bandit tasks with binary and probabilistic outcomes. The differences between datasets are reported in **Table 1**. Before analysis, the datasets were organized in a similar format, which have been made available on figshare.com. To ensure between-experiment commensurability of the ϕ and the β parameters and avoid the necessity of experiment-specific parameters priors, the best and the worst outcome of each study were rescaled between -1 and +1.

Model fitting

We fitted two models separately for all subjects. The ‘Full’ model is described in the text and contains both valence-specific update and gradual perseverance process. To compare the estimation of the learning rates with and without gradual perseveration

we also fitted ‘Confirmation’ model which only contains valence-specific update, but not the gradual perseveration process. Three experiments (C1, C2 and C3) contained ‘forced’ choice trials (also referred to as ‘observational’ trials). The models fitted to these datasets therefore included an additional learning rate dedicated to outcomes following observational trials (α_{OBS}).

We fitted each learning model to individual subject behavior by finding the parameters that minimize the negative log posterior probability (LPP) of the choices [51]:

$$LPP = LL + \log (P(\hat{\theta}_M|M))$$

Where LL is the log likelihood of the data (product of the probability of each choice) given the parameters, and the second term is the sum of log prior probabilities for each parameter. We used general broad priors from the literature [52]: Gamma(1.2,5) for the β , and Beta(1.1,1.1) for the learning rates and the choice accumulation rate τ and Normal(0,1)¹⁷ for the choice trace decision bias φ (see **Table 2** for all parameters priors and ranges).

We used a standard algorithm for function optimization, as implemented in Matlab’s fmincon function. We increased the maximum number of iterations and function evaluations to 10000 and we repeated the optimization procedure 5 times per each model and subjects by randomly picking starting points to avoid local minima. All codes are available here: github.com/spalminteri/conf-bias-meta-analysis.

¹⁷ S&K [23] modified the prior of the φ parameter to Normal(0,5). This may have artificially inflated the estimated size of the φ parameter and may therefore contribute explaining the small discrepancies with our results.

609

610 **Table 2:** computational model parameters' priors and ranges.

Parameter	Prior	Range
α_{CON}	Beta(1.1,1.1)	$0 \leq x \leq 1$
α_{DIS}	Beta(1.1,1.1)	$0 \leq x \leq 1$
α_{OBS}	Beta(1.1,1.1)	$0 \leq x \leq 1$
τ	Beta(1.1,1.1)	$0 \leq x \leq 1$
φ	Normal(0,1)	$-\infty \leq x \leq \infty$
β	Gamma(1.2,5)	$0 \leq x \leq \infty$

611

612 **Statistical analyses**

613 The subject-level parameters were used to calculate experiment-level averages and
614 their transformation, such as the average learning rate $(\alpha_{CON} + \alpha_{DIS})/2$, the confirmation
615 bias $(\alpha_{CON} - \alpha_{DIS})$ and the normalized bias $(\alpha_{CON} - \alpha_{DIS})/(\alpha_{CON} + \alpha_{DIS})$. Parameters and
616 their transformation were then compared among them or against zero using one-
617 sample two-tailed t-test.

618

619 **Data availability**620 figshare.com/articles/software/Confirmation_bias_gradual_perseveration/14877870

Acknowledgements

We thank Michiyo Sugawara for providing clarifications about S1 dataset and analyses. We thank Maël Lebreton for sharing the ‘multiple starting points’ optimization code. We thank Zoe Koopmans for checking the English and Fabien Cerrotti for taking care of the references. We thank Basile Garcia and Nahuel Salem-Garcia for checking the codes. We thank Mael Lebreton, Valerian Chambon, Valentin Wyart and Henri Vandendriessche for useful comments on the manuscript. SP is supported by an ATIP-Avenir grant (R16069JS), the Programme Emergence(s) de la Ville de Paris, the Fondation Fyssen, the Fondation Schlumberger pour l’Education et la Recherche and the Institut de Recherche en Santé Publique (IRESP, grant number : 20II138-00). The article was prepared in the framework of a research grant funded by the Ministry of Science and Higher Education of the Russian Federation (grant ID: 075-15-2020-928) and the French National Agency of Research (ANR; FrontCog ANR-17-EURE-0017).

