Dissociable mechanisms govern when and how strongly reward attributes affect decisions

Silvia U. Maier^{1,2,3 +*}, Anjali Raja Beharelle^{1,2 +*}, Rafael Polanía^{2,4}, Christian C. Ruff^{1,2#},
Todd A. Hare^{1,2# *}

¹Zurich Center for Neuroeconomics, Department of Economics, University of Zurich

²Neuroscience Center Zurich, University of Zurich, Swiss Federal Institute of Technology Zurich

³Translational Neuromodeling Unit, Institute for Biomedical Engineering, Swiss Federal Institute of Technology and University of Zurich

⁴Decision Neuroscience Lab, Department of Health Sciences and Technology, ETH, Swiss Federal Institute of Technology

* corresponding authors

+ Shared first authorship

Shared senior authorship

Abstract

Rewards usually have multiple attributes that are relevant for behaviour. For instance, even apparently simple choices between liquid or food rewards involve comparisons of at least two attributes, flavour and amount. Thus, in order to make the best choice, an organism will need to take multiple attributes into account. Theories and models of decision making usually focus on how *strongly* different attributes are weighted in choice, e.g., as a function of their importance or salience to the decision-maker. However, *when* different attributes impact on the decision process is a question that has received far less attention. Although one may intuitively assume a systematic relationship between the weighting strength and the timing with which different attributes impact on the final choice, this relationship is untested. Here, we investigate whether attribute timing has a unique influence on decision making using a time-varying drift diffusion model (tDDM) and data from four separate experiments. Contrary to expectations, we find only a modest correlation between how strongly and how quickly reward attributes impact on choice. Experimental manipulations of attention and neural activity demonstrate that we can dissociate at the cognitive and neural levels the processes that determine the relative weighting strength and timing of attribute consideration. Our findings demonstrate that processes determining either the weighting strengths or the timing of attributes in decision making can adapt independently to changes in the environment or goals. Moreover, they show that a tDDM incorporating separable influences of these two sets of processes on choice improves understanding and predictions of individual differences in decision behaviour.

Introduction

Decisions regularly involve comparisons of several attributes of the choice options. Consider the example of deciding between foods that differ in two key attributes, tastiness and healthiness. Often these attributes are misaligned, creating a conflict between the goal of eating healthy foods and the desire to experience pleasant tastes. Typically, we assume that choices for the healthier or better tasting food are determined by the values of these attributes, together with a subjective decision weight that the decision maker assigns to healthiness and taste. The assumption that reward attributes are subjectively weighted in the course of decision making applies not only to food choices, but also to many other types of decisions. In fact, it is a core feature of the standard analysis approaches for intertemporal, social, and risky decisions¹⁻⁴. Here, we show that this common approach is incomplete because it overlooks the possibility that reward attributes can enter into the decision process at different times (in addition to having different weighting strengths). Across several food choice paradigms, we find that there is considerable asynchrony in when tastiness and healthiness attributes enter into consideration. Furthermore, we demonstrate that the relative weighting strengths (i.e., the degree to which an attribute influences the evidence accumulation rate) and the onset times for tastiness and healthiness attributes in the decision process have separable influences on whether or not people choose to eat healthier foods.

We used an adapted time-varying sequential sampling model that allows for separate attribute consideration onset times to better understand the dynamic decision processes underlying choices between rewards with multiple attributes. This model allows us to draw inferences on latent aspects of the decision process from the observable choice outcomes and response times. It is well established that direct measures and estimates of information acquisition, evaluation, and comparison processes during choice provide a key means of testing predictions from different models of how stimulus and decision values are constructed or used. Uncovering such features of the decision process allows us to discriminate between and evaluate the plausibility of different models that seek to explain choice behaviour⁵. For example, choice models utilizing not only decision outcomes but

also response times and eye- or mouse-tracking data have provided insights into how and why decision-making is influenced by visual attention, time delays or pressure, additional alternatives, and earlier versus later occurring external evidence⁶⁻¹³. Moreover, it has been shown that dynamic accumulation models utilizing response-time data provide a deeper understanding of decisions and make better out-of-sample predictions than reduced form models such as logistic regressions^{14,15}. Here, we show that we can also use response-time data to determine *when* specific attributes enter into the decision process, in addition to *how strongly* they influence the evidence accumulation rate. Moreover, incorporating this information into the model improves predictions about individual decision-making behaviour.

An important implication of the finding that different attributes can enter into the choice process at separate times is that coefficients from traditional regression models (e.g., linear, logit, or probit) will represent a combination of both the true underlying weight or importance placed on each attribute and its relative (dis)advantage in processing time over the decision period. Therefore, any form of static or synchronous onset dynamic model will fail to fully capture the true underlying choice generating process. By static we mean models that treat values or value-differences as fixed rather than being actively constructed. As a consequence, even though such models may explain multi-attribute choice patterns relatively well if the relationship between attribute weighting and timing is fixed or sufficiently stable, they will fail to explain or predict alterations in decision behaviour if attribute weights and processing onset times can change independently in response to external environmental features or changes in internal cognitive strategies. The plausibility of this latter scenario is underlined by mouse-tracking experiments^{16,17} showing that different attributes (taste, healthiness) of the same food reward can enter into the decision process at separate times. However, the fundamental question of whether the relationship between attribute weighting strength and timing is stable or instead flexible and context-dependent has not yet been addressed.

We addressed this question using an adapted sequential sampling model that quantifies both the weight given to each attribute and its temporal onset during the

decision process. This allows us to explicitly measure whether the weighting strength and timing with which different attributes impact on choice are determined by a unitary process (or a set of consistently linked processes), or if, instead, attribute timing and weighting are the results of separable processes. By modelling choices from four separate datasets, which measured decision behaviour under different experimental manipulations (Figure 1), we show that attribute timing and weighting are determined by dissociable decision mechanisms. For example, we find that explicitly instructing individuals to consider either tastiness or healthiness during the choice process exerts separate effects on attribute weighting strength and timing. In another experiment, we show that transcranial direct current stimulation (tDCS) over the left dlPFC during food decisions has a selective effect on attribute weighting strength but not timing, demonstrating the separability of the underlying neural processes.

Figure 1. Details of the different food choice tasks used in each study. **(a)** On each trial in the mouse response trajectory (MRT) task used by Sullivan et al.¹⁶, participants first saw a start screen and had to respond by continuously moving the mouse toward the option they wanted to choose until they reached the box that contained the desired item. **(b)** In our gambles plus food choices (GFC) study, participants chose between two foods without ever being instructed to think about healthiness. They had up to 3 seconds to make their choice on a 4-point scale ranging from strongly prefer left to strongly prefer right. Intermixed between the food choices were trials on which participants had to select between decks of cards for monetary rewards. **(c)** In the instructed attention cues (IAC) task used by Hare et al.,¹⁸ cues to consider a specific attribute or choose naturally were depicted for 5 seconds before each choice block of 10 trials. Participants then had 3 seconds to make their choice on a 4-point scale from strong no to strong yes. **(d)** In our TDCS study, the reference food for the upcoming block was shown for 3 seconds before each block began. During each block, a series of 10 different foods were shown together with a 4-point scale from strong no to strong yes (in favour of eating the item shown over the reference). The identity of the reference food was written in text below each alternative shown on the screen as depicted in panel **d**.

Results

We adapted the traditional drift diffusion modelling (DDM) framework¹⁹⁻²¹ to allow for each attribute in a multi-attribute decision problem to enter into the evidence accumulation process at separate times (Figure 2a). We chose the DDM as a starting point because this type of sequential sampling model is relatively simple, yet has often been shown to be useful in explaining behaviour across many domains²¹ (see Supplemental section 2.1 for further discussion of this modelling choice). We refer to this modified DDM-like model as a time-varying DDM (tDDM) because the separate consideration onset times for each attribute cause the drift rate to vary over time within a choice. The full details of this tDDM are described in the Methods section, but briefly the relevant addition is a feature allowing for differences in how soon choice evidence related to each attribute begins to be considered. Specifically, we include a free parameter (relative start time (RST)) estimating how quickly one attribute begins to influence the rate of evidence accumulation relative to another. The drift rate determining the evidence update at each time step ($dt = 8$ ms) if taste enters first is

$$(Eq. 1a) \ E_t = E_{t-1} + \left(\omega_{taste} * TD + \left(t > \left| \frac{RST}{dt} \right| \right) * \omega_{health} * HD \right) * dt + noise$$

While if healthiness enters first it is

$$(Eq. 1b) \ E_t = E_{t-1} + \left(\left(t > \left| \frac{RST}{dt} \right| \right) * \omega_{taste} * TD + \omega_{health} * HD \right) * dt + noise$$

Thus, the times at which the weighted value differences in tastiness and healthiness attributes ($\omega_{taste} * TD$ and $\omega_{health} * HD$, respectively) begin to influence the evidence accumulation rate are determined by RST . When the conditional statement $\left(t > \left| \frac{RST}{dt} \right| \right)$ is false, it equals 0, while if true it equals 1. Multiplying one of the two weighted attribute values by 0 until $\left(t > \left| \frac{RST}{dt} \right| \right)$ is true means that that attribute does not factor into the evidence accumulation process for the initial time period

determined by $|RST|$. The RST parameter is defined as the consideration start time for healthiness minus the starting time for tastiness. Thus, RST will have a positive values when tastiness enters into consideration first and negative values when healthiness is considered first. Note that the standard, synchronous onset DDM is equivalent to the specific case of $RST = 0$, and then equations 1a and 1b are equivalent because t is always greater than $(|RST / dt|)$ and $(t > |\frac{RST}{dt}|)$ always equals 1.

In the food choices analysed here (Figure 1), we specified the RST such that it always measures the start time for healthiness relative to taste. In other words, it measures the amount of time that passes between the points at which healthiness versus tastiness attributes begin to influence the rate of evidence accumulation. Based on previous mouse-tracking results from Sullivan et al.¹⁶, our model formulation makes the simplifying assumption that once an attribute comes into consideration it continues to influence the rate of evidence accumulation until the choice is made. Model comparisons testing this assumption showed that the different starting time formulation fit the data better than DDMs that allowed for differential attribute consideration end times, both different start and end times, or starting point biases in favour of tastiness (see Supplemental Table S1).

Furthermore, we found that the attribute timing asynchrony estimated from response times by our tDDM was significantly associated with the results from Sullivan et al's mouse response trajectory analysis. Participants in that study made choices by moving a computer mouse from the bottom centre to the upper left or right corners of the screen to indicate their choices. Sullivan et al. used an analysis based on those computer-mouse response trajectories to determine the relative times at which health and taste attributes enter into the decision. We compared the estimates they obtained from their analysis of response trajectories with those we computed using the tDDM for the same data (see Table 1). We found that the time at which healthiness attributes entered into consideration was significantly correlated across the two analysis methods ($r = 0.503$, Posterior Probability of the

correlation being positive ($PP(r > 0) = 0.991$, 95% Highest Density Interval (HDI) = [0.157; 0.811]).

In total, we tested the tDDM in 272 participants across four datasets with different experimental conditions (mouse response trajectory choices (MRT), standard binary choices in a combined gambling and food choice (GFC) task that was repeated two weeks apart, choices following instructed attention cues (IAC) toward taste or healthiness, and choices under transcranial direct current stimulation (TDCS) (see Figure 1 and the Methods section for details of each study). The tDDM yielded a better fit to choices and response time (RT) distributions than the standard formulation of a DDM with a single, synchronous onset time (overall tDDM BIC = 280632.3, overall standard DDM BIC = 281909; Fig. S1). Parameter recovery tests demonstrated that choice and RT patterns simulated using known values of the standard and tDDM could be recovered in each case (see Supplemental Results section 1.1). Furthermore, the tDDM with separate onset times also generated significantly better out-of-sample predictions for food choices than the standard DDM. The mean squared error for this tDDM (0.204) was lower than that of the standard DDM (0.218) (posterior probability of greater accuracy for this tDDM versus a standard DDM = 0.98; see also Table S19).

