The effect of body image dissatisfaction on goal-directed decision making in a population marked by negative appearance beliefs and disordered eating

Jakub Onysk[1¶], Peggy Seriès[2&*]

[1] Computational and Biological Learning Laboratory, Department of Engineering, University of Cambridge, Cambridge, UK

[2] Institute for Adaptive and Neural Computation, School of Informatics, University of Edinburgh, Edinburgh, UK

* Corresponding author

E-email: pseries@inf.ed.ac.uk (PS)

[¶] Writing (original draft), visualisation, investigation, formal analysis, methodology, code.

[&] Conceptualisation, Supervision, Review and editing, project administration.

Abstract

Eating disorders are associated with one of the highest mortality rates among all mental disorders, yet there is very little research about them within the newly emerging and promising field of computational psychiatry. As such, we focus on investigating a previously unexplored, yet a core aspect of eating disorders – body image dissatisfaction. We continue a freshly opened debate about model-based learning and its trade-off against model-free learning – a proxy for goal-directed and habitual behaviour. We perform a behavioural study that utilises a two-step decision-making task and a reinforcement learning model to understand the effect of body image dissatisfaction on model-based learning in a population characterised by high scores of disordered eating and negative appearance beliefs, as recruited using Prolific. We find a significantly reduced model-based contribution in the body image dissatisfaction task condition in the population of interest as compared to a healthy control.

Introduction

Eating disorders (ED) form a group of different conditions that, in a detrimental way affects a person's relationship with food. This usually leads to physical and psychological problems, which severely decrease the quality of life. In fact, these problems can result in many deaths, as eating disorders are reported to have one of the highest, if not the highest, mortality rate among all mental disorders (1–4).

Generally, ED involve heightened preoccupation with food, such as restriction, or consumption of unusually large amounts of food, which in some cases is followed by compensatory behaviours such as vomiting, use of laxatives, or overexercising (5). In extreme cases, eating disorders result in death either due to severe malnutrition, morbid obesity or suicide (4,6). Moreover, the recovery from ED can be as low as 24% even after 10 years (7), which in combination with high mortality and a significant decrease in quality of life, calls for extended research into the roots and treatments of ED.

In addition to eating related behaviours, one of the core symptoms of ED is body image disturbance, which can be understood as a negative misrepresentation of one's body, body image preoccupation, usually involving disgust, shame and dissatisfaction (8). We here aim to investigate how body image dissatisfaction in eating disorders affects basic mechanisms of decision making, namely, habitual (repeating of the same action as a response to a stimulus) and goal-directed (intentional and deliberate decision-making) behaviour (9). On top of logistic regression analysis of the behavioural data (10), the mechanism will be captured in a computational, reinforcement learning (RL) framework.

Background

Eating disorders.

There are three main subtypes of eating disorders – Anorexia Nervosa (AN), Bulimia Nervosa (BN) and Binge-Eating Disorder (BED) as in 5th edition of The Diagnostic and Statistical Manual of Mental Disorders (5). All share low life satisfaction, greatly impaired quality of life, and increased mortality or suicide risk, and sometimes body image issues. A brief distinction between the types, but nowhere close to accurately describing the nuance and lived-experience of each subtype, would characterise: i) AN by extreme restriction of food and pursuit of weight loss; ii) BN by undergoing recurrent binge eating episode followed by compensatory behaviours; and iii) BED by the recurrence of uncontrollable binges, but usually without compensatory behaviours (5). In addition, there exists an array of atypical eating disorders that do not fit the description of the main (clinical) types (11), but nevertheless constitutes a majority of cases (12). In this paper, we do not focus on any particular type, but rather aim to explore and quantify the effect of body image dissatisfaction on decision making in broadly understood eating disorders (see Methods).

Computational psychiatry.

One possible avenue to understand and help devise treatments for eating disorders, alongside the traditional approach of psychotherapy, psychiatry and neuroimaging (13–15) is the emerging field of computational psychiatry (CP).This framework is based in the assumption that the brain's characteristic function is one of computation and information processing. As such, it offers an understanding of mental illness whereby the differences and/or errors within these computations may result in malfunctioning, maladaptive behaviours and mental states(16). In an attempt to describe these mechanisms, some researchers have focused on multiple modes of

decision making that can be gauged with a decision-making task and quantified with a computational model that captures individual and group differences (17,18). Mental disorders such as schizophrenia, anxiety, and obsessive-compulsive disorder (OCD) have been linked to various significant changes in decision-making processes (19–21). Moreover, employing computational methods to understand mental disorders promises bridging the neurobiological and psychiatric levels of descriptions of populations of interest, for example the role of dopaminergic signals in reward learning and it implications in addiction (22). Such link can potentially inform better and more advanced theories of mental disorders, as well as inspire new treatment approaches, and promises improved early detection leading to prevention (23).

Computational psychiatry of eating disorders.

As mentioned above, depression, OCD, anxiety and schizophrenia have received a lot attention in the field of CP, with very promising results and theories. Eating disorders have received less attention, with only a couple papers in the last nine years (24–34). To illustrate the difference a Google Scholar search for eating- disorder related papers:

("computational psychiatry" "OR" "computational" "OR" "reinforcement learning" "OR" "reinforcement" "OR" "bayesian" "OR" "decision making" "OR" "decision-making") "AND" ("anorexia" "OR" "anorexia nervosa" "OR" “bulimia nervosa" "OR" "binge" "OR" "binge eating” OR" "eating" "OR" "bulimia") (35),

returns 11 results, while a search for schizophrenia related papers:

("computational psychiatry" "OR" "computational" "OR" "reinforcement learning" "OR" "reinforcement" "OR" "bayesian" "OR" "decision making" "OR" "decision-making") "AND" ("schizophrenia") (36),

returns 466 results.

The main focus has been on two main aspects of decision making in eating disorders. One is related to how sensitive to punishments individuals with AN are, that is how well they learn from negative feedback from the environment (in this case a decision-making task). Unfortunately, the results prove to be quite contradictory (24,33).

The second main aspect of ED that has been of interest to computational psychiatrists is the trade-off between goal-directed and habitual system utilisation in decision-making. As mentioned in the introduction, goal-directed decision-making is related to acting in the environment with a goal in mind, intentionally and deliberately. It is usually characterised by forming an internal model of the environment that describes which states and actions will bring about the best results over prolonged time. On the other hand, habitual decision-making is associated with responding to stimuli in the environment in an automatic manner, without much deliberation, usually repeating those actions that immediately yield the best results. In this case, an agent does not create a model of the environment. As such, goal-directed behaviour utilises model-based learning, which allows building an accurate model of the states and actions and their associated values that takes into account a hidden probabilistic structure of the environment. However, this kind of learning is computationally more demanding, using more resources to support the process. The habitual behaviour employs model-free learning, which updates a running score of possible states and actions, based on the last experience, without registering any hidden structures in the environment. Such a process is computationally efficient as it relies mainly on the memory of the last events (9).

It has been shown that healthy participants employ both model-free and model-based learning that trade off against each other (10). Several studies attempted to see how this trade-off is different within eating disorders. The results are converging and suggest that reduced model

based learning is associated with high eating disorder questionnaire scores in a general population (30). The reduction is also evident clinical AN and BED groups (27,34). Moreover, a recent study by Foerde et al. (27) suggests even further reduced model-based learning contribution in AN when the task is strictly food related as compared to a monetary task. This is of particular significance for two reasons. Firstly, AN is characterised by extreme pursuit of weight-loss, which could intuitively be understood as extreme goal-directed behaviour. As such, one would expect to see increased model-based learning in AN, yet the results suggest otherwise – such "goal-pursuit" of weight-loss in AN is actually habitual, almost compulsive. Secondly, the study includes an additional food-related task, which aims to test whether the goal-directed deficits observed during the monetary task were due to monetary rewards not being perceived as motivating enough for the AN participants to employ model-based strategy (domain-specific). The results from this study show that such deficits are in fact domain-general, where model-based learning is even further reduced during the food-related task.

Problem Statement, Objectives and Hypothesis

Evidently, there is a need for more computational psychiatry research about eating disorders, given the high mortality. A particular aspect that has not been previously investigated is that of body image disturbance.

As a first step in this direction, we aim to explore the effect of body image dissatisfaction on decision making in population marked by negative appearance beliefs and disordered eating by implementing a two-step decision-making task that captures both model-based and model-free contribution (10). The task is given to a population that scores high on an eating disorder, and body image disturbance questionnaires. For comparison, a healthy control group (HC) is also

be recruited. We employ two task conditions, one of which attempts to target the dissatisfaction with body image and its manifestation in decision making. The introduction of a body image dissatisfaction condition aims to create a context similar to that in which ED finds themselves during everyday struggle - excessive worry about body image and how to change it through dieting. Secondly, the condition mirrors the food-related task of Foerde et al. (27), aiming to support the domain-general nature of the goal-directed deficit.

We hypothesise that: 1) In the monetary (neutral) condition and as described by previous studies, the group with ED will report significantly decreased model-based learning as compared to HC, and no difference in model-free learning as previously shown (27,30); 2) Model-based learning will be further decreased in the body image dissatisfaction condition in the group with ED. HC will not report any significant difference between conditions.