References

- [1] Sharot, T.; Garrett, N. (2016). Forming Beliefs: Why Valence Matters, *Trends in Cognitive Sciences*, Vol. 20, No. 1, 25–33. doi:10.1016/j.tics.2015.11.002
- [2] Sharot, T.; Korn, C. W.; Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality, *Nature Neuroscience*, Vol. 14, No. 11, 1475–1479. doi:10.1038/nn.2949
- [3] Palminteri, S.; Pessiglione, M. (2017). Opponent brain systems for reward and punishment learning: causal evidence from drug and lesion studies in humans, *Decision Neuroscience*, Elsevier, 291–303
- [4] Niv, Y.; Edlund, J. A.; Dayan, P.; Doherty, J. P. O. (2012). Neural Prediction Errors Reveal a Risk-Sensitive Reinforcement-Learning Process in the Human Brain, *The Journal of Neuroscience*, Vol. 32, No. 2, 551–562. doi:10.1523/JNEUROSCI.5498-10.2012
- [5] Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards ?, *Psychonomic Bulletin and Review*, Vol. 22, 1320–1327. doi:10.3758/s13423-014-0790-3

- 648 [6] Christakou, A.; Gershman, S. J.; Niv, Y.; Simmons, A.; Brammer, M.; Rubia, K. (2011). Neural
649 and Psychological Maturation of Decision-making in Adolescence and Young Adulthood,
650 *Journal of Cognitive Neuroscience*, Vol. 25, No. 11, 1807–1823. doi:10.1162/jocn
- 651 [7] Ljungman, L.; Boger, M.; Ander, M.; Ljótsson, B.; Cernvall, M.; von Essen, L.; Hovén, E.
652 (2016). Impressions That Last: Particularly Negative and Positive Experiences Reported by
653 Parents Five Years after the End of a Child's Successful Cancer Treatment or Death, *PLOS*
654 *ONE*, Vol. 11, No. 6, e0157076. doi:10.1371/journal.pone.0157076
- 655 [8] Frank, M. J.; Moustafa, A. A.; Haughey, H. M.; Curran, T.; Hutchison, K. E. (2007). Genetic
656 triple dissociation reveals multiple roles for dopamine in reinforcement learning, *Proceedings*
657 *of the National Academy of Sciences of the United States of America*, Vol. 104, No. 41,
658 16311–16316. doi:10.1073/pnas.0706111104
- 659 [9] Van Den Bos, W.; Cohen, M. X.; Kahnt, T.; Crone, E. A. (2012). Striatum-medial prefrontal
660 cortex connectivity predicts developmental changes in reinforcement learning, *Cerebral*
661 *Cortex*, Vol. 22, No. 6, 1247–1255. doi:10.1093/cercor/bhr198
- 662 [10] Worbe, Y.; Palminteri, S.; Hartmann, A.; Vidailhet, M.; Lehericy, S.; Pessiglione, M. (2011).
663 Reinforcement Learning and Gilles de la Tourette Syndrome, *Archives of General Psychiatry*,
664 Vol. 68, No. 12, 1257. doi:10.1001/archgenpsychiatry.2011.137
- 665 [11] Lefebvre, G.; Lebreton, M.; Meyniel, F.; Bourgeois-Gironde, S.; Palminteri, S. (2017).
666 Behavioural and neural characterization of optimistic reinforcement learning, *Nature Human*
667 *Behaviour*, Vol. 1, No. 4, 1–9. doi:10.1038/s41562-017-0067
- 668 [12] Palminteri, S.; Wyart, V.; Koechlin, E. (2017). The Importance of Falsification in Computational
669 Cognitive Modeling, *Trends in Cognitive Sciences*, Vol. 21, No. 6, 425–433.
670 doi:10.1016/j.tics.2017.03.011
- 671 [13] Lefebvre, G.; Summerfield, C.; Bogacz, R. (2020). A normative account of confirmation bias
672 during reinforcement learning, *BioRxiv*, 1–21. doi:10.1101/2020.05.12.090134
- 673 [14] Palminteri, S.; Khamassi, M.; Joffily, M.; Coricelli, G. (2015). Contextual modulation of value
674 signals in reward and punishment learning, *Nature Communications*, Vol. 6, No. 1, 8096.
675 doi:10.1038/ncomms9096
- 676 [15] Palminteri, S.; Lefebvre, G.; Kilford, E. J.; Blakemore, S. J. (2017). Confirmation bias in human
677 reinforcement learning: Evidence from counterfactual feedback processing, *PLoS*