Figure 2. Patterns of behaviour resulting from asynchronous evidence accumulation. (a) In this diagram of the drift-diffusion evidence accumulation process, we have assigned the tastier choice to the upper boundary (+1), the healthier choice to the lower boundary (-1), the starting point bias to zero, and the non-decision time to 0.7 s. The coloured lines indicate several example food choice response times from two distinct types of people. The traces that are initially coloured red are simulations of a decision-maker who considers tastiness alone for a period of time before beginning to consider healthiness. The traces that are initially blue are simulations from a decision-maker who considers healthiness alone for that same initial time window. This time-varying DDM includes an additional parameter, the relative start time (RST), that represents the average duration of this difference in consideration start times for each attribute (i.e. the distance in time between the two vertical dotted lines). If a boundary is not reached before the time-window in which one attribute has the initial advantage has elapsed, then both attributes begin to influence the evidence accumulation process. This is shown visually by the lines turning from red or blue to purple (i.e. red + blue). The histograms at each boundary show the distribution of response times for each choice outcome. The red or blue shading of the faster response times highlights the fact that those choices should be more dependent on the initially considered attribute because they were made before the second attribute began to be considered, on average. (b) This plot shows the results of a Hierarchical Bayesian logistic regression that tested for changes in the relative influence of tastiness and healthiness on choice outcomes as a function of which attribute was estimated to be considered first and when the response was made (see Eq. 2). The plot shows the mean plus 95% highest density interval (HDI) for estimates of the influence of taste versus healthiness in the empirical choice data as well as simulated choices generated by the best-fitting standard DDM and tDDM parameters for each participant. These influence estimates were derived from contrasts of coefficients in Eq. 2. Values to the left of zero (dotted line) indicate a stronger influence of healthiness than tastiness, while values greater than zero indicate the opposite. For each choice set (tDDM, Data, DDM), trials were divided into four categories. Those are, 1) HF_After: Healthiness is considered first, but the response is made after the average timing advantage has elapsed; 2) HF_Before: Healthiness is considered first and the response is made during the average initial advantage time window; 3) TF_After: Taste is considered first, but the response is made after the average timing advantage has elapsed; 2) TF_Before: Taste is considered first and the response is made during the average initial advantage time window. The empirical data show a difference in taste versus healthiness influences on choices made before versus after the average time advantage has elapsed, especially when healthiness is considered first. Critically, the tDDM generates choices that reproduce these differences, while the standard DDM does not.

Adding the separate onset time feature allows the model to capture important choice and RT patterns. Specifically, different onset times for the two attributes can

explain the fact that the relative contribution of tastiness and healthiness attributes to the evidence in favour of one option over another changes during the decision process. This change in the relative weighting of taste versus healthiness in our food choice data can be seen when computing a logistic regression model that estimates the influence of taste and healthiness on choices made before or after both attributes were estimated to have begun being considered on average (Table S2). Shared onset time DDMs cannot replicate the effect as shown in Figure 2b and Table S4. We note that separate attribute consideration onset timing is a general feature that could be added to many other types of sequential sampling models in addition to the DDM, e.g. ^{12,22-28}.

This feature of our tDDM differs in important ways from other types of multi-process sequential sampling models that include a combination of fast automatic processes and slower deliberate processes (e.g. dual process, fast guess, Ulrich DMC) ²⁹⁻³². These other frameworks can account for changes in the way evidence is accumulated over time in certain cognitive tasks, but are fundamentally inconsistent with our food choice data. First, responses made before the second attribute enters in to consideration are sensitive to the level of the first attribute indicating that they are not random guesses, or prepotent or habitual responses. Second, the data from the instructed attention cue experiment described below show that whether tastiness or healthiness is considered first is not automatic. Thus, modifying sequential sampling models to allow different attributes to enter into a deliberate consideration process at separate times is more appropriate to explain the outcomes and response times from the goal-directed choice process studied here.

Table 1. Fitted separate attribute onset tDDM parameters for each study and condition

Dataset	Parameter estimate					
<i>(1) MRT</i>	ω_{taste}	ω_{health}	Thr	nDT	RST	Bias
Keyboard trials	1.42 ± 0.45	0.12 ± 1.03	1.04 ± 0.28	0.65 ± 0.14	0.26 ± 0.36	0.11 ± 0.18
Mouse trials	0.94 ± 0.36	0.27 ± 0.32	1.36 ± 0.25	0.77 ± 0.16	0.14 ± 0.34	0.07 ± 0.19
<i>(2) IAC</i>	ω_{taste}	ω_{health}	Thr	nDT	RST	Bias
Natural Choice	1.37 ± 0.79	0.33 ± 1.33	1.27 ± 0.28	0.86 ± 0.12	0.42 ± 0.54	0 ± 0.37
Health Cue	0.98 ± 1.12	1.11 ± 0.60	1.39 ± 0.36	0.85 ± 0.14	-0.06 ± 0.55	-0.22 ± 0.33
Taste Cue	1.42 ± 0.96	0.47 ± 0.99	1.36 ± 0.36	0.83 ± 0.14	0.28 ± 0.49	0.00 ± 0.31
<i>(3) GFC</i>	ω_{taste}	ω_{health}	Thr	nDT	RST	Bias
Session 1	1.11 ± 0.35	-0.07 ± 0.65	1.29 ± 0.17	0.84 ± 0.13	0.3 ± 0.37	-0.01 ± 0.08
Session 2	1.19 ± 0.36	-0.29 ± 0.62	1.19 ± 0.22	0.75 ± 0.12	0.29 ± 0.37	-0.03 ± 0.12
<i>(4) tDCS</i>	ω_{taste}	ω_{health}	Thr	nDT	RST	Bias
Sham Baseline	0.74 ± 0.67	1.01 ± 0.5	1.29 ± 0.23	0.77 ± 0.16	-0.21 ± 0.4	-0.11 ± 0.26
Sham Stimulation	0.67 ± 0.61	1.03 ± 0.55	1.21 ± 0.21	0.71 ± 0.14	-0.1 ± 0.34	-0.14 ± 0.24
Cathodal Baseline	0.63 ± 0.81	1.01 ± 0.62	1.26 ± 0.24	0.75 ± 0.14	-0.05 ± 0.44	-0.07 ± 0.24
Cathodal Stimulation	0.92 ± 0.45	1.07 ± 0.73	1.19 ± 0.21	0.69 ± 0.11	-0.03 ± 0.38	-0.03 ± 0.22
Anodal Baseline	0.6 ± 0.84	1.15 ± 0.48	1.25 ± 0.2	0.75 ± 0.13	-0.09 ± 0.41	-0.08 ± 0.24
Anodal Stimulation	0.85 ± 0.6	1.19 ± 0.62	1.16 ± 0.2	0.7 ± 0.12	-0.09 ± 0.38	-0.04 ± 0.25

All parameters are reported as mean ± standard deviation. Abbreviations:

ω_{taste} = weighting factor determining how much the difference in taste attributes contributes to the evidence accumulation rate.

ω_{health} =weighting factor determining how much the difference in health attributes contributes to the evidence accumulation rate.

Thr = evidence threshold for responding.

nDT=non-decision time and corresponds to the starting time for taste in our model.

RST: relative start time for health (timing relative to start of taste processing, positive values denote that health enters the process later than taste).

Bias: starting point bias for the evidence accumulation process (zero = no bias).

Using the separate attribute consideration onset tDDM to test links between attribute weighting and timing

Having established the face validity, accuracy, and predictive utility of our modelling approach, we next used this tDDM to test several fundamental questions about how attribute timing and weighting work together, or potentially separately, to influence choice outcomes during healthy choice challenges.

Are more abstract attributes considered later in the choice process?

One may assume that for dietary choices, the relative start time of the more abstract attribute (healthiness) will lag behind the more concrete and immediately gratifying attribute of taste. However, our results indicate that this is not the case. Pooling the data across all studies, we found that the posterior probability that healthiness entered into consideration later than tastiness was only 0.48 (mean difference in starting times = 0.001 seconds, 95% HDI = [-0.05; 0.06]). In total, only 130 out of 272 participants (48 %) had relative-start-times for healthiness attributes that were delayed relative to those for tastiness. Figure 3 shows the relative start times for all participants by study. While more abstract or complicated attributes may reliably enter into the decision process later in some decision contexts, our results demonstrate that abstractness is not the main determinant of when an attribute will begin to be considered during a multi-attribute choice (see also supplemental results section 1.3).

Figure 3. Relative start times in seconds for healthiness compared to tastiness for all participants in each study. Positive values indicate that tastiness is considered before healthiness and negative values that healthiness is considered before tastiness. In each column every dot is a separate participant. The thick black horizontal bars represent within-study means and the rectangular bands indicate the 95% highest density intervals (HDIs). Dataset abbreviations: MRT = data from the computer-mouse response trials in Sullivan et al 2015; IAC = data from the natural choice condition in Hare et al 2011; GFC = newly collected data from the first session/day of an experiment combining gambles and food choices; TDCS = newly collected data from the pre-stimulation baseline choices in our tDCS experiment.

How are relative start times related to attribute weights and other tDDM parameters?

We addressed this question by computing a linear regression model that estimated the association between relative start times and all other tDDM parameters (see Eq. 3). The relative-start-time parameter was related to both the tastiness and healthiness weights as well as the starting point bias parameter (Table S5, Table S6), but overall the linear combination of other tDDM parameters explained only 30% of the variability in relative start times across participants. These results suggest that the onset times for choice attributes are not strictly determined by either their relative weights or other aspects of the decision process quantified by the tDDM parameters.

Effects of attention cues on attribute weights and relative-start-times

We also examined whether directing attention toward either healthiness or tastiness could change the time at which those attributes enter the decision process and if changes in timing were linked to changes in weighting strength. This analysis was motivated by previous findings¹⁸ that directing attention to the healthiness aspects of a food item resulted in substantial changes in choice patterns (Figure 4a). In this instructed attentional cues (IAC) experiment, instructive cues highlighted health, taste, or neither attribute for explicit consideration during the upcoming block of 10 food choices. We refer to these three block types as health-cued (HC), taste-cued (TC), and natural-cued (NC). The original analysis of these choice data focused on the regression weights for taste and health attributes in each choice condition but did not consider that the cues might change the relative times at which these attributes entered into the choice process. Our goal was to determine how potential alterations in attribute timing and weighting contributed to the observed changes in choice behaviour during health cue relative to natural blocks. Therefore, we reanalysed the choice data from¹⁸, fitting the tDDM to each attention condition separately.

First, we found that attention cues changed both the relative weighting and timing of taste and healthiness attributes. Compared to the natural choice blocks, 70% of participants reversed their relative weighting of taste and healthiness in taste or health cue blocks (i.e., went from taste > healthiness to taste < healthiness weight or vice versa), and 64% switched whether they considered tastiness or healthiness first. There was no significant difference in how often participants reversed the order of relative weights (switching from $\omega_{\text{taste}} > \omega_{\text{health}}$ to $\omega_{\text{taste}} < \omega_{\text{health}}$, or vice versa) compared to how often they reversed the order of relative onset times (switching from taste first to healthiness first, or vice versa) when moving from attribute-cued and natural-choice blocks (PP(weight reversal more prevalent than timing reversal = 0.70)).