Materials and Methods

Participants

An online study was performed. Participants for the study were recruited using Prolific, which is an online participant recruitment service used mainly for research and academic purposes (37). The platform provides a pool of participants that are reliable and allows to custom screen for different groups before recruitment.

Since two groups were needed – a healthy control (HC) and an eating disorder (ED) group - we applied two separate sets of pre-screening criteria on Prolific (pre-set filters that are available in the study designer). For HC we looked for people who:

a) “Have never gone on a diet in the past.”
b) “Do no currently for at least one week restrict food intake to manage weight.”

c) "Have no diagnosed mental health condition that is uncontrolled (by medication or intervention) and which has a significant impact on your daily life/activities."

d) "Do not have or have not had a diagnosed, on-going mental health illness or condition."

For ED, participants had to select that they have gone on a diet, as in criterion (a), and that they restrict food intake to either lose or maintain weight, as in criterion (b). Moreover, to ensure data of satisfactory quality both groups had additional criteria to meet:

(i) age between 18-38

(ii) have normal or corrected-to-normal vision

(iii) female as an assigned sex at birth

(iv) an approval rate of 98%

(v) a minimum of 20 previous submissions on Prolific.

Participants first completed a range of questionnaires (sub-study 1) to be then further selected to complete a decision-making task (sub-study 2). The questionnaires were hosted on Qualtrics (an online survey software) (38), while the decision-making task was hosted partly on Qualtrics and partly on Pavlovia (an online behavioural experiment platform) (39), designed using PsychoPy3 v. 2020.1.2 software (40). For each sub-study, participants were paid at a rate of £6.25 per hour. The questionnaire took on average 7 minutes to complete, while the task took on average 27 minutes. Subjects were based all around the world. The study was approved according to the University of Edinburgh's Informatics Research Ethics Process, with an RT number 2019/48215.

After further selection based on questionnaires and application of exclusion criteria for the decision-making task (see Supplementary material), 35 (ED group) and 32 (HC group) participants were selected for data analysis.

Power analysis

A prior power analysis was performed to determine the sample size needed to detect a significant difference (at the 5% level) between groups in the neutral condition. Following a similarly designed study by Foerde et al. (27), for calculation 1, we aimed for a medium effect size of d=0.75 (two-tailed t-test between two independent means), and the power $(1-\beta)=0.8$. The sample size calculation using G*Power software (41) suggested recruiting 29 participants per group. Moreover, we calculated (calculation 2) the sample size needed to detect the difference between conditions in the ED group (two-tailed t-test between two dependent means). Setting the effect size at d=0.75 and the power at $(1-\beta)=0.8$, resulted in n=16 for the ED group.

Additional post hoc power analysis was performed as above. Using the results from this study with 35 participants in the ED group and 32 in the HC group, calculation 1 revealed effect size of d=0.88 and power $(1-\beta)=0.94$, while calculation 2 revealed effect size of d=1.04 and power $(1-\beta)=0.99$.

Self-report questionnaires

In the first part of the study, participants answered questions about their age, gender, weight, and height. They also completed three questionnaires - Eating Attitudes Test (EAT-26) (42), Appearance Anxiety Inventory (AAI) (43) and The Obsessive–Compulsive Inventory (OCI-

R) (44) – to assess the spectrum of: an eating disorder, body image disturbance and obsessive-compulsive behaviours, respectively.

In order to divide participants into HC and ED groups, cut-off points were applied for the questionnaires as in the Table 1, chosen based on the literature. For EAT-26, it is suggested to use a cut-off of 11 for subclinical populations (45) who display disordered eating behaviours that could warrant further clinical diagnosis. Hence, for HC we set a threshold of ≤ 10 and for ED ≥ 14 to achieve stronger effects. For AAI, a few cut-off points have been suggested: a score above 6 may suggest some issues with body image (46), while scores between 15-20 is reported in an appearance-concerned adult population, and a score above 20 indicates a high-risk group for body image disturbance (47). We therefore chose as cut-off for ED a score ≥ 14, whereas for HC a score of ≤ 10 was chosen. Lastly, an additional criterion for HC group was applied so that they do not display worrying obsessive-compulsive behaviours as these have also been associated with reduced model-based learning (30). We selected a cut-off point of ≤ 10 on OCI-R questionnaire for HC, which is twice as low as the recommended optimal cut-off point (44). The slightly steeper cut-off points than those exactly recommended, should not however affect the results as we still obtain relevant populations of interest. Moreover, two attention checks were implemented in the questionnaires to filter out participants not taking part in the study in good faith.

Table 1. Cut-off points for ED and HC on EAT-26, AAI, and OCI-R questionnaires

	Group	
Measure	**HC**	**ED**
EAT-26	≤10	≥14
AAI	≤10	≥14
OCI-R	≤10	any

Two-step decision making task

Selected participants were asked to complete a two-step decision-making task in order to quantify model-based and model-free learning contribution (10). Before proceeding with the task, participants were asked to select a body type that is most similar to their own from a diverse range of body types in the S1 Fig in the Supporting Information. Furthermore, they were instructed on how to complete the task by reading through an illustrated tutorial. After reading the instructions, they completed 25 trials of the task as a practice.

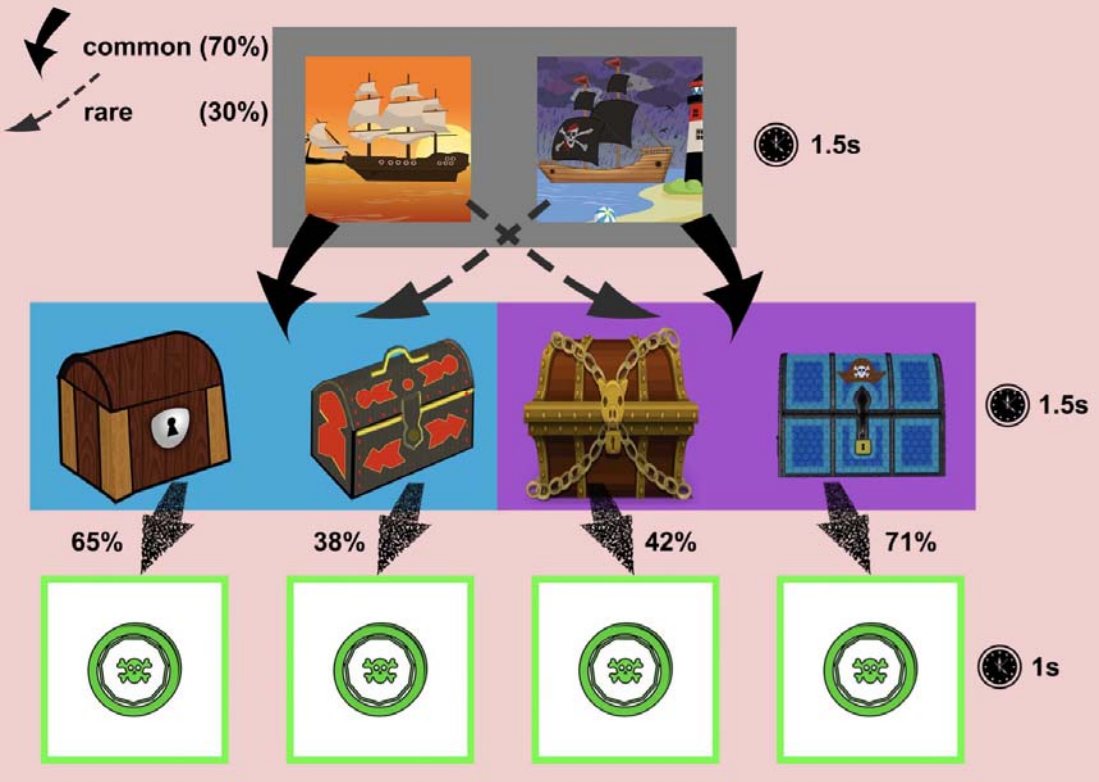

Fig 1. The structure of the two-step decision-making task.

In the first stage (grey), a dashed arrow corresponds to a rare transition with the probability of 30%, while the bold arrow corresponds to a common transition with the probability of 70%. In the second stage, the speckled arrow is associated with the probability of receiving a reward. The duration of each stage is noted to the right (stage one, stage two, reward).

The task reproduces that of Daw et al. (10) but was conceived as a treasure game to make it as engaging and easy to understand as possible (Fig 1). In the first stage (grey stage), participants are presented with a choice between two stimuli - on the left, a ship against a sunset, while on the right, a pirate ship near a lighthouse. They are asked to play a role of a treasure hunter and decide which ship to board. Each ship can sail to two destinations - a blue island and a purple island. The game has a hidden structure, which participants tend to learn with time, such that the normal ship sails to the blue island (left) 70% of the time (common transition) and to the purple island (right) 30% of the time (rare transition), while the pirate ship sails to the blue island 30% of the time and to the purple island 70% of the time.