- 678 *Computational Biology*, Vol. 13, No. 8, e1005684. doi:10.1371/journal.pcbi.1005684
- 679 [16] Chambon, V.; Théro, H.; Vidal, M.; Vandendriessche, H.; Haggard, P.; Palminteri, S. (2020).
 680 Information about action outcomes differentially affects learning from self-determined versus
 681 imposed choices, *Nature Human Behaviour*, Vol. 4, No. 10, 1067–1079. doi:10.1038/s41562-
 682 020-0919-5
- 683 [17] Farashahi, S.; Donahue, C. H.; Hayden, B. Y.; Lee, D.; Soltani, A. (2019). Flexible combination
 684 of reward information across primates, *Nature Human Behaviour*, Vol. 3, No. 11, 1215–1224.
 685 doi:10.1038/s41562-019-0714-3
- 686 [18] Gagne, C.; Zika, O.; Dayan, P.; Bishop, S. J. (2020). Impaired adaptation of learning to
 687 contingency volatility in internalizing psychopathology, *ELife*, Vol. 9, 1–51.
 688 doi:10.7554/ELIFE.61387
- 689 [19] Steinke, A.; Lange, F.; Kopp, B. (2020). Parallel model-based and model-free reinforcement
 690 learning for card sorting performance, *Scientific Reports*, Vol. 10, No. 1, 1–18.
 691 doi:10.1038/s41598-020-72407-7
- 692 [20] Katahira, K.; Yuki, S.; Okanoya, K. (2017). Model-based estimation of subjective values using
 693 choice tasks with probabilistic feedback, *Journal of Mathematical Psychology*, Vol. 79, 29–43.
 694 doi:10.1016/j.jmp.2017.05.005
- 695 [21] Akaishi, R.; Umeda, K.; Nagase, A.; Sakai, K. (2014). Autonomous Mechanism of Internal
 696 Choice Estimate Underlies Decision Inertia, *Neuron*, Vol. 81, No. 1, 195–206.
 697 doi:10.1016/j.neuron.2013.10.018
- 698 [22] Gershman, S. J.; Pesaran, B.; Daw, N. D. (2009). Human reinforcement learning subdivides
 699 structured action spaces by learning effector-specific values, *Journal of Neuroscience*, Vol. 29,
 700 No. 43, 13524–13531. doi:10.1523/JNEUROSCI.2469-09.2009
- 701 [23] Sugawara, M.; Katahira, K. (2021). Dissociation between asymmetric value updating and
 702 perseverance in human reinforcement learning, *Scientific Reports*, Vol. 11, No. 1, 3574.
 703 doi:10.1038/s41598-020-80593-7
- 704 [24] Wilson, R. C.; Collins, A. G. (2019). Ten simple rules for the computational modeling of
 705 behavioral data, *ELife*, Vol. 8. doi:10.7554/eLife.49547
- 706 [25] Collins, A. G. E.; Frank, M. J. (2012). How much of reinforcement learning is working memory,
 707 not reinforcement learning? A behavioral, computational, and neurogenetic analysis, *European*