Figure 4. Choice patterns and separate attribute consideration onset tDDM parameter estimates for the IAC study by condition. **(a)** Proportion of times that subjects chose to eat the food (i.e., they responded “yes” or “strong yes”) as a function of attention cue type (Health, Natural, or Taste) and taste-health combination of food under consideration (Tasty or Untasty crossed with Healthy or Unhealthy). In terms of mean choice proportions, directing attention towards healthiness decreased the proportion of choosing healthy-untasty items and decreased the proportion of choosing unhealthy-tasty items compared to the natural condition. The changes in choices during Health blocks were accompanied by higher weights and faster relative-start-times for healthiness. **(b)** Compared to the natural condition, attention cues to health resulted in a higher relative drift weight (arbitrary units) for the corresponding attribute health attribute compared to the natural condition (blue shading). There was no significant change in the relative drift weights of the two attributes during the taste cue blocks (red-shading). **(c)** Attention cues to health also led to a faster relative start time (seconds) for health compared to taste attributes compared to natural blocks (green shading). Once again, there was no significant difference in relative timing between natural and taste blocks. For all plots, the dots within each column represent the value for a single participant in the sample. Darker shading indicates that multiple participants share the same value for that parameter. Black horizontal bars indicate condition means and white, blue, red, or green shaded rectangles indicate the 95% HDIs for each measure. The grey shaded bars in each plot serve to visually separate the columns for each condition and highlight the zero-points on the y-axes.

Focusing our analyses on the health cue blocks that showed a significant change in choice outcomes compared to natural blocks (Fig. 4a), we found that, on average, cuing attention to health attributes both significantly increased the magnitude of participants’ weights for healthiness and sped up the time at which health entered into the accumulation process (relative to taste, i.e., relative start times) (Figure 4b-c; Table 2). These results demonstrate that both the timing and weighting of taste and healthiness attributes can be flexibly and rapidly changed in response to the attention cues preceding every block of 10 choices.

Table 2. Changes in separate attribute consideration onset tDDM parameters between attention cued conditions.

	Mean difference	95% HDI	Posterior Probability
<i>A. Taste weighting (ω_{taste})</i>			
Natural - Health	0.354	[-0.113 0.832]	0.933
Taste - Health	0.455	[-0.088 1.003]	0.951
<i>B. Health weighting (ω_{health})</i>			
Health - Natural	0.746	[0.188 1.325]	0.995
Health - Taste	0.633	[0.245 1.028]	0.999
<i>C. Relative start time of health (RST)</i>			
Natural - Health	0.469	[0.2 0.748]	0.999
Taste - Health	0.336	[0.121 0.548]	0.999

This table shows the effects of attention cues on the tDDM parameters estimated from choice data in the IAC study. Changes in relative starting times (RST) or weighting parameters (ω_{taste} , ω_{health}) induced by the experimental conditions that are shown in bold were significantly different from zero. Mean differences and their 95% highest density intervals (HDI) were computed based on 100,000 samples drawn from the posterior distributions of each parameter³³. The third column displays the posterior probabilities that differences are greater than zero. All comparisons were made so that *a priori* predicted effects would be positive. Abbreviations: ω_{taste} = weighting factor determining how much the difference in taste attributes contributes to the evidence accumulation rate. ω_{health} = weighting factor determining how much the difference in health attributes contributes to the evidence accumulation rate.

Dissociating attribute weighting strengths and timing at the neural level

We next addressed the question of whether attribute weighting strength and timing are implemented by dissociable neural processes. We did so by analysing data from an experiment applying cathodal, anodal, or sham transcranial direct current stimulation (tDCS) over left dlPFC during food choices (see Methods section for details). Numerous neuroimaging and electrophysiological studies have reported correlational evidence for a role of the dlPFC in multi-attribute choice³⁴⁻³⁷. There is also ample evidence that applying brain stimulation (both transcranial direct current and magnetic) over multiple different sub-regions of the left or right dlPFC is associated with changes in several different forms of multi-attribute decision making³⁸⁻⁴⁴. Here, we applied tDCS over a region of the left dlPFC that has been shown to correlate with individual differences in health challenge success rates and multi-

attribute decisions more generally^{18,45-53} in order to uncover the mechanistic changes in the choice process caused by tDCS over this particular region. Stimulation over this region of left dlPFC did not significantly change measures of working memory, response inhibition, or monetary temporal discounting in our participants (Supplemental Results section 1.2).

Previous studies suggest that the effects of stimulation over left dlPFC are strongest on trials in which the participant does not strongly favour one outcome over the other (i.e., stimulation effects are greatest in difficult choices) and depend on baseline preferences over the rewards^{39,43}. Therefore, we restricted our analysis of health challenge success to trials in which the predicted probability of choosing the healthier food was between 0.2 and 0.8 and focused on the difference in behaviour between baseline and active-stimulation choice sessions. Specifically, we computed a Bayesian hierarchical logistic regression model that accounted for both stimulation type and the healthiness and tastiness differences on each trial in the tDCS dataset (see Eq. 6 in the Methods for details). We compared the interaction effects measuring changes in each participant's health challenge success from the pre-stimulation baseline to the active stimulation condition for cathodal and anodal versus sham stimulation groups. This revealed a greater decrease in health challenge success under cathodal relative to sham stimulation (Table S7; regression coef. = -0.32 ± 0.15 , 95% HDI = $[-0.58; -0.08]$, PP(cathodal polarity X active stimulation interaction coef. < 0) = 0.98), but no change in health challenge success for anodal relative to sham stimulation (regression coef. = -0.03 ± 0.15 , 95% HDI = $[-0.28; 0.22]$, PP(anodal polarity X active stimulation interaction coef. > 0) = 0.4). There was also a main effect within the cathodal stimulation group indicating that these individuals had fewer health challenge successes when making food choices under active stimulation compared to their pre-stimulation baseline choices (regression coef. = -0.32 ± 0.15 , 95% HDI = $[-0.63; -0.03]$, PP(active stimulation < 0) = 0.9999). Thus, we find that inhibitory stimulation over left dlPFC leads to fewer health challenge successes (see Figure 5 as well).

Figure 5. Changes in health challenge success following tDCS over left dlPFC. This plot shows the raw changes in health challenge success under stimulation compared to baseline across stimulation groups. Unlike the regression summarized in Table S7, the effects shown in this plot do not account for the taste or health differences in each choice. Each dot represents the difference between active stimulation or sham and baseline in one participant. Left dlPFC-targeted cathodal stimulation significantly decreased health challenge success (mean decrease = $5.7\% \pm 0.02\%$, 95% HDI = $[-10\%; -1.5\%]$, $PP(\text{active cathodal stimulation} < 0) = 0.995$). Note that we tested the change under cathodal stimulation using a method³³ that is robust to outliers such as the two extreme participants near -1. There were no significant differences in healthy choices under anodal or sham stimulation. Black horizontal bars indicate group means and rectangular bands depict the 95% HDIs. The grey shaded bars in each plot serve to visually separate the columns for each condition and highlight the zero-points on the y-axes.

In order to elucidate the changes in choice processes caused by the stimulation, we fit the separate attribute consideration onset tDDM to dietary choices made during the pre-stimulation baseline and active or sham tDCS sessions. When testing how the tDDM parameters changed between baseline and active stimulation sessions in each group, we found that the cathodal group had increased weighting of taste attributes under stimulation compared to baseline choices (mean difference = 0.14, HDI = $[0.03; 0.25]$, $PP(\text{Cath active} > \text{Cath baseline}) = 0.99$; Figure 6a) and that the change from baseline was greater under cathodal stimulation than sham (mean difference = 0.21, 95% HDI = $[0.01; 0.42]$, $PP(\Delta \text{Cath} > \Delta \text{Sham}) = 0.98$). Crucially, the relative-start-time parameters were unaffected during left dlPFC-targeted cathodal tDCS (Table 3, Figure 6b). Moreover, the tDCS-induced changes in taste relative to health weighting parameters and relative start times were not significantly correlated ($r = -0.07$, 95% HDI = $[-0.325; 0.188]$, $PP(r > 0) = 0.30$). Consistent with the

lack of significant change in choice behaviour under anodal tDCS, we found no significant changes in any tDDM parameter under anodal stimulation (Table 3). In summary, we found that cathodal tDCS over left dlPFC changed the relative decision weight placed on taste attributes, but not the speed with which taste, relative to healthiness, began to influence the choice process (Table 3).

Figure 6. Changes in the separate attribute consideration onset tDDM parameters during tDCS over left dlPFC. **(a)** Cathodal stimulation increased the weighting of taste attributes (ω_{taste} , red shading on right) relative to baseline choices, $PP(\text{Cath ST} > \text{Cath BL}) = 0.99$. This change from baseline was greater under cathodal stimulation than sham, $PP((\text{Cath ST} - \text{Cath BL}) > (\text{Sham ST} - \text{Sham BL})) = 0.98$. The red lines and stars highlight this main effect and interaction. Anodal stimulation did not lead to significant changes in attribute weighting parameters, and neither tDCS protocol affected drifts weights for healthiness (ω_{health} , blue shading on left). The weighting strength parameters are plotted in arbitrary units. **(b)** tDCS had no significant effect on the RST parameters (plotted in seconds, green shading). Black horizontal bars indicate group means and blue, red, or green shaded rectangles indicated the 95% HDIs for each parameter. The grey shaded bars in each plot serve to visually separate the columns for each condition and highlight the zero-points on the y-axes.

Table 3. Effects of tDCS over left dlPFC on tDDM parameters.

	Mean difference	95% HDI	Posterior probability
A. Taste weighting (ω_{taste})			
Baseline - Anodal tDCS	-0.094	[-0.23 0.043]	0.081
Cathodal tDCS - Baseline	0.138	[0.027 0.248]	0.993
Baseline - Sham tDCS	0.072	[-0.087 0.23]	0.815
Δ Sham - Δ Anodal	-0.185	[-0.41 0.047]	0.053
Δ Cathodal - Δ Sham	0.215	[0.014 0.42]	0.982
B. Health weighting (ω_{health})			
Anodal tDCS - Baseline	0.098	[-0.023 0.221]	0.941
Baseline - Cathodal tDCS	-0.074	[-0.246 0.102]	0.197
Sham tDCS - Baseline	0.025	[-0.082 0.13]	0.685
Δ Anodal - Δ Sham	0.063	[-0.093 0.223]	0.787
Δ Sham - Δ Cathodal	-0.039	[-0.241 0.164]	0.349
C. Relative start time of health (RST)			
Baseline - Anodal tDCS	-0.002	[-0.098 0.093]	0.484
Cathodal tDCS - Baseline	0.021	[-0.095 0.135]	0.648
Baseline - Sham tDCS	-0.103	[-0.225 0.018]	0.047
Δ Sham - Δ Anodal	0.098	[-0.057 0.251]	0.895
Δ Cathodal - Δ Sham	-0.081	[-0.25 0.086]	0.171

This table reports changes in the separate attribute consideration onset tDDM relative starting times (RST) or weighting parameters (ω_{taste} , ω_{health}) as a result of tDCS over the left dlPFC. Rows in bold indicate changes that are significantly different from zero. The Δ symbol always indicates a difference score equal to the value in the stimulation minus the baseline session within a given condition. Rows containing this symbol report differences of differences across conditions. Mean differences (or differences of differences) and their 95% highest density intervals (HDI) were computed based on 100,000 samples drawn from the posterior distributions of each parameter. The third column displays the posterior probabilities that differences are greater than zero. All comparisons were made so that *a priori* predicted effects would be positive.

Discussion

We have shown that separable mechanisms determine the degree to which an attribute affects the evidence accumulation rate (weighting strength) and the relative speed with which it begins to do so (timing). Thus, both attribute timing and weighting strength must be examined if we seek to better understand decision making at the mechanistic level.

The clearest evidence that timing and weighting strength are dissociable comes from our tDCS experiment showing that stimulation over the left dlPFC caused a change in the weights placed on the taste factor, but not the timing of taste versus healthiness attributes during dietary choices. Moreover, changes in the relative weighting and the relative timing of each attribute between baseline and cathodal stimulation sessions were not significantly correlated, further indicating that the neural mechanisms altered by our tDCS protocol were specifically related to attribute weighting.