Once the participant boards the chosen ship, they are taken to one of the islands (stage two), where two chests await them. Here, they are asked to quickly select the chest to try their chances at finding a treasure (a pirate coin). However, each chest is assigned a probability of containing the coin, which evolves over time according to a Gaussian random walk with a standard deviation $\sigma = 0.0275$, as in the Fig 2 (more details in the Supporting Information). One of the strategies is to track, over time, which chest is the most favourable, i.e. has the highest chance of yielding the coin. After the chest is opened and the participant receives the coin (or not), they are taken back to the first stage to repeat the trial.

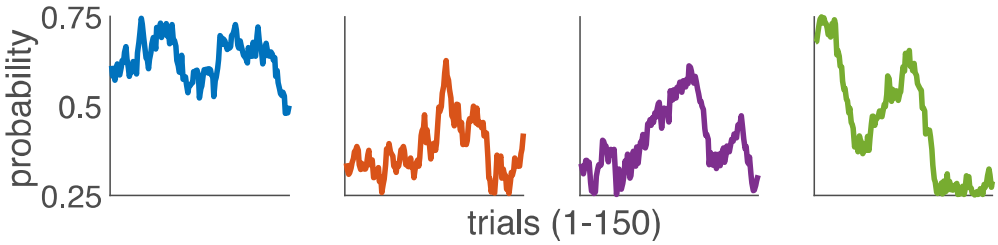

Fig 2. The sample evolution of the reward probability for each of the chest in stage 2. The probabilities (in the order as in Fig 1) evolve over 150 trials, according to a Gaussian random walk with σ=0.0275. One set of two evolutions always starts randomly in a range [0.58,

0.72], while the other set of two starts in [0.31, 0.45]. The values of probabilities are bounded in [0.25, 0.75].

There are two conditions, 150 trials each. Each condition consists of two series of 75 trials, with a short break in between to counteract a possible loss of attention or tiredness. The neutral (NT) condition is exactly like in the Fig 1, with a pirate coin as a reward (Fig 3A). The other condition, the body image dissatisfaction (BID) condition, differs from the neutral in that the reward is a pirate coin next to a body type that the participant selected as the most similar to their own (Fig 3C). The aim of the BID condition is to gauge the effects of body image dissatisfaction on decision-making.

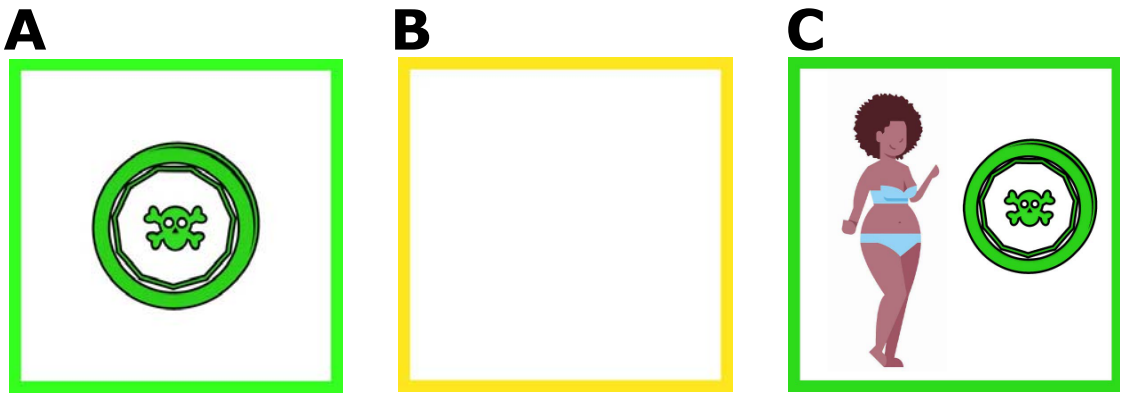

Fig 3. Possible reward outcomes during the task.

In the neutral condition, participants could receive either (A) or (B), while in the BID condition they could receive either (C) or (B). An empty box (B) indicates no reward. The body type was selected by the participant before the task, more body type examples can be found in the S1 Fig.

For a balanced design half of the participants in each group completed the neutral condition first, followed by the BID condition. The other half completed the task in a reverse order. All results are then based on the average across two subgroups.

Model-based and model-free learning in the task.

The task allows to capture the contribution of model-based and model-free learning. Thanks to its stage-like structure and probabilistic nature, we can easily distinguish between two strategies a participant can have, usually a mix of the two.

For the purpose of this example, we focus on two extreme cases (30). On one hand, an agent could utilise only model-free learning strategy. This means that they would make their decision as to which ship and chest to select purely based on the reward they received (or lack thereof), usually repeating the rewarded action, regardless of the fact the transition type. On the other hand, a participant who uses only model-based strategy, takes into account the learnt transition structure of the task, on top of the knowledge about previous rewards. Moreover, they track the probability of receiving the reward to know which chest is the most favourable. In other words, they chose the same ship if the previous trial had a common transition with a reward or a rare transition and no reward.

Mixed-effects logistic regression of raw choice data

First, a mixed-effects logistic regression analysis of raw choice data was performed to quantify the model-based and mode-free learning effects (10). The analysis focuses only on the choices made in the first stage of the task and how these are influenced by received a reward and transition type in the previous trial. The regression models the probability of repeating the same choice, p(stay) and how it’s influenced by reward (rewarded=1, unrewarded=-1) and transition (common=1, rare =-1), and their interaction. The main effect of reward is interpreted as model-free learning contribution and the reward x transition interaction as model-based learning

contribution. These effects were regressed alongside with group, condition, and age (z-scored) as fixed effects as below:

glmer(stay ~ r * transition * (group * condition + age_z) + (r * transition + 1 | sub), family = binomial),

The estimates were obtained with the lme4 package in R, using Bound Optimization by Quadratic Approximation (bobyqa) with 1e5 functional evaluations (30).

Stay probabilities calculation (frequency-based)

As an additional measure, for each participant we calculated the following probabilities p(stay|rewarded, common), p(stay|unrewarded, common), p(stay|rewarded, rare), p(stay|unrewarded, rare) based on the frequency of repeated choice in each of the four (reward, transition) cases. These were used for additional regression analyses and figures.

Model-based score calculation (frequency-based)

Yet another measure that helps to capture model-based learning contribution that is independent of the reinforcement learning is the model-based score (MB score) (48). This is calculated as follows:

MBscore = p(stay|rewarded, common) - p(stay|rewarded, rare) - p(stay|unrewarded, common) + p(stay|unrewarded, rare)

Model-based and model free learning - an RL model

Finally, the full computational RL model, described and quantified below (27,30,49), incorporates choice data from two stages of the task, For simplicity, we begin with the update

equation of the state-action value function, $Q^{II}(s_2, c_2)$ for the second stage (II) states and actions. There are two possible stage two states s_2: 1 - the blue island; 2 - the purple island, such that $s_2 \in \{1,2\}$. On each island, an agent can make two separate choices c_2: L - open the left chest; R - open the right chest, such that $c_2 \in \{L, R\}$. Moreover, after opening the chest, the agent can receive the reward, $r \in \{0,1\}$, where 0 corresponds to an empty chest, and 1 corresponds to the pirate coin. At the start of the task all value functions are initialised at 0.5. On any trial, t, we update the value function $Q^{II}(s_2, c_2)$ of the visited state and action taken as in the Eq. 1. The subscript, t, indicates the current trial values, whereas $t + 1$ indicates the values at the following trial. α is the learning rate.

$$Q_{t+1}^{II}(s_{2,t}, c_{2,t}) = (1 - \alpha)Q_t^{II}(s_{2,t}, c_{2,t}) + r_t \tag{1}$$

To calculate the probability of making a choice $c \in \{L, R\}$ at stage two, we use the *softmax* function as in Eq. 2, with an inverse temperature parameter β_2 quantifying the influence of the value function on making the choice.

$$P(c_{2,t} = c) = \frac{\exp\{\beta_2 Q_t^{II}(s_{2,t}, c)\}}{\sum_{c' \in \{L,R\}} \exp\{\beta_2 Q_t^{II}(s_{2,t}, c')\}} \tag{2}$$

In stage one (I), we directly see how model-free and model-based learning play their part. Here, we have two sets of update equations. The first set, Eq. 3-5 is model-based, where the agent uses their knowledge about the environment - which stage two state-action is the best as well as the structure of the task (common vs. rare transition) to update the values of the ships in stage one. In this case, the value functions for both ships are updated simultaneously, such that the value of choosing a ship is the weighted sum (by transition probability) of maximums over stage two actions values.

$$\boldsymbol{Q_t^{MB}} = \left[Q_t^{MB}(c_{1,t} = L), Q_t^{MB}(c_{1,t} = R)\right] \tag{3}$$

$$Q_t^{MB}(c_{1,t} = L) = 0.7 \times \max_{c_2}\{Q_t^{II}(s_2 = 1, c_2)\} + 0.3 \times \max_{c_2}\{Q_t^{II}(s_2 = 2, c_2)\} \quad (4)$$

$$Q_t^{MB}(c_{1,t} = R) = 0.7 \times \max_{c_2}\{Q_t^{II}(s_2 = 2, c_2)\} + 0.3 \times \max_{c_2}\{Q_t^{II}(s_2 = 1, c_2)\} \quad (5)$$