- Journal of Neuroscience*, Vol. 35, No. 7, 1024–1035. doi:10.1111/j.1460-9568.2011.07980.x
- [26] Palminteri, S.; Pessiglione, M. (2017). Chapter 23 - Opponent Brain Systems for Reward and Punishment Learning: Causal Evidence From Drug and Lesion Studies in Humans, J.-C. Dreher; L. Tremblay (Eds.), *Decision Neuroscience*, Academic Press, San Diego, 291–303. doi:https://doi.org/10.1016/B978-0-12-805308-9.00023-3
- [27] Dabney, W.; Kurth-Nelson, Z.; Uchida, N.; Starkweather, C. K.; Hassabis, D.; Munos, R.; Botvinick, M. (2020). A distributional code for value in dopamine-based reinforcement learning, *Nature*, Vol. 577, No. 7792, 671–675. doi:10.1038/s41586-019-1924-6
- [28] Klayman, J. (1995). Varieties of Confirmation Bias, J. Busemeyer; R. Hastie; D. L. Medin (Eds.), (Vol. 32), Academic Press, 385–418. doi:https://doi.org/10.1016/S0079-7421(08)60315-1
- [29] Raymond S. Nickerson. (1998). Confirmation bias: a ubiquitous phenomenon in many guises, *Review of General Psychology*, Vol. 2, No. 2, 175–220
- [30] Paleari, M.; Lisetti, C.; Grizard, A. (2007). *Adapting Psychologically Grounded Facial Emotional Expressions to Different Anthropomorphic Embodiment Platforms*
- [31] Tarantola, T.; Folke, T.; Boldt, A.; Pérez, O. D.; Martino, B. De. (2021). Confirmation bias optimizes reward learning, *BioRxiv*, 2021.02.27.433214
- [32] van Rooij, I.; Baggio, G. (2021). Theory Before the Test: How to Build High-Verisimilitude Explanatory Theories in Psychological Science, *Perspectives on Psychological Science*. doi:10.1177/1745691620970604
- [33] Gold, J. I.; Law, C. T.; Connolly, P.; Bennur, S. (2008). The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning, *Journal of Neurophysiology*, Vol. 100, No. 5, 2653–2668. doi:10.1152/jn.90629.2008
- [34] POLDRACK, R. A. (2021). *Hard to Break*, Princeton University Press. doi:10.2307/j.ctv191kx44
- [35] Miller, K. J.; Shenhav, A.; Ludvig, E. A. (2019). Habits without values., *Psychological Review*, Vol. 126, No. 2, 292–311. doi:10.1037/rev0000120
- [36] Lally, P.; van Jaarsveld, C. H. M.; Potts, H. W. W.; Wardle, J. (2010). How are habits formed: Modelling habit formation in the real world, *European Journal of Social Psychology*, Vol. 40, No. 6, 998–1009. doi:10.1002/ejsp.674