The use of analysis strategies that quantify the separate effects on relative timing and weighting is important for interpreting brain stimulation data on the role of dlPFC and other brain regions in value-based choices. Previous studies have reported that stimulation targeted over various brain regions causes changes in several forms of decision making including choices over trade-offs between monetary amounts and risk or time, or between rewards for oneself and others^{38-44,54,55}. Notably, all of these choices involve multi-attribute stimuli and, frequently, conflict between the different attributes. In light of our modelling results, we can speculate that the mechanistic change caused by stimulation over the dlPFC is in the attribute weighting process in some cases. However, the different studies have targeted a range of dlPFC coordinates across both the left and right hemispheres and have used various forms of brain stimulation with potentially different local and widespread effects. Therefore, one should not assume that altered attribute weighting is the mechanistic result of every dlPFC-targeted stimulation protocol. Fortunately, asynchronous evidence accumulation modelling methods, such as the separate attribute consideration onset tDDM used here, could be applied to most of the existing datasets cited above or newly acquired data to gain further insights into how and why brain stimulation causes changes in choice behaviour. Moreover, such

analyses are by no means limited to brain stimulation and can be applied to any set of response-time and choice data on multi-attribute decisions (e.g. self/other, amount/delay, risk/magnitude) under different biological or experimental conditions. This may elucidate other neural regions that are involved in determining the relative timing of attribute consideration.

In our current work, for example, we found that the relative importance given to a specific attribute, as well as its speed in entering into the choice process, could be altered by instructions that directed attention to that attribute. Although a large body of work has established that value construction and comparison processes are malleable and subject to attention, perceptual constraints, and other contextual factors^{9-11,56,57}, the influence of attribute consideration timing within a given decision is rarely discussed or directly tested. Query theory^{58,59} is a notable exception in that it explicitly posits that the order in which attribute values are queried from memory or external sources will bias value construction and choice processes because the recall of initial attributes reduces the accessibility of subsequent attributes. Although the current data cannot be used to address the question directly, future experiments may address the important mechanistic question of whether or not memory retrieval is a driving factor in the consideration onset asynchronies revealed by the separate attribute consideration onset tDDM.

Despite open questions about the relationship between memory and relative starting times, our finding that attribute consideration start times are asynchronous lends strong support to the idea that choices are made based on comparisons of both separate attribute values as well as overall option values. Hunt and colleagues¹¹ demonstrated that a hierarchical sequential sampling process that operates over both separate attribute and overall option values explains risky choice behaviour and brain activity better than models operating only on integrated values. Reeck and colleagues¹⁰ have shown that individual variation in temporal discounting can be explained by patterns of information acquisition that support attribute-wise or option-wise comparisons; moreover, their study shows that choices can be made more patient by an experimental manipulation that promotes attribute-wise comparisons compared to one promoting option-wise comparisons. Together, these results and others (e.g. Roe, et al.²⁷ and Bhatia⁵⁷) indicate that attribute-level

comparisons play an important role in determining choice outcomes. Hierarchical attribute and option-level comparisons are implicit in our specification of the separate attribute consideration onset tDDM because the choice outcome and response time are determined by a weighted sum of the differences in attribute values. However, we show that attribute-level comparisons do not all begin at the same point in time, and that the magnitude of the difference in relative start times across attributes influences option-level comparisons and choice outcomes.

Our results raise important questions about how attribute weighting strengths and onset timing jointly influence choice outcomes: How should we interpret choices in which the outcome is determined by the advantage in relative timing as opposed to weighted evidence? Could this be strategic use of cognitive flexibility to align decision making with current goals or should we consider such outcomes to be mistakes? Traditionally, a weighted combination of all attribute values is assumed to yield the “correct” choice⁶⁰. If the weighting strength on each attribute is appropriate, then any asynchrony in onset timing could produce suboptimal choices (i.e., choices in favour of options with a lower weighted sum over all attribute values than another available alternative). In that sense, it is surprising that we find substantial attribute onset asynchrony in healthy young adults and that, in individuals striving to maintain a healthy lifestyle (i.e., the sample recruited for our tDCS experiment), a higher level of asynchrony is associated with better health challenge success. However, this view is predicated on the assumption that the attribute weighting strengths are appropriate for the current goal or context.

On the other hand, it is possible that shifts in the timing of attribute consideration can be used to achieve the desired outcome. Suppose that a decision maker knows (not necessarily explicitly) that her standard attribute weights are inconsistent with her current decision context or goal, and adjusting those weights by the necessary amount is costly or unlikely. In that case, shifting the relative onset timing could be an effective means of reducing effort and improving the chances of making a goal-consistent choice. In other words, altering the relative starting times may be a form of proactive control^{61,62}. For example, a decision maker who goes on a diet may find it difficult to convince herself that she does not like the taste of ice cream and/or to constantly trade off this delicious taste against the downsides of

excess sugar and fat. An alternative way to bring about health challenge success in this situation may be to adjust the process(es) that determine relative start times for healthiness and tastiness, to focus on the healthiness of each alternative option alone for a brief period in order to forgo extremely unhealthy options (without putting in time or effort to compare taste benefits to health costs).

The use of timing differences as described above would be consistent with at least two existing theories on the role of attention in cognition. First, it is consistent with the idea that rational inattention strategies⁶³⁻⁶⁵ can be employed as a means of reducing effort costs. Specifically, if the time advantage for healthiness is large enough, then one could *theoretically* decide against eating an unhealthy food before even considering its tastiness and thus not experience temptation or conflict. Second, the idea that distinct processes determine consideration onset times and weights for different attributes is paralleled in theories of emotion and food-craving regulation that posit separate attention deployment and stimulus appraisal steps (e.g.^{35,66,67}). However, we don't yet know if strategic use of attribute consideration onsets or related processes actually happen, nor if adjusting the process determining relative onset times is, in fact, less effortful or more likely to succeed than strategies that attempt to alter the attribute weighting strengths.

Altering the processes that determine the relative onset times could be a means or a result of delaying and reducing attention. However, although we found that both cueing attention to healthiness and having the goal of maintaining a healthy lifestyle (tDCS sample vs. all others) were associated with faster average onset times for healthiness attributes, we don't know yet if relative onset times can be manipulated as part of a deliberate strategy. It is also important to note that the response to healthiness cues was heterogeneous in the sense that, although most participants made healthy choices more often following those cues, some participants changed only attribute weights or only attribute start times in favour of healthy choices, rather than both. Further research is needed to understand why individuals responded to these cues in different ways.

The ability to understand or predict how an intervention or policy change will affect choice processes and their outcomes for specific individuals or groups of people is important for any program hoping to promote behavioural change, for

example in domains such as health, crime, or financial stability. Greater knowledge of the cognitive and neural mechanisms that drive choices in specific individuals is an important step toward this understanding⁶⁸. Our findings demonstrate that *when* a specific attribute begins to influence the decision process - a factor that has been generally neglected - is an important determinant of choice outcomes. They also suggest that examining relative differences in attribute start times may prove to be useful in understanding why interventions and policies work in some cases (e.g., for specific individuals or groups) but not in others, and may help to increase their effectiveness. Overall, the work we present here provides both a concrete advancement in our knowledge of multi-attribute choice processes and a functional set of computational modelling tools that can be applied to extract deeper mechanistic insights from data on choice outcomes and response times.

Methods

For all data sets in which we relied on published studies, we included the final reported sample in our analyses. For these studies, we will describe the methodological details relevant for our analyses and refer the reader to the published papers for any further details. All participants provided written informed consent in accordance with the procedures of the Institutional Review Board of the California Institute of Technology, the Institutional Review Board of the Faculty of Business, Economics and Informatics at the University of Zurich, or the Ethics Committee of the Canton of Zurich. All participants received a flat fee to compensate for their time in addition to the food they chose.

Data set 1 – Mouse response trajectories (MRT):

We use the choice and response time data from the study of Sullivan et al.¹⁶ to test the face validity of our time-varying sequential sampling model. These data are openly available at <https://osf.io/jmiwn/>. All participants in the MRT sample were healthy adults and had no specific dietary restrictions. Before making any choices, they were reminded of the importance of healthy eating by reading a short excerpt from WebMD.com before starting the choice task.

Participants. The experiment was approved by the Institutional Review Board of the California Institute of Technology. Twenty-eight (7 female) healthy adult participants completed the study.

Procedure. Participants were asked to fast for 4 hours prior to the study. They first rated 160 foods for taste and health on a 5-point Likert scale with values from -2 (“very little”) to +2 (“very much”). After these ratings, participants were asked to read a short text from WebMD.com on the beneficial effect of healthy eating, in order to increase the frequency with which they used health challenge success in the following dietary choice task. In the choice paradigm, participants made 280 choices between two foods on the screen (see Fig. 1a). The selection ensured that food pairs would represent all possible combinations of taste and health ratings equally. After each block of 40 choices, participants could take a short break. In 240 trials,

participants used the mouse to answer, while in the remaining 40 trials, they answered with the keyboard. In mouse trials, participants had to click the “Start” box at the bottom of the screen to initiate the trial. The cursor reappeared after a random waiting period of 0.2 to 0.5 seconds. From this point on, participants had to move the mouse continuously towards the food they wanted to select. They were instructed to answer as quickly and accurately as possible. A random fixation time of 0.4 to 0.7 seconds separated the trials. In keyboard trials, participants selected food items by pressing the left or right choice keys. At the end of the study, one randomly selected trial was paid out and participants were asked to stay in the lab for 30 minutes or until they had eaten their obtained food.

Data set 2 – Gambles plus food choices (GFC)

Data for this behavioural study (gamble plus food choice, GFC) were collected from the same individuals in two testing sessions two weeks apart. The two sessions were run on the same weekday and daytime in a two-hour visit in the afternoon. Participants in this study were healthy and did not have any specific dietary restrictions. During the study, they chose naturally and were neither reminded about eating a healthy diet nor encouraged to eat healthy in any way.

Participants. The Study was approved by the Institutional Review Board of the University of Zurich’s Faculty of Business, Economics and Informatics. Thirty-seven participants (17 female, mean age = 22.6 ± 3 years SD) were included in this study. A pre-screening procedure ensured that all participants regularly consumed sweets and other snack foods and were not currently following any specific diet or seeking to lose weight. All participants were healthy and had no current or recent acute illness (e.g., cold or flu) at the time of the study. All participants complied with the following rules to ensure comparability across the study sessions: They got a good night’s sleep and did not consume alcohol the evening before the study. On the study day, they took a photograph of the small meal that they consumed 3 hours before the appointment, and sent this photo to the experimenter. One day before the second study session, participants received a reminder about the rules above and were asked to consume a small meal before their second appointment that was

equivalent to their meal before the first test session. Participant received 37.5 CHF (approx. 39 USD) for each session.

Procedure. Participants were asked to eat a small meal of approximately 400 calories 3 hours prior to their appointment and to consume nothing but water in the 2.5 hours before the study started. In the laboratory, participants first rated 180 food items for taste and health. They then made 150 food choices, one of which was randomly selected to be realised at the end of the experiment. On each trial, the screen showed 2 foods next to each other and participants chose the food they wanted to eat using a 4-point scale, picking either “strong left”, “left”, “right”, or “strong right” (Fig. 1b). The pairing order and positions of the foods on the screen (left vs. right) were completely randomized, and the allocation algorithm ensured that one of the foods would be rated as healthier than the other. Participants had 3 seconds to make their choice, with a jittered interval of 1-3 seconds fixation between trials. Between blocks of dietary decisions, participants played a game in which they had to guess cards for monetary rewards. We ignore the card guessing choices for the analyses presented here. At the end of the experiment, participants stayed in the laboratory for an additional 30 minutes during which they ate the food they obtained during the study. Note that participants on the second day saw a new set of choice options that was created based on the taste and health ratings they gave on that second day, using the same allocation algorithm as in session 1.

Data set 3 – Instructed attention cues (IAC)

In order to determine how attention cues affected attribute timing and weighting, we re-analysed data from Hare, et al. ¹⁸. Participants in this study were not following a specific health or dietary goal in their everyday life, but received a cue to think about the healthiness or tastiness of the foods before deciding on a subset of choices in the study.