The second, model-free update in Eq. 6, updates the values of the ships based only on the reward received at the end of the trial.

$$Q_{t+1}^{MF}(c_{1,t}) = (1 - \alpha)Q_t^{MF}(c_{1,t}) + r_t \quad (6)$$

Model-free and model-based contributions are joined together in a weighted value function over two ship choices, Q^I, as in Eq. 7.

$$Q_t^I(c) = \beta_{MB}Q_t^{MB}(c) + \beta_{MF}Q_t^{MF}(c) + \rho I(c = c_{1,t-1}) \quad (7)$$

The contribution of each system is captured with β_{MB} (model-based) and β_{MF} (model-free) parameters. An additional indicator $I(c = c_{1,t-1})$ tells if the choice made on the current trial is repeated as in the previous one, with a parameter ρ describing how much switching or staying is done regardless of the feedback. Finally, the probability of choosing either ship is calculated as in Eq. 8, analogously to the second stage.

$$P(c_{1,t} = c) = \frac{\exp\{Q_t^I(c)\}}{\sum_{c' \in \{L,R\}} \exp\{Q_t^I(c')\}} \quad (8)$$

The model has a total of five parameters: $\beta_{MB}, \beta_{MF}, \beta_2, \alpha, \rho$.

Model fitting

Model parameters for each group and condition were estimated using hierarchical Bayesian approaches, which provides the best test-retest reliability for this particular model, and follows the procedure as in Brown et al. (50). The estimation was performed using RStan package (v.2.21.0) (51) in R (v. 4.0.2) based on Markov Chain Monte Carlo techniques (No-U-Turn Hamiltonian Monte Carlo). Each parameter was estimated with a mean, scale, and

individual error estimates. The learning rate α was constrained to (0,1) with an inverse logit transformation, while the means of β_{MB}, β_{MF}, were bounded below by 0. We used weakly informative and uninformative priors for the means: $\alpha \sim N(0,2.5), \beta_{MB} \sim N(0,100)$, $\beta_{MF} \sim N(0,100), \beta_2 \sim N(0,100), \rho \sim N(0,100)$. For the scales of each parameter, we used Cauchy(0,2.5) distribution, all constrained to be greater than 0. Individual error terms were all $N(0,1)$, where for β_{MB}, β_{MF} these were constrained to be greater than 0. For each of the four chains we ran 2000 samples (after discarding 2000 warm-up ones). Chains were inspected for convergence and their $\hat{R}$ values were all around 1 (below 1.1). The mean value (across chains) of each parameter for each participant were used in the analysis.

Mixed-effects linear regression of parameter estimates

The estimates for each parameter were analysed using mixed-effects linear regression to compare them between groups and conditions, with group, condition and age (z-scored) as fixed effects per subject:

```
lmer(parameter ~ group * condition + age_z + (1 | sub))
```

Parameter recovery

To quantify the reliability of the parameter estimates, parameter recovery was performed after data collection (see Supporting Information). For the parameters range compatible with the collected data, the model and fitting procedure described above provide fair to excellent reliability, with average (across parameters, groups and conditions) parameter recovery Pearson correlation coefficient (PCC), r = 0.834.

Results

Demographic and questionnaires summary

The screening procedures resulted in recruiting two significantly different groups - healthy control (HC) and eating disorder (ED). The summarised information, along with two-sample t-tests, can be found in Table 2. The groups are mainly characterised by average total scores on the three questionnaires, with scores significantly different between groups across all questionnaires ($p<0.001$). As the groups also significantly differ in age, the variable (z-scored) was included as a covariate in the regression analysis.

Table 2. Summary of demographic information and questionnaire scores in each group.

Measure	**HC (n=32) Mean ±SD**	**ED (n=35) Mean ±SD**	**t value**	**p value**
EAT-26	2.94 ± 2.72	25.83 ± 10.67	-11.78	**<0.001***
AAI	4.59 ± 2.87	23.94 ± 7.26	-14.10	**<0.001***
OCI-R	5.5 ± 3.59	22.83 ± 11.09	-8.44	**<0.001***
BMI (kg/m^2)	21.83 ± 4.27	26.07 ± 6.21	-3.01	**0.004***
Age	26.38 ± 4.61	30.57 ± 4.45	-3.79	**<0.001***

This includes means and standard deviations (SD) of EAT-26, AAI, OCI-R scores, age, and BMI, as well as t- and p- values of the two-sample t-tests.

Task performance - rewards and reaction times

The performance in the task was analysed independently of the reinforcement learning model. Average characteristics were calculated for each group (detailed results in S3 Table). These measures include: total reward in the neutral and BID condition, and the total reward after completing the full task, as well as mean reaction times (RT) during the neutral, BID, and across both conditions. There were no differences in the above measures between groups.

Collected raw choice data analysis

Mixed-effects logistic regression analysis – collected data.

To assess the model-based and model-free contribution during the task, raw choice data was analysed with mixed-effects logistic regression of the probability of repeating the same choice in stage one (Table 3). HC group showed a significant contribution of both model-free system as indicated by the reward effect (p-value<0.001), and model-based system as indicated by the reward x transition interaction effect (p-value=0.018). According to our hypothesis, we would expect a significantly lower estimate for the reward × transition × groupED effect to detect a difference between groups in the neutral condition; as well as a significantly lower estimate for the reward × transition × groupED × conditionBID effect to detect a difference between condition in the ED group. However, the analysis revealed no differences between groups or conditions in terms of model-based and model-free learning. The estimated probabilities (frequency-based) for each case in each group and condition are depicted in Fig 4.

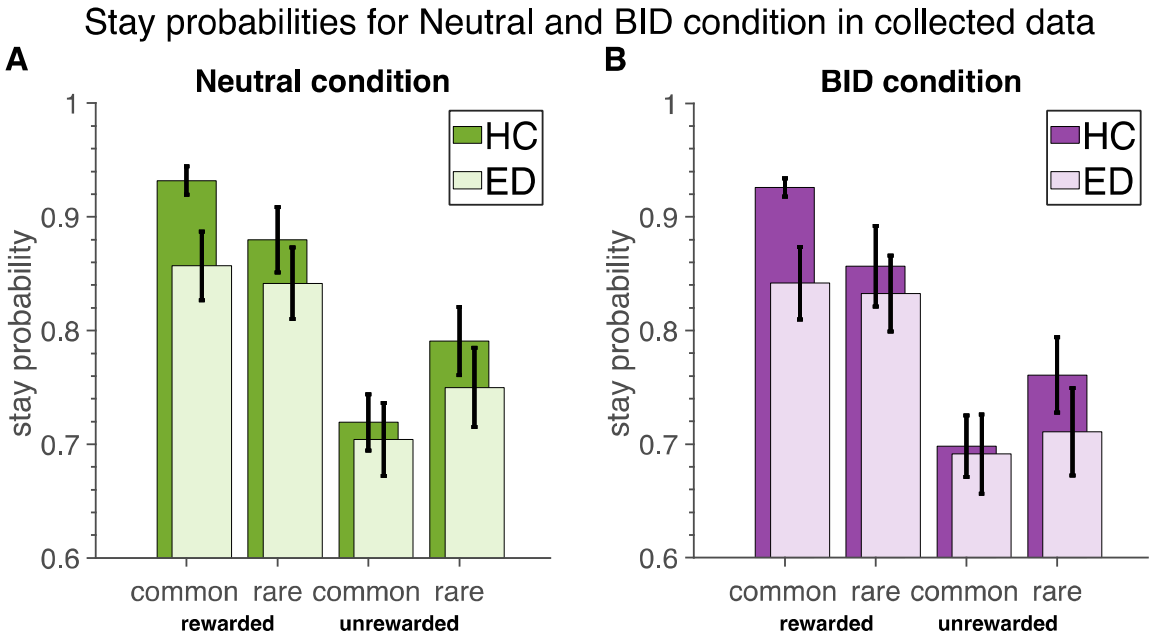

Fig 4. Stay probabilities in the collected data. (A) Neutral condition. (B) BID condition.

Table 3. Random effects logistic regression for probability of staying (collected data).

Effects	Estimate	SE	z value	p value
Intercept (HC, conditionNT)	1.68	0.24	6.86	**<0.001***
reward	0.92	0.33	2.77	**0.006***
transition	-0.57	0.16	-3.68	**<0.001***
groupED	-0.46	0.35	-1.33	0.183
conditionBID	-0.21	0.13	-1.56	0.120
age_z	0.18	0.17	1.07	0.285
reward × transition	0.84	0.36	2.37	**0.018***
groupED × conditionBID	-0.07	0.18	-0.36	0.718
reward × groupED	-0.19	0.47	-0.40	0.689
reward × conditionBID	-0.03	0.22	-0.13	0.900
reward × age_z	-0.05	0.22	-0.21	0.833
transition × groupED	0.26	0.21	1.23	0.219
transition × conditionBID	0.09	0.16	0.57	0.569
transition × age_z	-0.05	0.09	-0.57	0.566
reward × groupED × conditionBID	0.21	0.30	0.72	0.473
transition × groupED × conditionBID	0.11	0.22	0.51	0.612
reward × transition × groupED	-0.43	0.50	-0.87	0.387
reward × transition × conditionBID	0.05	0.27	0.21	0.837
reward × transition × age_z	0.00	0.23	0.00	0.999
reward × transition × groupED × conditionBID	-0.30	0.36	-0.85	0.397

Note. Group, condition, age (z-scored) are treated as fixed-effect covariates per subject.