- 738 [37] Festinger, L. (1962). Cognitive Dissonance, *Scientific American*, Vol. 207, No. 4, 93–106.
739 doi:10.1038/scientificamerican1062-93
- 740 [38] Izuma, K.; Murayama, K. (2013). Choice-Induced Preference Change in the Free-Choice
741 Paradigm: A Critical Methodological Review, *Frontiers in Psychology*, Vol. 4, No. April 2014.
742 doi:10.3389/fpsyg.2013.00041
- 743 [39] Chen, M. K.; Risen, J. L. (2010). How Choice Affects and Reflects Preferences: Revisiting the
744 Free-Choice Paradigm, *Journal of Personality and Social Psychology*, Vol. 99, No. 4, 573–594.
745 doi:10.1037/a0020217
- 746 [40] Nakao, T.; Kanayama, N.; Katahira, K.; Odani, M.; Ito, Y.; Hirata, Y.; Nasuno, R.; Ozaki, H.;
747 Hiramoto, R.; Miyatani, M.; Northoff, G. (2016). Post-response $\beta\gamma$ power predicts the degree of
748 choice-based learning in internally guided decision-making, *Scientific Reports*, Vol. 6, No.
749 October. doi:10.1038/srep32477
- 750 [41] Izuma, K.; Matsumoto, M.; Murayama, K.; Samejima, K.; Sadato, N.; Matsumoto, K. (2010).
751 Neural correlates of cognitive dissonance and choice-induced preference change,
752 *Proceedings of the National Academy of Sciences*, Vol. 107, No. 51, 22014–22019.
753 doi:10.1073/pnas.1011879108
- 754 [42] Marciano-Romm, D.; Romm, A.; Bourgeois-Gironde, S.; Deouell, L. Y. (2016). The Alternative
755 Omen Effect: Illusory negative correlation between the outcomes of choice options, *Cognition*,
756 Vol. 146, 324–338. doi:10.1016/j.cognition.2015.10.003
- 757 [43] Marciano, D.; Bentin, S.; Deouell, L. Y. (2018). Alternative outcomes create biased
758 expectations regarding the received outcome: Evidence from event-related potentials,
759 *Neuropsychologia*, Vol. 113, 126–139. doi:10.1016/j.neuropsychologia.2018.03.010
- 760 [44] McGuire, J. T.; Nassar, M. R.; Gold, J. I.; Kable, J. W. (2014). Functionally Dissociable
761 Influences on Learning Rate in a Dynamic Environment, *Neuron*, Vol. 84, No. 4, 870–881.
762 doi:10.1016/j.neuron.2014.10.013
- 763 [45] Behrens, T. E. J.; Woolrich, M. W.; Walton, M. E.; Rushworth, M. F. S. (2007). Learning the
764 value of information in an uncertain world, *Nature Neuroscience*, Vol. 10, No. 9, 1214–1221.
765 doi:10.1038/nn1954
- 766 [46] Harada, T. (2020). Learning From Success or Failure? – Positivity Biases Revisited, *Frontiers*
767 *in Psychology*, Vol. 11, No. July, 1–11. doi:10.3389/fpsyg.2020.01627

- 768 [47] Drugowitsch, J.; Wyart, V.; Devauchelle, A.-D.; Koechlin, E. (2016). Computational Precision
 769 of Mental Inference as Critical Source of Human Choice Suboptimality, *Neuron*, Vol. 92, No. 6,
 770 1398–1411. doi:10.1016/j.neuron.2016.11.005
- 771 [48] Findling, C.; Skvortsova, V.; Dromnelle, R.; Palminteri, S.; Wyart, V. (2019). Computational
 772 noise in reward-guided learning drives behavioral variability in volatile environments, *Nature*
 773 *Neuroscience*, Vol. 22, No. 12, 2066–2077. doi:10.1038/s41593-019-0518-9
- 774 [49] Bavard, S.; Lebreton, M.; Khamassi, M.; Coricelli, G.; Palminteri, S. (2018). Reference-point
 775 centering and range-adaptation enhance human reinforcement learning at the cost of irrational
 776 preferences, *Nature Communications*, Vol. 9, No. 1, 4503. doi:10.1038/s41467-018-06781-2
- 777 [50] Bavard, S.; Rustichini, A.; Palminteri, S. (2021). Two sides of the same coin: Beneficial and
 778 detrimental consequences of range adaptation in human reinforcement learning, *Science*
 779 *Advances*, Vol. 7, No. 14. doi:10.1126/sciadv.abe0340
- 780 [51] Daw, N. D. (2011). Trial-by-trial data analysis using computational models, *Decision Making,*
 781 *Affect, and Learning*, Oxford University Press, 3–38.
 782 doi:10.1093/acprof:oso/9780199600434.003.0001
- 783 [52] Daw, N. D.; Gershman, S. J.; Seymour, B.; Dayan, P.; Dolan, R. J. (2011). Model-based
 784 influences on humans' choices and striatal prediction errors., *Neuron*, Vol. 69, No. 6, 1204–15.
 785 doi:10.1016/j.neuron.2011.02.027
- 786 [53] Schüller, T.; Fischer, A. G.; Gruendler, T. O. J.; Baldermann, J. C.; Huys, D.; Ullsperger, M.;
 787 Kuhn, J. (2020). Decreased transfer of value to action in Tourette syndrome, *Cortex*, Vol. 126,
 788 39–48. doi:10.1016/j.cortex.2019.12.027
- 789