Participants. The study was approved by the Institutional Review Board of the California Institute of Technology. Thirty-three participants (23 female, mean age

24.8 ± 5.1 years SD) were included. Screening ensured that they were not currently following any specific diet or seeking to lose weight. All participants were healthy, had no history of psychiatric diagnoses or neurological or metabolic illness, were not taking medication, had normal or corrected-to-normal vision, and were right-handed.

Procedure. Participants were instructed to fast and drink only water in the 3 hours prior to the study. In this experiment, participants made a series of 180 choices within an MRI scanner while BOLD fMRI was acquired. The experiment had three conditions with 60 trials each that were presented in blocks of 10, with the order of blocks and foods shown within blocks fully randomized for each participant. Each food was shown only once (Fig. 1c). In condition one, participants were asked to attend to the tastiness of the food when making their choices, in the second condition, to attend to the healthiness of the food, and in the third condition, to choose naturally. The instructions emphasized that participants should always choose what they preferred to eat regardless of the attention/consideration cues. Before each block, the attention condition cue was displayed for 5 seconds. On each choice trial, participants had 3 seconds to answer and were shown feedback on their choice for 0.5 seconds after responding. Trials were separated by a variable fixation period of 4 to 6 seconds. Most participants responded on a 4-point scale “strong yes”, “yes”, “no” or “strong no” to indicate if they preferred to eat or to not eat the food shown on the current trial. Five out of 33 participants completed a version of the task including a fifth option that allowed them to signal indifference between eating and not eating the food. We followed the original analysis procedures in IAC and analysed all 33 subjects as one set. After the scan, participants rated the 180 food items for taste (regardless of health) and health (regardless of taste), with the order of rating types randomized across participants. After both the choice task and ratings were complete, one trial from the choice task was randomly chosen to be realised. Participants were required to eat the food if they answered “yes” or “strong yes”. If they answered “no” or “strong no”, they still had to stay in the laboratory for the 30-minute waiting period; however, they were not allowed to eat any other

food. Participants were fully informed of these choice incentivization procedures before beginning the study.

Data set 4 – Transcranial direct current stimulation study (TDCS)

All participants in this study were pre-screened during recruitment to ensure that they were actively following a healthy lifestyle. They were specifically asked if they would agree to do their best to choose the healthier option whenever possible on the day of the study. Participants who indicated that they would not do so were still allowed to complete the experiment and were reimbursed for their time, but we did not analyse their data. All participants received a flat fee of 100 CHF (approx. 104 USD).

Participants. The Ethics Committee of the Canton of Zurich approved the study protocol and all participants provided written informed consent. In total, 199 participants were enrolled in the study. No participants reported any history of psychiatric or neurological conditions or had any acute somatic illness. Participants were pre-screened in telephone interviews to ensure they did not suffer from any allergies, food intolerances, or eating disorders. To ensure that the snacks in the food choice task would present a temptation, participants were only eligible if they reported regularly consuming snack foods (at a minimum 2-3 times per week) while at the same time trying to maintain an overall balanced and healthy diet. Data from 25 participants were excluded because they failed to meet *a priori* inclusion criteria or data quality checks. Within the study we requested a written statement of compliance with a health goal for the time of the experiment (see below). Seven men and 1 woman indicated they would not comply with the health goal; their data were excluded from all analyses. Note that these participants still completed the experimental procedures and received the same compensation through food and monetary incentives as those who complied, so there was no incentive for the participants to lie about following the health goal. Data from 8 participants had to be excluded because they confused the response keys or forgot the identity of the reference item during the task. Four participants were excluded on site due to safety precautions regarding tDCS. Three participants were excluded

on site because a re-check of the inclusion criteria revealed that they did not actually like snacks or only consumed them on 1-2 occasions per month instead of the minimum 2 times per week. One additional participant had to be excluded because the choice set could not be constructed due to the fact that he reported only the most extreme values on all health and taste ratings. Lastly, data from one participant was excluded because she never made a healthy choice when taste and healthiness were in conflict in the baseline condition, precluding inference about within-subject changes due to stimulation. This left 87 men and 87 women in the final dataset. Participants were randomly allocated to stimulation conditions. The anodal (58 participants, 30 female), cathodal (57 participants, 30 female), and sham (59 participants, 27 female) stimulation groups did not differ from each other with regard to age, body mass index (BMI), or self-reported eating patterns (as assessed by the Three Factor Eating Questionnaire, German validated version by Pudel and Westenhöfer⁶⁷) (see SI Table S8). The groups also did not differ with regard to impulse control (in the stop signal reaction time, SSRT), working memory capacity (digit span test), or time discounting preferences. Finally, the groups did not differ in the level of hunger that they reported before the choice task (see SI Tables S9-16).

tDCS stimulation protocol. The target electrode (5 x 7 cm) was placed on the left dlPFC (see SI Figure S2a). The reference electrode (10 x 10 cm) was placed over the vertex, off-centred to the contralateral side in such a way that a 5 x 7 cm area of the reference electrode was centred over the vertex while the remaining area was placed more to the right side. The target electrode covered the two dlPFC regions depicted in SI Figure S2b (MNI peak coordinates = [-46 18 24] and [-30 42 24]). These targets were selected because they both showed greater activity for health challenge success > failure in two previous fMRI studies^{51,53}. The coordinates for both dlPFC and vertex were identified in each participant's individual T1-weighted anatomical MR image using a neuronavigation system (Brainsight, Rogue Research, RRID:SCR_009539, <https://www.rogue-researcher.com/>; see insert in Figure S2b). We applied anodal, cathodal, or sham tDCS over this dlPFC site using a commercially available multi-channel stimulator (neuroConn GmbH). Between a ramp-up and ramp-down phase of 20 seconds, active stimulation with 1 milliamperes (mA) took

place for 30 minutes (anodal and cathodal group) or 5 seconds (sham). Sham stimulation was delivered with either the anode or the cathode over the dlPFC, counterbalanced over the whole sham group. Both the participants and the experimenters mounting the tDCS electrodes were blind to the stimulation condition.

Procedure. Participants first rated 180 food items for health and taste. They were instructed to rate taste regardless of the healthiness and vice versa for each of our 180 food items on a continuous scale that showed visual anchor points from -5 (“not at all”) to +5 (“very much”). Before or after these ratings, participants completed a battery of control tasks in randomized order. All control tasks were performed both before and after stimulation: a stop signal reaction time task (SSRT), a self-paced digit span working memory (WM) test, and a self-paced monetary inter-temporal choice task (ITC). In order to test for stimulation effects on taste and health ratings, participants also re-rated a subset of foods after stimulation (see supplemental information section 1).

After all pre-stimulation tasks had been completed, but before any food choices were made, we asked participants to sign a health goal statement in which they indicated whether they would commit to maintaining a health goal during the following food choice task or not (see SI section 1.2 for an English translation of the health goal text). Participants indicated that they would or would not commit to the goal, dated, and signed the document, and then handed it back to the experimenter. Participants could not see which option others in the room had selected and the experimenter randomizing the tDCS conditions was blind to the participants’ responses to the health goal.

Just prior to beginning the food choice task, participants indicated their current hunger levels. They then completed a series of food choices. The first 101 participants made 60 food choices at baseline, however we increased the number of baseline choices to 80 for the final 98 participants in order to have an even number at baseline and under stimulation. All other experimental factors were kept the same for all 199 participants. The baseline choices allowed us to make within-subject

comparisons of health challenge success before and during stimulation. Once participants had finished making the baseline choices, stimulation was applied. Participants did not make any choices for the first 3 minutes of stimulation to allow the current to stabilize. Following the stabilization period, they completed another set of food choices ($n = 120$ for participants 1:101 and $n = 80$ for participants 102:199). No choice pairs were repeated between the baseline and stimulation choice sets. However, the difficulty in terms of taste difference was balanced across the two choice sets (see SI).

Participants completed the set of food choices under stimulation (or sham) in a maximum of 16 minutes. In the remaining 8-14 minutes of stimulation (or sham) time, participants completed several control tasks. We randomized the order of the post-stimulation control tasks so that all tasks had an equal chance of being run in the period when current was still being applied versus the 5-10 minute window immediately after stimulation (during which physiological aftereffects of the tDCS were still present, see ^{69,70}). Once they had completed all post-stimulation control tasks, participants filled in a questionnaire battery (Three Factor Eating Questionnaire (TFEQ), Cognitive Reflection Test (CRT), “Big Five” personality dimensions (NEO-FFI), socio-economic status). They also indicated whether and to what degree they had tried to comply with the health goal throughout the study, whether they had felt the stimulation and how strongly, and whether they had any problems understanding or following the instructions. Finally, participants received and ate their selected food 30 minutes after they made their final decision in the food choice task.

Food choice paradigm. Participants were asked to eat a small meal of approx. 400 kcal 3 hours prior to the study and consume nothing but water in the meantime. In the health challenge paradigm, participants chose which food they wanted to eat at the end of the study. In order to comply with their health goal, they had to choose the healthier item as often as they could. However, the paradigm was engineered such that health and taste of the food options always conflicted based on the participant’s ratings, so they would always have to forgo the tastier food in order to

choose healthy. Participants knew that one of their choices would be realised in the end, and they would have to eat whatever they chose on the trial that was randomly selected.

Participants were shown the picture of a reference food for 3 seconds at the beginning of each block. This reference food was either healthier and less tasty than all 10 items shown in the upcoming block or tastier and less healthy than all 10 upcoming items. On each of the 10 trials within a block, participants had to decide if they preferred to eat the food currently shown on the screen or the reference food at the end of the study. The identity of the reference food was written in text on the screen so that participants did not need to remember it (see Figure 1d). During each choice trial, participants had 3 seconds to make their decisions, and each trial was separated by a jittered inter-trial interval of 2-6 seconds. One trial was selected at random to be realised after all experimental procedures were completed. At the end of the study, participants stayed in the lab for 30 minutes to eat the food they obtained in the study.

Statistical Analyses

All analyses presented in this paper were performed with the R (“R Core Team,” 2015), STAN⁷¹ and JAGS⁷² statistical software packages. For all Bayesian modelling analyses, we used the default, uninformative priors specified by the brms package⁷³ or BEST³³ R-packages). Throughout the paper, the notation PP() indicates the posterior probability that the relation stated within the parentheses is true. Whenever we analysed previously published data, we applied the same subject- and trial-level exclusion criteria described in the original papers.

Time-varying Drift Diffusion Model with separate attribute consideration onset times

We fit a drift diffusion model that allowed for differential onset times for taste and health attributes during evidence accumulation to participants choice outcome and reaction time data. Several of the food choice tasks used a 4-point decision strength scale, and for these data we collapsed choices into a binary yes/no or left/right

choice. The following six free parameters were estimated separately for each participant and condition:

Thr: evidence threshold for responding (symmetric around zero)

Bias: starting point bias for the evidence accumulation process

nDT: non-decision time

RST: relative start time for health (positive values mean that health enters the process after taste, negative values mean health enters before taste)

ω_{taste} : weighting factor determining how much taste contributes to the evidence accumulation rate.

ω_{health} : weighting factor determining how much healthiness contributes to the evidence accumulation rate.

The values of these six parameters were used to simulate choices and response times using the sequential sampling model described in the equation below to update the relative evidence level at each subsequent time step t .