We further verified whether model-based and model-free learning strategies are used in our sample, regardless of group and condition. As such, these fixed effects were removed from the regression analysis (Table 4). The simpler model revealed significant contribution of both learning systems in the joint population (Reward effect p-value<0.001, Reward x Transition effect p-value=0.008). Moreover, there was a significant effect of transition (p-value <0.001)

Table 4. Random effects logistic regression for probability of staying (collected data).

Effects	Estimate	SE	z value	p value
Intercept	1.31	0.15	8.45	**<0.001***
r	0.87	0.20	4.30	**<0.001***
transition	-0.36	0.09	-4.07	**<0.001***
age_z	0.08	0.16	0.48	0.628
r:transition	0.56	0.21	2.64	**0.008***
r:age_z	-0.07	0.20	-0.34	0.732
transition:age_z	0.02	0.09	0.18	0.857
r:transition:age_z	-0.13	0.21	-0.60	0.552

Note. Age (z-scored) is treated as fixed-effect covariates per subject.

Frequency based-probability regression – collected data.

Furthermore, we estimated the probabilities of staying after common and rewarded trials using the frequency calculation and used them in mixed-effects linear regression analysis (Table 5). This simpler comparison revealed reduced probability of staying after common, rewarded trials between groups in the neutral condition, which was just short of significance (p-value = 0.054), indicating a potentially weaker model-based learning capacity in ED. There were no differences between conditions in either group.

Table 5. Random effects for probability of staying after common and rewarded trials based on collected data (frequency estimate).

Effects	Estimate	SE	t value	p value
Intercept (HC, conditionNT)	0.93	0.03	35.79	**<0.001***
groupED	-0.07	0.04	-1.96	0.054
conditionBID	-0.01	0.01	-0.46	0.649
age_z	0.00	0.02	-0.09	0.930
groupED:conditionBID	-0.01	0.02	-0.50	0.622

Note. Age is treated as fixed-effect covariates per subject.

MB score analysis – collected raw choice data

Lastly, we calculated MB scores in the collected data (Fig *5*). These were regressed against group and condition variables (

Table **6**). There were no differences in MB scores between groups and conditions.

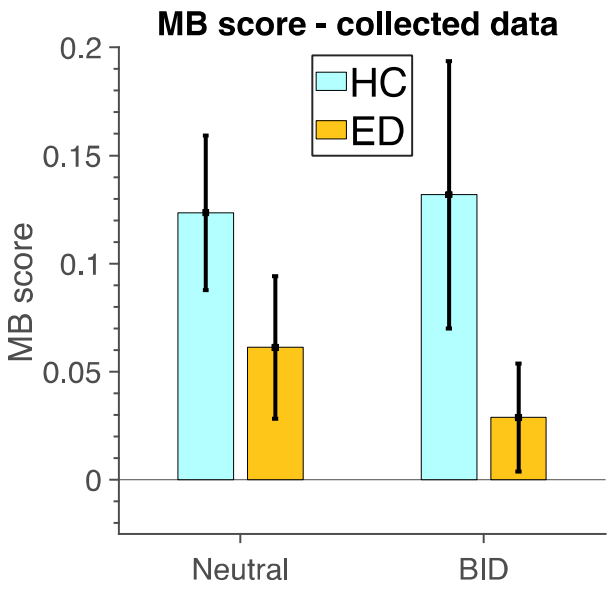

Fig 5. MB scores per group and condition in collected data.

Table 6. Random effects linear regression for MB score (collected data).

Effects	Estimate	SE	t value	p value
Intercept (HC, conditionNT)	0.12	0.04	2.80	**0.006***
groupED	-0.06	0.06	-0.93	0.354
conditionBID	0.01	0.04	0.20	0.845
age_z	-0.01	0.03	-0.20	0.843
groupED:conditionBID	-0.04	0.06	-0.69	0.490

RL parameter estimates - model-based and model-free learning

We next performed a more sensitive analysis, using a RL model, which, unlike the previous analyses which are based on trial averages, takes into account incremental learning over many trials. Model parameters were fit as described in the Methods. A comparison of average

parameter estimates with standard errors (SE) of β_{MB}, β_{MF} between groups and conditions can be found in Fig 6.

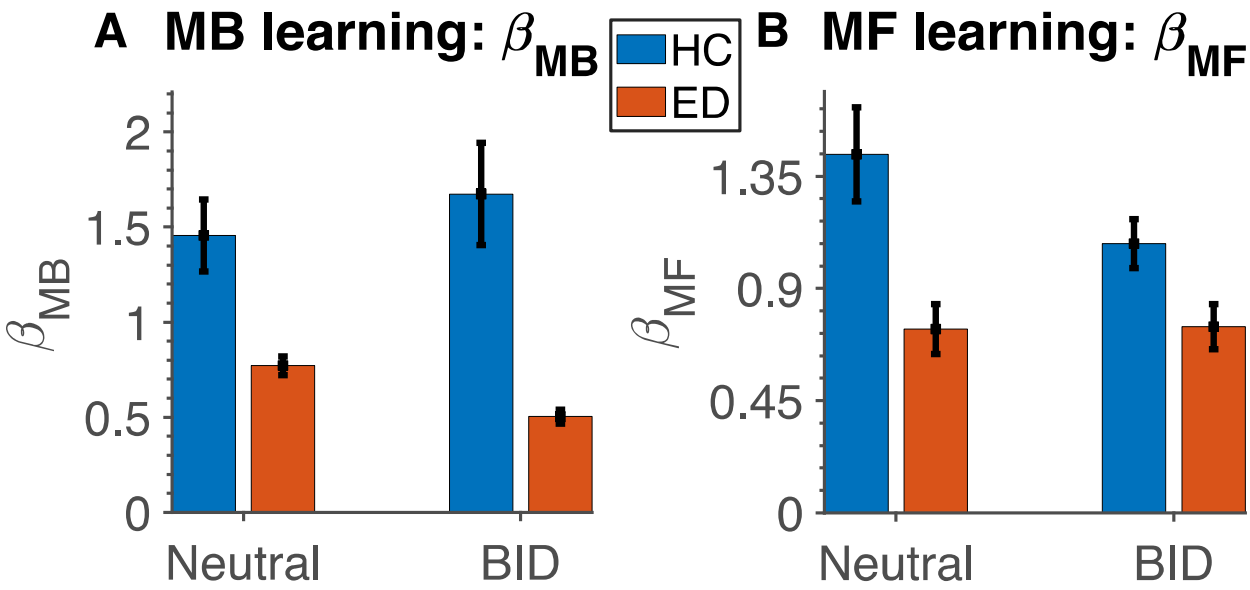

Fig 6. Estimated model parameters.

Mean ± standard error (SE) of the estimated model parameters in HC (blue) and ED (red) during each condition (neutral and BID). (A) β_{MB}, (B) β_{MF}.

The results from the mixed effects linear regression model for each parameter, showing a significant difference between groups and conditions can be found in Tables 7-10.

Table 7. Mixed effects linear regression analysis of model-based learning parameter $\boldsymbol{\beta_{MB}}$

Effects	**Estimate**	**SE**	**t value**	**p value**
Intercept (HC, conditionNT)	1.42	0.17	8.23	**<0.001***
ED group	-0.61	0.25	-2.46	**0.016***
BID condition	0.22	0.11	1.92	0.060
Age	-0.09	0.12	-0.75	0.454
ED group × BID condition	-0.49	0.16	-3.08	**0.003***

Group, condition, age (z-scored) are treated as fixed-effect covariates per subject.

Table 8. Mixed effects linear regression analysis of model-free learning parameter $\boldsymbol{\beta_{MF}}$

Effects	Estimate	SE	t value	p value
Intercept (HC, conditionNT)	1.56	0.14	10.94	**<0.001***
ED group	-0.70	0.21	-3.81	**<0.001***
BID condition	-0.32	0.14	-2.62	**0.011***
*Age	0.02	0.09	0.30	0.764
ED group × BID condition	0.52	0.20	1.94	0.057

Group, condition, age (z-scored) are treated as fixed-effect covariates per subject.

Table 9. Mixed effects linear regression analysis of learning rate parameter $\boldsymbol{\alpha}$

Effects	Estimate	SE	t value	p value
Intercept (HC, conditionNT)	0.61	0.05	12.60	**<0.001***
ED group	-0.09	0.07	-1.33	0.188
BID condition	0.15	0.05	3.02	**0.004***
Age	0.02	0.03	0.50	0.619
ED group × BID condition	-0.10	0.07	-1.43	0.159

Group, condition, age (z-scored) are treated as fixed-effect covariates per subject.