If taste enters first, the update equation is

$$\text{(Eq. 1a)} \quad E_t = E_{t-1} + \left(\omega_{\text{taste}} * \text{TD} + \left(t > \left| \frac{RST}{dt} \right| \right) * \omega_{\text{health}} * \text{HD} \right) * dt + \text{noise}$$

While if healthiness enters first it is

$$\text{(Eq. 1b)} \quad E_t = E_{t-1} + \left(\left(t > \left| \frac{RST}{dt} \right| \right) * \omega_{\text{taste}} * \text{TD} + \omega_{\text{health}} * \text{HD} \right) * dt + \text{noise}$$

Thus, the times at which the weighted value differences in tastiness and healthiness attributes ($\omega_{\text{taste}} * \text{TD}$ and $\omega_{\text{health}} * \text{HD}$, respectively) begin to influence the evidence accumulation rate are determined by RST . When the conditional statement

$\left(t > \left| \frac{RST}{dt} \right| \right)$ is false, it equals 0, while if true it equals 1. Multiplying one of the two

weighted attribute values by 0 until $\left(t > \left|\frac{RST}{dt}\right|\right)$ is true means that that attribute does not factor into the evidence accumulation process for the initial time period determined by $|RST|$. The RST parameter is defined as the consideration start time for healthiness minus the starting time for tastiness. Thus, RST will have a positive values when tastiness enters into consideration first and negative values when healthiness is considered first. Note that the standard, synchronous onset DDM is equivalent to the specific case of $RST = 0$, and then equations 1a and 1b are equivalent because t is always greater than $(|RST| / dt)$ and $\left(t > \left|\frac{RST}{dt}\right|\right)$ always equals 1.

Evidence accumulation proceeds according to equation 1a or 1b in the following manner. The evidence accumulation process begins with an initial value (E_0) that is equal to the value of the **Bias** parameter. This value is then updated in discrete time-steps of $dt = 0.008$ s until $|E_t|$ is greater than the threshold (**Thr**) parameter value. The noise at each step of the accumulation process is drawn from a Gaussian distribution with mean 0 and SD = 1.4. The differences in taste and healthiness ratings between Food1 and Food2 (or Food1 vs. 0 for the single item choices in the IAC dataset) on a given trial are denoted by TD and HD, respectively. Once the threshold is crossed, the response time is computed as $t^*dt + nDT$, where **nDT** is a free parameter for a non-decision time that accounts for the time required for any initial perceptual or subsequent motor processes that surround the period of active evidence accumulation and comparison.

We estimated the best values for all six free parameters described above separately for each participant and condition using the Differential Evolution algorithm described in Mullen, et al.⁷⁴ with a population size of 60 members run over 150 iterations. On every iteration, we simulated 3000 decisions and response times for all unique combinations of taste and healthiness trade-offs in the participant's choice set using each population member's six tDDM parameters. We then computed the likelihood of the observed data given the distribution generated by the 3000 simulated choices for a given set of parameters. On each subsequent iteration, the population evolves toward a set of parameters that maximize the likelihood of the observed data using the procedures described by Mullen and

colleagues⁷⁴. We examined the evolution of the population over the 150 iterations/generations and found that the Differential Evolution algorithm settled on set of best-fitting parameters well before 150 iterations in our data sets. The upper and lower bounds on the search space for each of the 6 parameters are listed in Table S17. The ratings for taste and healthiness were z-scored across all available ratings of each type for the whole set of participants in each study.

Lastly, we also fit a standard DDM to all datasets using the same procedures as the tDDM, except that the equation omitted the relative-start-time parameter, which meant that both tastiness and healthiness were constrained to enter into the evidence accumulation process at the same time.

We also note that we fit the tDDM using two levels of resolution for the tastiness and healthiness ratings in the GFC and TDCS studies. The tastiness and healthiness ratings from these two studies were collected on a 426-point visual analogue scale. We initially fit the tDDM using the 426-point ratings scale. We also estimated the fits after first reducing the resolution to 10 equally-sized bins (i.e., 42.6 points per bin) for both taste and health. Both versions yielded very similar results, but the estimation proceeded considerably faster when using the binned ratings because this reduced the number of unique combinations of attributes and therefore the number of simulations required for the fitting procedure. We report the parameter values and results from the model with binned ratings for these studies.

Tests of parameter recovery

We generated simulated choices and reaction times by parameterizing the standard and tDDMs using the best fitting parameters for each model estimated from the choices made by participants in the baseline condition for all four studies. The simulated choice sets were based on these parameters and the tastiness and healthiness differences participants faced on every decision trial. Thus, the simulated choice sets matched the empirical data in terms of trial numbers and attribute difference distributions. Fitting these simulated choices allowed us to quantify both models' ability to recover known parameter values within the context of our

experimental datasets and the ability to distinguish between these models (see Supplemental Figs. S1-S3 and Table S1).

Testing taste versus healthiness influence by response time

In addition to parameter recovery tests, we used the simulated choices to test how well each model reproduced choice and response-time characteristics observed in the empirical data. A hierarchical Bayesian logistic regression analysis showed that the influence of taste and healthiness on choice outcomes differed as a function of response times. (Eq. 2; Tables S2-4). More specifically, this analysis tested the influence of each attribute on trials in which the response was made before versus after the relative-starting-time advantage of the first attribute had elapsed, on average. The population-level regressors are listed in Equation 2 below.

$$\text{Eq. (2) Left} = \beta_0 + \beta_1 \text{HFirst} + \beta_2 \text{InitAdv} + \beta_3 \text{TD} + \beta_4 \text{HD} + \beta_5 \text{HFirst} * \text{InitAdv} + \beta_6 \text{HFirst} * \text{TD} + \beta_7 \text{HFirst} * \text{HD} + \beta_8 \text{InitAdv} * \text{TD} + \beta_9 \text{InitAdv} * \text{HD} + \beta_{10} \text{HFirst} * \text{InitAdv} * \text{TD} + \beta_{11} \text{HFirst} * \text{InitAdv} * \text{HD} + e$$

In this equation, Left is a binary indicator of the choice outcome. HFirst is a dummy variable (1 = healthiness, 0 = taste) indicating which attribute is considered first – as determined by the tDDM. InitAdv is a dummy variable (1 = before, 0 = after) indicating whether the response was made before the median value of the sum of relative-starting-time difference plus non-decision time across participants had elapsed. This sum was equal to 1 second. One second was also the cut-off for first quartile of the response time range, meaning that 25% of choices were made in 1 second or less. The abbreviations TD and HD stand for the differences in tastiness and healthiness, respectively, on each trial. Subject-specific coefficients were estimated for all regressors except HFirst, because each participant had only one level of that regressor in his or her in baseline condition.

We computed this regression using four different subsets of empirical or simulated data. Initially, we analysed the baseline choice trials pooled over all 4 studies (baseline = mouse response trials from MRT, day 1 trials from GFC, no-cue trials from IAC, pre-stimulation trials from TDCS). We then compared this model to

two simpler models that omitted either 1) the dependency on response time (i.e. InitAdv dummy variable) or 2) both the dependency on response time and the indicator for which attribute a participant considered first (i.e. InitAdv and HFirst dummy variables). The full model explained the data better (see supplemental results Table S3), and therefore, we used it to examine choice patterns generated by the standard and tDDM. The means and 95% HDIs for regression coefficients plotted in Figure 2b come from estimating the hierarchical logistic regression in Eq. 2 to observed, standard, or tDDM-simulated choices for all participants in whose |RST| parameter fell into the third quartile. We subset the data into this quartile so that timing differences between taste and healthiness would be big enough to have a clear effect in both the real and simulated data, as well as to reduce the computational burden of fitting the hierarchical logistic regression. The estimates plotted in Fig. 2b were computed by combining the beta coefficients for specific regressors as listed below.

$$TF_After = TD - HD,$$

$$TF_Before = InitAdv * TD - InitAdv * HD,$$

$$HF_After = (HFirst * TD + TD) - (Hfirst * HD + HD),$$

$$HF_Before = (HFirst * InitAdv * TD + InitAdv * TD) - (Hfirst * InitAdv * HD + InitAdv * HD).$$

Essentially, these combinations yield the relative influences of taste and healthiness for each separate category of choices rather than representing differences in the influence across choice categories.

Correspondence of tDDM health delay estimates with MRT estimates

With their mouse response trajectory analysis, Sullivan and colleagues were able to estimate to within a fraction (1/101) of each response time when health first became and remained significant in each choice (their Figure 4b). In order to compare our estimate (which was given in seconds and represents a mean value across all of a given set of choices) to the MRT estimates, we transformed the MRT estimates of start times for health into a mean estimate in seconds as well. Specifically, we took the mean of the estimated trial-wise health start time bins for each participant and multiplied it by the participant's mean RT, then divided by 101. The MRT method was only able to estimate health start times for N=18 (out of 28)

participants and, therefore, we calculated the Bayesian equivalent of Pearson's correlation coefficient between tDDM and mouse-tracking estimates of health start times in this subset of participants. Unless otherwise noted, all correlation coefficients reported in this paper represent the mean of the posterior distribution from a Bayesian correlation analysis. These Bayesian correlations were implemented in R and JAGS based on code published on the blog, doingbayesiandataanalysis.blogspot.com, that accompanies the "Doing Bayesian Data Analysis" book by Kruschke ⁷⁵.

Relationship between relative-start-times and other tDDM parameters

To explain how individual differences in the relative-start-time for healthiness were related to the other tDDM parameters, we estimated the model specified in Eq. 3 (Table S5) below:

$$\text{Eq. (3) RST} = \beta_0 + \beta_1 \omega_{\text{taste}} + \beta_2 \omega_{\text{health}} + \beta_3 \text{nDT} + \beta_4 \text{Thr} + \beta_5 \text{Bias} + \beta_6 \text{Study IAC} + \beta_7 \text{Study MRT} + \beta_8 \text{Study TDCS} + \beta_9 \text{Bias*Study IAC} + \beta_{10} \text{Bias*Study MRT} + \beta_{11} \text{Bias*Study TDCS} + e$$

Note that we interacted the Bias parameter from the tDDM with a dummy variable indicating the Study, because the bias measures different answers across studies given the task designs (e.g. left/right, eat/do not eat). The GFC study served as the baseline in this regression.

Out of sample tests for comparing the standard and tDDMs

We fit the standard and tDDM with separate attribute consideration onsets to the odd-numbered choices from each participant and then compared the accuracy of the two models when predicting even-numbered choice outcomes. We used the squared error in predicting choice outcomes as our measure of accuracy ⁷⁶. The predicted outcome for each choice was computed as the mean outcome over 1000 simulations from the standard and tDDM. Choices for the food on the left or to eat the food in single option decisions were set to a value of 1, and the alternative

choice was set to a value of 0. Thus, the mean outcome from the 1000 simulations for each choice represented the probability of a given outcome. The scoring rule for accuracy on each trial was then computed as: $(\text{True_Outcome} - \text{Prediction})^2$. We computed the squared error separately for tastier and less tasty choice outcomes and then took the mean error across these trials types to obtain a measure of balanced error.

Changes in tDDM parameters between instructed attention conditions

We compared tDDM parameters fit to choices during health-cued (HC), taste-cued (TC), and natural-cued (NC) blocks using a Bayesian t-like test (implemented in the R Package, BEST version 3.1.0³³), which in turn relies on JAGS (version 3.3.0⁷²).

Modelling changes in behaviour under tDCS

We first fit the hierarchical regression model specified in Eq. 6 to the odd-numbered baseline trials in our tDCS dataset. Based on those fitted parameters, we generated predictions about the probability of health challenge success in even-numbered trials as a function of tDCS polarity (anodal, cathodal, sham), stimulation session (baseline, active), health difference, taste difference, and participant identity. We then estimated Eq. 6 on all even-numbered trials for which the probability of health challenge success was predicted to be between 0.2 and 0.8.

To examine whether stimulation over left dlPFC caused changes in health challenge success, we fit a Bayesian hierarchical logistic regression model to the tDCS dataset. The population-level regressors for this model are given in Eq. 4.

$$\begin{aligned} \text{Eq. (4) Health Challenge Success} = & \beta_0 + \beta_1 \text{TD} + \beta_2 \text{HD} + \beta_3 \text{stimulationON} + \beta_4 \text{Cathodal} \\ & + \beta_5 \text{Anodal} + \beta_6 \text{TD} * \text{stimulationON} + \beta_7 \text{HD} * \text{stimulationON} + \beta_8 \text{TD} * \text{Cathodal} + \beta_9 \\ & \text{TD} * \text{Anodal} + \beta_{10} \text{HD} * \text{Cathodal} + \beta_{11} \text{HD} * \text{Anodal} + \beta_{12} \text{stimulationON} * \text{Cathodal} + \beta_{13} \\ & \text{stimulationON} * \text{Anodal} + \beta_{14} \text{TD} * \text{stimulationON} * \text{Cathodal} + \beta_{15} \\ & \text{TD} * \text{stimulationON} * \text{Anodal} + \beta_{16} \text{HD} * \text{stimulationON} * \text{Cathodal} + \beta_{17} \\ & \text{HD} * \text{stimulationON} * \text{Anodal} + e \end{aligned}$$

Here, TD and HD denote the absolute value of taste and healthiness difference between foods on each trial, stimulationON was a dummy variable taking the value 1 under stimulation and 0 at baseline, and Anodal and Cathodal were the active stimulation conditions. The Sham condition was the baseline in this regression. The model included the main effects of all regressors as well as the two and three-way interactions between attribute differences and stimulation Type and session (i.e., baseline vs. stimulation on). The model also included subject-specific intercepts, stimulation effects, and slopes for HD and TD (see Table S7).

References

- 1 Mazur, J. E. in *Quantitative Analyses of Behavior V: The Effect of Delay and of Intervening Events on Reinforcement Value* (eds ML Commons, JE Mazur, & JA Nevin) 55-73 (Erlbaum, 1987).
- 2 Fehr, E. & Schmidt, K. M. A theory of fairness, competition, and cooperation. *Q J Econ* **114**, 817-868 (1999).
- 3 Laibson, D. Golden eggs and hyperbolic discounting. *Q J Econ* **112**, 443-477, doi:10.1162/003355397555253 (1997).
- 4 Kahneman, D. & Tversky, A. Prospect theory: An analysis of decisions under risk. *Econometrica* **47**, 278 (1979).
- 5 Palminteri, S., Wyart, V. & Koehlin, E. The Importance of Falsification in Computational Cognitive Modeling. *Trends Cogn Sci* **21**, 425-433, doi:10.1016/j.tics.2017.03.011 (2017).
- 6 Carland, M. A., Marcos, E., Thura, D. & Cisek, P. Evidence against perfect integration of sensory information during perceptual decision making. *J. Neurophysiol.* **115**, 915-930, doi:10.1152/jn.00264.2015 (2016).
- 7 Summerfield, C. & Tsetsos, K. Building Bridges between Perceptual and Economic Decision-Making: Neural and Computational Mechanisms. *Front Neurosci* **6**, 70, doi:10.3389/fnins.2012.00070 (2012).
- 8 Gluth, S., Hotaling, J. M. & Rieskamp, J. The Attraction Effect Modulates Reward Prediction Errors and Intertemporal Choices. *J. Neurosci.* **37**, 371-382, doi:10.1523/JNEUROSCI.2532-16.2016 (2017).
- 9 Krajbich, I., Armel, C. & Rangel, A. Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* **13**, 1292-1298, doi:10.1038/nn.2635 (2010).
- 10 Reeck, C., Wall, D. & Johnson, E. J. Search predicts and changes patience in intertemporal choice. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 11890-11895, doi:10.1073/pnas.1707040114 (2017).
- 11 Hunt, L. T., Dolan, R. J. & Behrens, T. E. Hierarchical competitions subserving multi-attribute choice. *Nat. Neurosci.* **17**, 1613-1622, doi:10.1038/nn.3836 (2014).

- 12 Kunar, M. A., Watson, D. G., Tsetsos, K. & Chater, N. The influence of attention on value integration. *Atten Percept Psychophys* **79**, 1615-1627, doi:10.3758/s13414-017-1340-7 (2017).
- 13 Forstmann, B. U., Ratcliff, R. & Wagenmakers, E. J. Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions. *Annu. Rev. Psychol.* **67**, 641-666, doi:10.1146/annurev-psych-122414-033645 (2016).
- 14 Clithero, J. A. Improving out-of-sample predictions using response times and a model of the decision process. *J Econ Behav Organ* **148**, 344-375, doi:10.1016/j.jebo.2018.02.007 (2018).
- 15 Webb, R. The (neural) dynamics of stochastic choice. *Manage Sci* (2018).
- 16 Sullivan, N., Hutcherson, C., Harris, A. & Rangel, A. Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychol Sci* **26**, 122-134, doi:10.1177/0956797614559543 (2015).
- 17 Lim, S. L., Penrod, M. T., Ha, O. R., Bruce, J. M. & Bruce, A. S. Calorie Labeling Promotes Dietary Self-Control by Shifting the Temporal Dynamics of Health- and Taste-Attribute Integration in Overweight Individuals. *Psychol Sci* **29**, 447-462, doi:10.1177/0956797617737871 (2018).
- 18 Hare, T. A., Malmaud, J. & Rangel, A. Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **31**, 11077-11087, doi:10.1523/JNEUROSCI.6383-10.2011 (2011).
- 19 Ratcliff, R. A note on modeling accumulation of information when the rate of accumulation changes over time. *J. Math. Psychol.* **21**, 178-184 (1980).
- 20 Ratcliff, R. A theory of memory retrieval. *Psychol. Rev.* **85**, 59-108 (1978).
- 21 Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. Diffusion Decision Model: Current Issues and History. *Trends Cogn Sci* **20**, 260-281, doi:10.1016/j.tics.2016.01.007 (2016).
- 22 Holmes, W. R., Trueblood, J. S. & Heathcote, A. A new framework for modeling decisions about changing information: The Piecewise Linear Ballistic Accumulator model. *Cogn Psychol* **85**, 1-29, doi:10.1016/j.cogpsych.2015.11.002 (2016).
- 23 Cisek, P., Puskas, G. A. & El-Murr, S. Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* **29**, 11560-11571, doi:10.1523/JNEUROSCI.1844-09.2009 (2009).
- 24 Usher, M. & McClelland, J. L. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550-592 (2001).
- 25 Brown, S. D. & Heathcote, A. The simplest complete model of choice response time: linear ballistic accumulation. *Cogn Psychol* **57**, 153-178, doi:10.1016/j.cogpsych.2007.12.002 (2008).
- 26 Busemeyer, J. R. & Townsend, J. T. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432-459 (1993).
- 27 Roe, R. M., Busemeyer, J. R. & Townsend, J. T. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol. Rev.* **108**, 370-392 (2001).

- 28 Stewart, N., Chater, N. & Brown, G. D. Decision by sampling. *Cognit. Psychol.* **53**, 1-26 (2006).
- 29 Ulrich, R., Schröter, H., Leuthold, H. & Birngruber, T. Automatic and controlled stimulus processing in conflict tasks: Superimposed diffusion processes and delta functions. *Cognit. Psychol.* **78**, 148-174 (2015).
- 30 Ratcliff, R. & Rouder, J. N. Modeling response times for two-choice decisions. *Psychological science* **9**, 347-356 (1998).
- 31 Alós - Ferrer, C. A Dual - Process Diffusion Model. *Journal of Behavioral Decision Making* **31**, 203-218 (2018).
- 32 Caplin, A. & Martin, D. The Dual - Process Drift Diffusion Model: Evidence from Response Times. *Econ. Inq.* **54**, 1274-1282 (2016).
- 33 Kruschke, J. K. Bayesian estimation supersedes the t test. *J. Exp. Psychol. Gen.* **142**, 573-603, doi:10.1037/a0029146 (2013).
- 34 Krawczyk, D. C. Contributions of the prefrontal cortex to the neural basis of human decision making. *Neurosci. Biobehav. Rev.* **26**, 631-664 (2002).
- 35 Han, J. E., Boachie, N., Garcia-Garcia, I., Michaud, A. & Dagher, A. Neural correlates of dietary self-control in healthy adults: A meta-analysis of functional brain imaging studies. *Physiol. Behav.* **192**, 98-108, doi:10.1016/j.physbeh.2018.02.037 (2018).
- 36 Kable, J. W. & Glimcher, P. W. The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625-1633, doi:10.1038/nn2007 (2007).
- 37 Rushworth, M. F. & Behrens, T. E. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* **11**, 389-397, doi:10.1038/nn2066 (2008).
- 38 Ruff, C. C., Ugazio, G. & Fehr, E. Changing social norm compliance with noninvasive brain stimulation. *Science (New York, N.Y.)* **342**, 482-484, doi:10.1126/science.1241399 (2013).
- 39 Shen, B. *et al.* High-definition tDCS alters impulsivity in a baseline-dependent manner. *Neuroimage* **143**, 343-352 (2016).
- 40 Essex, B. G., Clinton, S. A., Wonderley, L. R. & Zald, D. H. The impact of the posterior parietal and dorsolateral prefrontal cortices on the optimization of long-term versus immediate value. *J. Neurosci.* **32**, 15403-15413 (2012).
- 41 Cho, S. S. *et al.* Effect of continuous theta burst stimulation of the right dorsolateral prefrontal cortex on cerebral blood flow changes during decision making. *Brain Stimul* **5**, 116-123, doi:10.1016/j.brs.2012.03.007 (2012).
- 42 Camus, M. *et al.* Repetitive transcranial magnetic stimulation over the right dorsolateral prefrontal cortex decreases valuations during food choices. *Eur. J. Neurosci.* **30**, 1980-1988 (2009).
- 43 Figner, B. *et al.* Lateral prefrontal cortex and self-control in intertemporal choice. *Nat. Neurosci.* **13**, 538-539, doi:10.1038/nn.2516 (2010).
- 44 Knoch, D. *et al.* Studying the neurobiology of social interaction with transcranial direct current stimulation--the example of punishing unfairness. *Cerebral cortex (New York, N.Y. : 1991)* **18**, 1987-1990, doi:10.1093/cercor/bhm237 (2008).
- 45 Kool, W., McGuire, J. T., Wang, G. J. & Botvinick, M. M. Neural and behavioral evidence for an intrinsic cost of self-control. *PloS one* **8**, e72626, doi:10.1371/journal.pone.0072626 (2013).