Table 10. Mixed effects linear regression analysis of inverse temperature parameter $\boldsymbol{\beta_2}$

Effects	Estimate	SE	t value	p value
Intercept (HC, conditionNT)	1.63	0.12	13.06	**<0.001***
ED group	-0.60	0.18	-3.34	**0.001**
BID condition	-0.03	0.11	-0.25	0.802
Age	0.10	0.08	1.23	0.224
ED group × BID condition	0.09	0.16	0.56	0.576

Group, condition, age (z-scored) are treated as fixed-effect covariates per subject.

As hypothesised the model-based contribution (as quantified with β_{MB} parameter**)** in the neutral condition is decreased in ED as compared to HC (Table 7; p-value = 0.016). Moreover, there is a further reduction in model-based learning in the BID condition in ED (p-value = 0.003), which is not present in the HC group (p-value = 0.060)

Surprisingly, model-free learning is significantly reduced in the neutral condition in the ED group (compared to HC p-value <0.001, Table 8). Furthermore, model-free learning is slightly attenuated in the BID condition in HC (p-values = 0.011).

In addition to model-based and model-free learning, there was an increase in the learning rate between conditions for HC (p-value = 0.004; Table 9), as well as a reduction in the β_2 inverse temperature parameter estimate in the ED group in the neutral condition (p-value = 0.001; Table 10), potentially indicating a more exploratory choice strategy in the second stage.

Correlations of $\Delta\beta_{MB}$ with other covariates

To check how the difference in model-based learning between condition correlates with questionnaire scores and demographics (age, EAT-26, AAI, and OCI-R scores all z-scored), we introduced $\Delta\beta_{MB}$ that captures the difference in β_{MB} between conditions:

$$\Delta\beta_{MB} = \beta_{MB,BID} - \beta_{MB,NT}$$

when negative, this would indicate the BID condition had a reducing effect on model-based learning.

The correlation plots with Pearson correlation coefficient (PCC), ρ, and p-values for the hypothesis of no relationship between $\Delta\beta_{MB}$ and the covariates can be found in Fig 7. We found significant negative correlations with EAT-26 and AAI scores and OCI-R scores as in Table 11 (r=-0.312, r=-0.314, r=-0.316 with p-value=0.01, p-value=0.01, p-value=0.009). Correlation between questionnaire scores can be found in the Supporting Information (S7 Table).

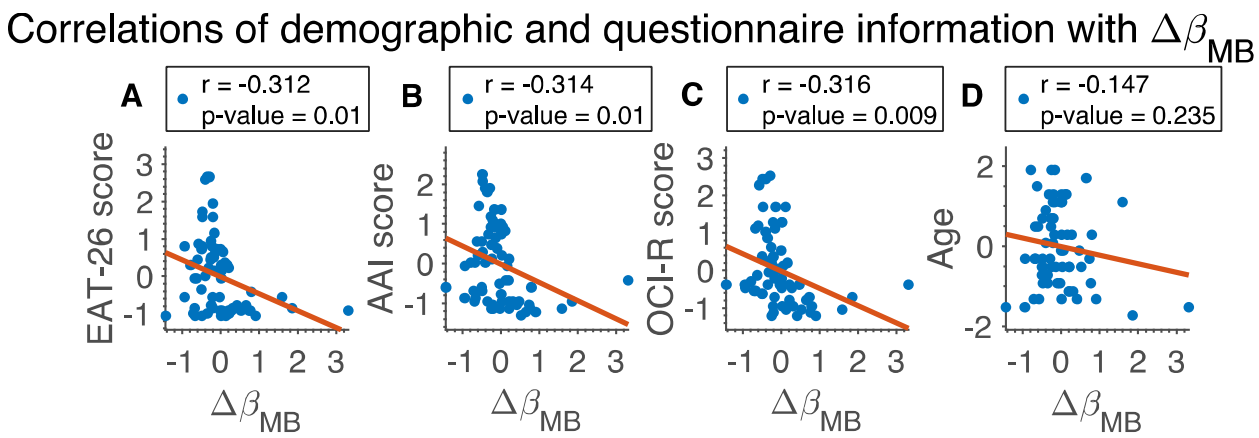

Fig 7. Correlation of EAT-26, AAI, OCI-R and age measurements (z-scored) with $\Delta\beta_{MB}$. Correlation is across both groups. Pearson's correlation coefficient (r) and p-values are included.

Table 11. PCC, r, between $\Delta\beta_{MB}$ and the covariates across HC and ED groups.

	Covariates			
	EAT-26	**AAI**	**OCI-R**	**Age**
r	-0.312	-0.314	-0.316	-0.147
p-value	**0.01***	**0.01***	**0.009***	0.235

P-values for the hypothesis of no relationship between the $\Delta\beta_{MB}$ and covariates (EAT-26, AAI, OCI-R scores and age - all z-scored) are included.

Discussion

The current study focused on a strikingly missing element in computational psychiatry research on eating disorders - body image dissatisfaction. In particular, model-based learning was explored to identify whether deficits in goal-directed learning manifest alongside negative appearance beliefs, with the hypotheses of decreased model-based learning between groups (ED vs. HC) and between conditions in ED.

First of all, the findings from a two-step decision-making task show a significantly smaller contribution of model-based learning in a population characterised by high body image dissatisfaction and disordered eating as compared to the HC group. This replicates previous

findings that suggest a decreased model-based learning, in a monetary condition of the two-step decision-making task, in groups characterised by disordered eating and compulsive behaviours (27,30,34). We found a further model-based learning reduction in the population of interest in the condition that was meant to involve body image dissatisfaction (as implemented in the BID condition). We did not find such an effect in the healthy control. These results confirm the set hypotheses regarding model-based learning. Additionally, the reduction in model-based learning between groups, and between conditions in the ED group was not associated with age.

Surprisingly, contrary to one of our hypotheses, we found decreased (BID vs Neutral) model-free learning in the HC group, as well as reduced model-free learning between groups in the neutral condition. This is also contradictory with previous studies that find no difference in model-free learning between groups in the monetary two-step task (27,30). One potential reason for the difference of results is the high heterogeneity in the estimates of model-free learning parameter both in ED and HC group, which may have skewed the results. In fact, HC group, in general, has more heterogeneous estimates of parameters than ED. This may be due to the HC group being in fact composed of different subgroups, possibly even along the eating disorder continuum, for example if they were not entirely sincere when answering the questionnaires. Alternatively, ED may exhibit a generally reduced learning capability (see below for further discussion).

Furthermore, we found a significant correlation of the difference between the measures of model-based learning in the neutral and BID condition with EAT-26, AAI scores and OCI-R scores, with a similar correlation coefficient. This may suggest that the model-based difference between conditions captures a similar psychiatric dimension common to all three questionnaires, which would also be in line with a significant correlation between EAT-26, AAI, and OCI-R

scores as well as a previously reported results(52). As such, our study may highlight dimensionally-grounded approach to mental illness, similarly to a previous online behavioural study implicating reduced goal-directed learning in compulsive behaviour and intrusive thought dimension (30).

Lastly, the logistic regression of the raw choice data showed the two populations employ both learning systems. However, to detect the differences between groups and conditions we used a more sensitive RL model analysis that takes into account the incremental learning from many trials as well as stage two choices, which are not part of the logistic regression model(27,30,49,53).

Implications

First of all, as far as we know, this is the first study that looked at the body image dissatisfaction from a computational perspective. The fact this phenomenon is associated with decreased contribution of model-based learning provides some support for the mechanism of extreme habitual body preoccupation. The fear of gaining weight and associated with it goal of either losing or not gaining weight set as a goal, evolves over time into a rigid and habitual body checking. When an individual's body type is displayed onto a screen it may act as a trigger for the habitual behaviour of body comparison and concern. As such, a body-preoccupied state leads to significant reduction in the model-based capabilities that could be allocated towards task completion in a more goal-directed manner, as is done in the neutral condition. Our result adds to the discussion reopened by Foerde et al. (27) about domain-general/specific deficits in goal-directed learning. It provides support for the view that goal-directed control impairments are not specific to purely monetary tasks but suggests that it could be aggravated in conditions that trigger body image preoccupations.

The high heterogeneity of parameter estimates (model-based/free) in the healthy control may suggest that some healthy controls share some traits with the ED group. This might be related to traits that we have not directly explored in our study but that have been associated with deficits in model-based learning, e.g. alcohol addiction, or impulsivity (30), despite partially controlling for the effect of compulsivity on model-based learning (30) by excluding participants with high OCI-R scores in the healthy control group. Alternatively, it is possible that some participants from a healthy control group do in fact have body image preoccupation issues. A potential factor in group misclassification could be a wide-spread and widely accepted societal preoccupation with dieting, looks as well as the phenomenon of body shaming that could render the recruitment of a ‘healthy’ population difficult (54,55). This could manifest in participants as selecting the option of “no past diet experience” in the screening stage despite ‘clean eating’ they might engage in as a widely accepted ‘health standard’, which is actually emotionally distressing and linked with functional impairments (56). As such, the heterogeneity in the parameter estimates in the healthy control may highlight the blurred boundary between health and dieting, and indeed call into question the possibility of a truly healthy control group within this field of research.

Limitations

We recognise that the effects may be subtle, and only detectable with a sensitive RL model. As such, it would be beneficial to perform a similar study on a larger online population and/or in the lab on a clinical group to strengthen the evidence for the hypothesis of this paper and test its replicability. This should also further address any possible concerns related to the varied nationalities of the participants.

A possible limitation presents itself in the selection of body types/silhouettes by the participants. These may be seen as overly expressive and caricature-like, failing to accurately capture their body type, which may cast doubt on the source of the reduction in model-based learning. As such, we are currently working on a paradigm that offers a more rigorously defined range of body types for the participants to select from, where a selected image is compared to the participant's basic body parameters (e.g. weight, height), capturing the discrepancy between reality and perception (57–59). That being said, the results of this study offer an new point for further examination: the significant change in goal-directed decision-making exclusively in the ED group, caused by the mere presence of said silhouettes.

An alternative explanation of the results may expose potential issues with the design of the BID condition. Since the silhouette is displayed with the reward, this may induce aversive behaviour (due to high-body dissatisfaction), whereby learning from the rewards is reduced, which results in lower values of inverse temperate parameters both for model-free and model-based learning. Though a valid concern, we did not observe any reduction in model-free learning in the BID condition (vs. neutral) in the ED group. This may suggest that the silhouettes did not have an aversive effect on reward learning, at least in the simple temporal difference learning. As such this may offer support for our case that reward learning within the model-based strategy is similarly unaffected, since we do not see a reason why the negative effect silhouettes would be selective to one type of reward learning. An additional interpretation of the reduced inversive temperature parameters may indicate that the ED group employs a slightly different strategy where they rely more on action exploration rather than exploitation to perform well in the task (19). Nevertheless, it would be useful to more concretely verify the findings by re-running the study using a different design.

Despite employing the most reliable fitting method for this particular task and model (50), as well as applying strict exclusion criteria to participants' behavioural data and questionnaire attention checks, our parameter recovery resulted only in fair reliability of the model-based parameter. However, this is not uncommon, as studies using similar models achieve comparable measures for the model-based parameters (60,61). Having informally tested various models that allow for constraining of the β_{MB}, β_{MF} parameters to be greater than 0, the only way to improve the parameter recovery would be to increase the sample size or obtain higher quality lab-based data.

Kool et al. (61) suggest that the structure of the two step decision making task used here does not accurately estimate the trade-off between model-based and model-free learning. The authors propose certain modifications to the task to increase that accuracy such as changing the drift rate in the Gaussian random walk of reward probabilities, reducing the number of stage two choices to one per state, or introducing a deterministic transition structure. However, it has been shown that in some cases of slight variations to the task structure, the reinforcement learning model will not be able to distinguish between model-based and model-free actions as efficiently as in the standard task (62), while the task in the current form has proven to yield reliable and consistent results (27,30).

Lastly, there are a couple of steps that could be taken in order to further explore the mechanisms and effects of body image dissatisfaction on decision-making in eating disorders. A similar study could be performed on a clinical population, expanding to fMRI data collection and analysis (24). This should allow to correlate the behavioural-computational changes in the ED group with neural signatures providing neurobiological basis of some of the decision-making mechanisms associated with body image dissatisfaction. Ideally, in the future, a more whole

understanding of eating disorders and accompanying body image issues could help devise new and more effective treatment options.

Conclusions

Given the high mortality of eating disorders, this study expands the sparse field of computational psychiatry of eating disorders that so far has focused on the general perception of reward (24,33), and quantification of model-based learning in a neutral setting (30,34) or in relation to food choices (27). Since one of the prevalent aspects of eating disorders is body image preoccupation, we explored its effect on model-based learning in comparison with a healthy control. The results from the online study on a population characterised by high scores on eating disorder and body image dissatisfaction questionnaires show a significantly negative effect of body image dissatisfaction on model-based learning that is not present in the healthy control. This finding offers additional insight into the mechanisms of the disorder and the effect that the core element of the disorder, such as body image dissatisfaction, has on decision-making. Directed by the above result, further inquiry can be performed into more nuanced treatment strategies that could help to break from rigid habits and strengthen model-based capabilities of patients related to the perception of their body through well-tailored therapeutic activities.

References

1. Arcelus J, Mitchell AJ, Wales J, Nielsen S. Mortality Rates in Patients With Anorexia Nervosa and Other Eating Disorders: A Meta-analysis of 36 Studies. Arch Gen Psychiatry. 2011;68(7):724–31.

2. Fichter MM, Quadflieg N. Mortality in eating disorders - results of a large prospective clinical longitudinal study. Int J Eat Disord. 2016;49(4):391–401.

3. Godier LR, Park RJ. Compulsivity in anorexia nervosa: a transdiagnostic concept. Front Psychol. 2014;5.

4. Smink FRE, van Hoeken D, Hoek HW. Epidemiology of Eating Disorders: Incidence, Prevalence and Mortality Rates. Curr Psychiatry Rep. 2012;14(4):406–14.

5. American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders (DSM-5®). American Psychiatric Pub; 2013. 1520 p.

6. Morris J, Twaddle S. Anorexia nervosa. BMJ. 2007;334(7599):894–8.

7. Berkman ND, Lohr KN, Bulik CM. Outcomes of eating disorders: A systematic review of the literature. International Journal of Eating Disorders. 2007;40(4):293–309.

8. Henn AT, Taube CO, Vocks S, Hartmann AS. Body Image as Well as Eating Disorder and Body Dysmorphic Disorder Symptoms in Heterosexual, Homosexual, and Bisexual Women. Front Psychiatry. 2019;10:531.

9. Dolan RJ, Dayan P. Goals and Habits in the Brain. Neuron. 2013;80(2):312–25.

10. Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. Model-Based Influences on Humans' Choices and Striatal Prediction Errors. Neuron. 2011;69(6):1204–15.

11. Garcia F, Delavenne H, Déchelotte P. Atypical eating disorders: a review. NDS. 2011;67.

12. Fairburn CG, Bohn K. Eating disorder NOS (EDNOS): an example of the troublesome "not otherwise specified" (NOS) category in DSM-IV. Behav Res Ther. 2005;43(6):691–701.

13. Kessler RM, Hutson PH, Herman BK, Potenza MN. The neurobiological basis of binge-eating disorder. Neuroscience & Biobehavioral Reviews. 2016;63:223–38.

14. Linardon J, Kothe EJ, Fuller□Tyszkiewicz M. Efficacy of psychotherapy for bulimia nervosa and binge-eating disorder on self-esteem improvement: Meta-analysis. European Eating Disorders Review. 2019;27(2):109–23.

15. Simon JJ, Skunde M, Walther S, Bendszus M, Herzog W, Friederich H-C. Neural signature of food reward processing in bulimic-type eating disorders. Soc Cogn Affect Neurosci. 2016;11(9):1393–401.

16. Huys QJM, Browning M, Paulus MP, Frank MJ. Advances in the computational understanding of mental illness. Neuropsychopharmacology. 2021;46(1):3–19.

17. Huys QJM, Moutoussis M, Williams J. Are computational models of any use to psychiatry? Neural Networks. 2011;24(6):544–51.

18. Seriès P, editor. Computational psychiatry: a primer. Cambridge, Massachusetts: The MIT Press; 2020.

19. Raymond JG, Steele JD, Seriès P. Modeling Trait Anxiety: From Computational Processes to Personality. Front Psychiatry. 2017;8.

20. Rocha FF da, Alvarenga NB, Malloy-Diniz L, Corrêa H. Decision-making impairment in obsessive-compulsive disorder as measured by the Iowa Gambling Task. Arq Neuro-Psiquiatr. 2011;69(4):642–7.

21. Valton V, Romaniuk L, Douglas Steele J, Lawrie S, Seriès P. Comprehensive review: Computational modelling of schizophrenia. Neuroscience & Biobehavioral Reviews. 2017;83:631–46.

22. Huys QJM, Tobler PN, Hasler G, Flagel SB. Chapter 3 - The role of learning-related dopamine signals in addiction vulnerability. In: Diana M, Di Chiara G, Spano P, editors. Progress in Brain Research. Elsevier; 2014. p. 31–77. (Dopamine; vol. 211).

23. Huys QJ, Pizzagalli DA, Bogdan R, Dayan P. Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. Biol Mood Anxiety Disord. 2013 Dec;3(1):12.

24. Bernardoni F, Geisler D, King JA, Javadi A-H, Ritschel F, Murr J, et al. Altered Medial Frontal Feedback Learning Signals in Anorexia Nervosa. Biological Psychiatry. 2018;83(3):235–43.

25. Chan TWS, Ahn W-Y, Bates JE, Busemeyer JR, Guillaume S, Redgrave GW, et al. Differential impairments underlying decision making in anorexia nervosa and bulimia nervosa: A cognitive modeling analysis: Decision-Maing In Anorexia And Bulimia. Int J Eat Disord. 2014;47(2):157–67.

26. DeGuzman M, Shott ME, Yang TT, Riederer J, Frank GKW. Association of Elevated Reward Prediction Error Response With Weight Gain in Adolescent Anorexia Nervosa. AJP. 2017;174(6):557–65.

27. Foerde K, Daw ND, Rufin T, Walsh BT, Shohamy D, Steinglass JE. Deficient Goal-Directed Control in a Population Characterized by Extreme Goal Pursuit. Journal of Cognitive Neuroscience. 2020;1–19.

28. Frank GKW, Reynolds JR, Shott ME, O’Reilly RC. Altered Temporal Difference Learning in Bulimia Nervosa. Biological Psychiatry. 2011;70(8):728–35.

29. Giannunzio V, Degortes D, Tenconi E, Collantoni E, Solmi M, Santonastaso P, et al. Decision-making impairment in anorexia nervosa: New insights into the role of age and decision-making style. European Eating Disorders Review. 2018;26(4):302–14.

30. Gillan CM, Kosinski M, Whelan R, Phelps EA, Daw ND. Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. eLife. 2016;5:e11305.

31. Patzelt EH, Kool W, Millner AJ, Gershman SJ. Incentives Boost Model-Based Control Across a Range of Severity on Several Psychiatric Constructs. Biological Psychiatry. 2019;85(5):425–33.

32. Reiter AMF, Heinze H-J, Schlagenhauf F, Deserno L. Impaired Flexible Reward-Based Decision-Making in Binge Eating Disorder: Evidence from Computational Modeling and Functional Neuroimaging. Neuropsychopharmacol. 2017;42(3):628–37.

33. Verharen JPH, Danner UN, Schröder S, Aarts E, van Elburg AA, Adan RAH. Insensitivity to Losses: A Core Feature in Patients With Anorexia Nervosa? Biological Psychiatry: Cognitive Neuroscience and Neuroimaging. 2019;1–9.

34. Voon V, Derbyshire K, Rück C, Irvine MA, Worbe Y, Enander J, et al. Disorders of compulsivity: a common bias towards learning habits. Mol Psychiatry. 2015;20(3):345–52.

35. Google Scholar. Computational Psychiatry of Eating Disorders Search - Google Scholar [Internet]. [cited 2021 Sep 12]. Available from: https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&q=%28%22computational+psychiatry%22+%22OR%22+%22computational%22+%22OR%22+%22reinforcement+learning%22+%22OR%22+%22reinforcement%22+%22OR%22+%22bayesian%22+%22OR%22+%22decision+making%22+%22OR%22+%22decision-making%22%29+%22AND%22+%28%22anorexia%22+%22OR%22+%22anorexia+nervosa%22+%22OR%22+%E2%80%9Cbulimia+nervosa%22+%22OR%22+%22binge%22+%22OR%22+%22binge+eating%E2%80%9D+OR%22+%22eating%22+%22OR%22+%22bulimia%22%29&btnG=

36. Google Scholar. Computational Psychiatry of Schizophrenia Search - Google Scholar [Internet]. [cited 2021 Sep 12]. Available from: https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&q=%28%22computational+psychiatry%22+%22OR%22+%22computational%22+%22OR%22+%22reinforcement+learning%22+%22OR%22+%22reinforcement%22+%22OR%22+%22bayesian%22+%22OR%22+%22decision+making%22+%22OR%22+%22decision-making%22%29+%22AND%22+%28%E2%80%9Cschizophrenia%E2%80%9D%29&btnG=

37. Prolific. Prolific | Online participant recruitment for surveys and market research [Internet]. [cited 2020 Aug 4]. Available from: https://www.prolific.co/

38. Qualtrics. Qualtrics - Leading Experience Management & Survey Software [Internet]. Qualtrics. [cited 2020 Aug 5]. Available from: https://www.qualtrics.com/uk/

39. Pavlovia. Pavlovia [Internet]. [cited 2020 Aug 5]. Available from: https://pavlovia.org/

40. PsychoPy. Home — PsychoPy v3.0 [Internet]. [cited 2020 Apr 14]. Available from: https://www.psychopy.org/

41. Faul F, Erdfelder E, Lang A-G, Buchner A. G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behavior Research Methods. 2007;39(2):175–91.

42. Garner DM, Olmsted MP, Bohr Y, Garfinkel PE. The Eating Attitudes Test: psychometric features and clinical correlates. Psychological Medicine. 1982;12(4):871–8.

43. Veale D, Eshkevari E, Kanakam N, Ellison N, Costa A, Werner T. The Appearance Anxiety Inventory: Validation of a Process Measure in the Treatment of Body Dysmorphic Disorder. Behavioural and Cognitive Psychotherapy. 2014;42(5):605–16.

44. Foa EB, Huppert JD, Leiberg S, Langner R, Kichic R, Hajcak G, et al. The Obsessive-Compulsive Inventory: Development and validation of a short version. Psychological Assessment. 2002;14(4):485–96.

45. Orbitello B, Ciano R, Corsaro M, Rocco PL, Taboga C, Tonutti L, et al. The EAT-26 as screening instrument for clinical nutrition unit attenders. International Journal of Obesity. 2006;30(6):977–81.

46. Bjureberg J, Enander J, Andersson E, Ivanov VZ, Rück C, Fernández de la Cruz L. Sudden Gains in Internet-Based Cognitive Behavior Therapy for Body Dysmorphic Disorder. Behavior Therapy. 2019;

47. Mastro S, Zimmer-Gembeck MJ, Webb HJ, Farrell L, Waters A. Young adolescents' appearance anxiety and body dysmorphic symptoms: Social problems, self-perceptions and comorbidities. Journal of Obsessive-Compulsive and Related Disorders. 2016;8:50–5.

48. Chen H, Mojtahedzadeh N, Belanger MJ, Nebe S, Kuitunen-Paul S, Sebold M, et al. Model-Based and Model-Free Control Predicts Alcohol Consumption Developmental Trajectory in Young Adults: A 3-Year Prospective Study. Biological Psychiatry. 2021;89(10):980–9.

49. Decker JH, Otto AR, Daw ND, Hartley CA. From Creatures of Habit to Goal-Directed Learners: Tracking the Developmental Emergence of Model-Based Reinforcement Learning. Psychol Sci. 2016;27(6):848–58.

50. Brown VM, Chen J, Gillan CM, Price RB. Improving the Reliability of Computational Analyses: Model-Based Planning and Its Relationship With Compulsivity. Biological Psychiatry: Cognitive Neuroscience and Neuroimaging. 2020;

51. Stan Development Team. RStan: the R interface to Stan [Internet]. 2020. Available from: http://mc-stan.org/

52. Roberts ME. Disordered eating and obsessive-compulsive symptoms in a sub-clinical student population. New Zealand journal of psychology (Christchurch 1983). 2006;35(1):45–54.

53. Otto AR, Raio CM, Chiang A, Phelps EA, Daw ND. Working-memory capacity protects model-based learning from stress. Proceedings of the National Academy of Sciences. 2013;110(52):20941–6.

54. Kar P. Partha Kar: Dieting and body shaming. BMJ [Internet]. 2019 [cited 2020 Aug 16];364. Available from: https://www.bmj.com/content/364/bmj.l1222

55. Sujoldzić A, De Lucia A. A cross-cultural study of adolescents--BMI, body image and psychological well-being. Coll Antropol. 2007;31(1):123–30.

56. Ambwani S, Shippe M, Gao Z, Austin SB. Is #cleaneating a healthy or harmful dietary strategy? Perceptions of clean eating and associations with disordered eating among young adults. Journal of Eating Disorders. 2019;7(1):17.

57. Alexi J, Dommisse K, Cleary D, Palermo R, Kloth N, Bell J. An Assessment of Computer-Generated Stimuli for Use in Studies of Body Size Estimation and Bias. Frontiers in Psychology. 2019;10:2390.

58. Ralph-Nearman C, Filik R. Development and validation of new figural scales for female body dissatisfaction assessment on two dimensions: thin-ideal and muscularity-ideal. BMC Public Health. 2020;20(1):1114.

59. Hudson GM, Lu Y, Zhang X, Hahn J, Zabal JE, Latif F, et al. The Development of a BMI-Guided Shape Morphing Technique and the Effects of an Individualized Figure Rating Scale on Self-Perception of Body Size. European Journal of Investigation in Health, Psychology and Education. 2020;10(2):579–94.

60. Shahar N, Hauser TU, Moutoussis M, Moran R, Keramati M, NSPN consortium, et al. Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling. Gershman SJ, editor. PLoS Comput Biol. 2019;15(2):e1006803.

61. Kool W, Cushman FA, Gershman SJ. When Does Model-Based Control Pay Off? O'Reilly JX, editor. PLoS Comput Biol. 2016;12(8):e1005090.

62. Silva CF da, Hare TA. A note on the analysis of two-stage task results: How changes in task structure affect what model-free and model-based strategies predict about the effects of reward and transition on the stay probability. PLOS ONE. 2018;13(4):e0195328.