- 46 Rudolf, S. & Hare, T. A. Interactions between dorsolateral and ventromedial prefrontal cortex underlie context-dependent stimulus valuation in goal-directed choice. *J. Neurosci.* **34**, 15988-15996, doi:10.1523/JNEUROSCI.3192-14.2014 (2014).
- 47 Zysset, S. *et al.* The neural implementation of multi-attribute decision making: a parametric fMRI study with human subjects. *Neuroimage* **31**, 1380-1388, doi:10.1016/j.neuroimage.2006.01.017 (2006).
- 48 Kahnt, T., Heinzle, J., Park, S. Q. & Haynes, J. D. Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. *Neuroimage* **56**, 709-715, doi:10.1016/j.neuroimage.2010.05.058 (2011).
- 49 Enax, L., Hu, Y., Trautner, P. & Weber, B. Nutrition labels influence value computation of food products in the ventromedial prefrontal cortex. *Obesity (Silver Spring)* **23**, 786-792, doi:10.1002/oby.21027 (2015).
- 50 Luo, S., Ainslie, G., Pollini, D., Giragosian, L. & Monterosso, J. R. Moderators of the association between brain activation and farsighted choice. *Neuroimage* **59**, 1469-1477, doi:10.1016/j.neuroimage.2011.08.004 (2012).
- 51 Hare, T. A., Camerer, C. F. & Rangel, A. Self-control in decision-making involves modulation of the vmPFC valuation system. *Science (New York, N.Y.)* **324**, 646-648, doi:10.1126/science.1168450 (2009).
- 52 Hare, T. A., Hakimi, S. & Rangel, A. Activity in dlPFC and its effective connectivity to vmPFC are associated with temporal discounting. *Front Neurosci-Switz* **8**, 50, doi:10.3389/fnins.2014.00050 (2014).
- 53 Maier, S. U., Makwana, A. B. & Hare, T. A. Acute Stress Impairs Self-Control in Goal-Directed Choice by Altering Multiple Functional Connections within the Brain's Decision Circuits. *Neuron* **87**, 621-631, doi:10.1016/j.neuron.2015.07.005 (2015).
- 54 Marechal, M. A., Cohn, A., Ugazio, G. & Ruff, C. C. Increasing honesty in humans with noninvasive brain stimulation. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 4360-4364, doi:10.1073/pnas.1614912114 (2017).
- 55 Brevet-Aeby, C., Brunelin, J., Iceta, S., Padovan, C. & Poulet, E. Prefrontal cortex and impulsivity: Interest of noninvasive brain stimulation. *Neurosci. Biobehav. Rev.* **71**, 112-134, doi:10.1016/j.neubiorev.2016.08.028 (2016).
- 56 Caplin, A. Measuring and modeling attention. *Annual Review of Economics* **8**, 379-403 (2016).
- 57 Bhatia, S. Associations and the accumulation of preference. *Psychol. Rev.* **120**, 522-543, doi:10.1037/a0032457 (2013).
- 58 Johnson, E. J., Haubl, G. & Keinan, A. Aspects of endowment: a query theory of value construction. *J. Exp. Psychol. Learn. Mem. Cogn.* **33**, 461-474, doi:10.1037/0278-7393.33.3.461 (2007).
- 59 Weber, E. U. & Johnson, E. J. Mindful judgment and decision making. *Annu. Rev. Psychol.* **60**, 53-85, doi:10.1146/annurev.psych.60.110707.163633 (2009).
- 60 Payne, J. W., Bettman, J. R. & Johnson, E. J. *The adaptive decision maker.* (Cambridge University Press, 1993).
- 61 Aron, A. R. From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biol. Psychiatry* **69**, e55-e68 (2011).

- 62 Castiglione, A., Wagner, J., Anderson, M. & Aron, A. R. Preventing a Thought from Coming to Mind Elicits Increased Right Frontal Beta Just as Stopping Action Does. *Cereb. Cortex* **29**, 2160-2172 (2019).
- 63 Sims, C. A. Implications of rational inattention. *J Monetary Econ* **50**, 665-690 (2003).
- 64 Caplin, A. & Dean, M. Revealed preference, rational inattention, and costly information acquisition. *Am. Econ. Rev.* **105**, 2183-2203 (2015).
- 65 Krajchich, I. & Dean, M. How can neuroscience inform economics? *Current Opinion in Behavioral Sciences* **5**, 51-57, doi:10.1016/j.cobeha.2015.07.005 (2015).
- 66 Giuliani, N. R. & Berkman, E. T. Craving is an affective state and its regulation can be understood in terms of the extended process model of emotion regulation. *Psychological inquiry* **26**, 48-53 (2015).
- 67 Gross, J. J. Emotion regulation: Current status and future prospects. *Psychological Inquiry* **26**, 1-26 (2015).
- 68 Eldar, E., Bae, G. J., Kurth-Nelson, Z., Dayan, P. & Dolan, R. J. Magnetoencephalography decoding reveals structural differences within integrative decision processes. *Nature Human Behaviour* **2**, 670 (2018).
- 69 Nitsche, M. A. & Paulus, W. Sustained excitability elevations induced by transcranial DC motor cortex stimulation in humans. *Neurology* **57**, 1899-1901 (2001).
- 70 Fritsch, B. *et al.* Direct current stimulation promotes BDNF-dependent synaptic plasticity: potential implications for motor learning. *Neuron* **66**, 198-204, doi:10.1016/j.neuron.2010.03.035 (2010).
- 71 Carpenter, B. *et al.* Stan: A probabilistic programming language. *Journal of Statistical Software* **20**, 1-37 (2016).
- 72 Plummer, M. in *Proceedings of the 3rd international workshop on distributed statistical computing*.
- 73 Bürkner, P.-C. brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software* **80**, 1-28, doi:10.18637/jss.v080.i01 (2017).
- 74 Mullen, K., Ardia, D., Gil, D. L., Windover, D. & Cline, J. DEoptim: An R package for global optimization by differential evolution. *Journal of Statistical Software* **40**, 1-26 (2011).
- 75 Kruschke, J. K. *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*. Second edn, (Academic Press / Elsevier, 2015).
- 76 Gneiting, T. & Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* **102**, 359-378 (2007).

Code and data availability

The data analysed in this paper and code for fitting the diffusion models and running the other analyses are openly available on the Open Science Framework at

<https://osf.io/g76fn/>. Additional data for the mouse response trajectory experiments from Sullivan et al. (2015) are available at <https://osf.io/jmiwn/>

Acknowledgements

This study has been supported by a Swiss National Science Foundation grant 320030_143443 (CCR and TAH), a Marie Curie International Incoming Fellowship PIIF-GA-2012-327196 (ARB), and EU FP7 grant 607310 (TAH). The funders had no role in the conceptualization, design, data collection, analysis, decision to publish, or preparation of the manuscript. The authors thank Aidan Makwana and Ana Cubillo for contributing code for the food choice and control tasks, Jonathan Price for help with data collection and documentation, and Gaia Lombardi for help with data collection and documentation as well as useful discussions about analysis strategies and implementation.

Author contributions

All authors designed one or more aspects of the research; S.U.M. and A.R.B. collected the novel data in studies GFC and TDCS; R.P. and T.A.H. designed the time-varying drift diffusion model with separate attribute consideration onset times; S.U.M., A.R.B., R.P., and T.A.H. analysed the data; All authors wrote the paper.

Financial interests

The authors declare no competing financial interests.

Supplementary information

Dissociable mechanisms govern when and how strongly reward attributes affect decisions

Silvia U. Maier^{1,2,3*}, Anjali Raja Beharelle^{1,2*}, Rafael Polanía^{2,4}, Christian C. Ruff^{1,2#},
Todd A. Hare^{1,2#}

¹Zurich Center for Neuroeconomics, Department of Economics, University of Zurich

²Neuroscience Center Zurich, University of Zurich, Swiss Federal Institute of Technology Zurich

³Translational Neuromodeling Unit, Institute for Biomedical Engineering, Swiss Federal Institute of Technology and University of Zurich

⁴Decision Neuroscience Lab, Department of Health Sciences and Technology, ETH, Swiss Federal Institute of Technology

* Shared first authorship

Shared senior authorship

1. Supplemental Results

1.1 Model comparison and parameter recovery tests for the standard DDM and time-varying DDM with separate attribute consideration onset times.

Figure S1. Response time (RT) patterns in the empirical data, and data simulated using the best-fitting standard and tDDM parameters. **a)** These density plots show the RT distributions in the empirical data (green), as well as choices simulated using the standard (red) and tDDM (purple) parameters. The columns separate responses into healthy and unhealthy choice outcomes. In both columns, the response time is shown on the x-axis, with positive RTs indicating left (in GFC, MRT, and TDCS) or eat (in IAC) choices and negative RTs indicating right or don't eat choices. **b)** This panel shows the cumulative distribution function for RTs by choice outcome in each study. The columns separate outcomes into healthy and unhealthy choice categories. In both **(a)** and **(b)**, it is evident that participants respond faster than predicted by the standard DDM. On the other hand, choices generated from the best-fitting tDDM parameters more closely match the empirical data.

The time-varying DDM with separate consideration onset times for taste and healthiness (tDDM) yielded a better fit to choice and reaction time distributions across all subjects ($N = 272$) than the standard formulation of a DDM with a single, synchronous onset time (Figure S1; tDDM BIC = 280632.3, standard DDM BIC = 281909). Furthermore, the improvement in the fit for the tDDM over the standard DDM is proportional to the absolute value of the estimated relative start time for a given participant (Table S18). This relationship is expected because the greater the difference in onset times between taste and healthiness, the more behaviour will deviate from the predictions of the standard, synchronous onset DDM.

We also tested the out-of-sample accuracy for standard versus tDDM fits to our human participants' choices. Specifically, we fit the standard and tDDMs to the odd-numbered choices from each participant and then quantified the accuracy of each model when predicting even-numbered choice outcomes. The mean squared error for the tDDM (0.204) was lower than that of the standard DDM (0.218). Table S18 shows the results of a hierarchical Bayesian linear regression on the trial-wise improvement in prediction error for the tDDM over the standard DDM as a function of the taste and healthiness differences in each choice and the best fitting tDDM parameters for each human participant. Thus, overall the tDDM better explains response-time dependent influences of taste and healthiness on choice outcomes, response time distributions regardless of whether the choice was healthy or unhealthy, and makes more accurate out-of-sample predictions for choice outcomes.

We performed parameter recovery tests for both the standard and tDDM. These tests were based on the actual food choices each of the 272 participants faced in the lab experiments. Specifically, we simulated one set of choices using the best-fitting standard DDM parameters and a second set using the best-fitting tDDM parameters for each participant using the choice options (i.e. taste and healthiness differences) the participant faced in the lab. We then fit both simulated choice sets using the standard and tDDM. This gave us four new sets of fitted parameters that we could compare with the known generating parameters and models.

Parameter recovery was very good for both the standard and tDDM when fitting to the corresponding generating model (Fig. S2 and S3; Table S20). Furthermore, when fitting the tDDM to choices generated by the standard DDM, there was no correlation between errors in the relative weights versus errors in the relative timing (mean $r = 0.02$, PP ($r < 0$) = 0.36, 95% HDI = [-0.09;0.15]). However, when fitting the standard DDM to choices generated by the tDDM, the errors in relative weight versus relative timing parameter recovery were negatively correlated (mean $r = -0.44$, PP ($r < 0$) > 0.99995, 95% HDI = [-0.54;-0.35]). This pattern of biased errors is expected because a standard DDM is forced to account for the difference in attribute consideration onset timing by over-estimating the difference in weights between the two attributes.

Figure S2. Parameter recovery for the time-varying DDM with separate consideration onset times for tastiness and healthiness attributes. The plots in the first column show the distributions of all 272 generating and recovered relative weighting (a) and timing parameters (b). There was no significant difference between generating and recovered relative weighting (mean difference = 0.01, 95% HDI = [-0.36, 0.54], posterior probability of a difference > 0 = 0.662) or relative timing parameters (mean difference = -0.01, 95% HDI = [-0.03, 0.01], posterior probability of a difference > 0 = 0.105). The panels in the second column show the correlations between the generating and recovered relative weighting (c) and timing parameters (d). The red dotted line indicates the $x = y$ identity line. Panel e) plots the error in relative weight recovery against the error in relative timing recovery. This plot shows that there is no significant correlation between the two types of error when fitting the model ($r = -0.1$, 95% HDI = [-0.215; 0.018], posterior probability of observing a negative correlation = 0.95).

Figure S3. Parameter recovery for the standard DDM with simultaneous consideration onset times for tastiness and healthiness attributes. **a)** shows the distributions of all 272 generating and recovered relative weighting parameters. Recall that the relative consideration start time is fixed at zero for the standard DDM. There was no significant difference between generating and recovered relative weighting (mean difference = 0.004, 95% HDI = [-0.05, 0.05], posterior probability of a difference > 0 = 0.58). **b)** shows the correlation between the same generating and recovered parameters from the histogram in panel **a**. The red dotted line indicates the $x = y$ identity line.

In addition to these parameter recovery tests, we performed model identification tests on simulated data. At the model level, we were able to determine whether the generating model was the standard or a time-varying DDM with separate consideration onsets for taste and healthiness in each case. These recovery tests were based on out-of-sample predictive accuracies. To compute the out-of-sample prediction accuracy we generated 1000 choices for 150 new food pairs from 272 real-subject-based simulated agents that used either the generating or recovered standard and tDDM parameters. We then quantified out-of-sample accuracy by computing the squared error between new choices (i.e. the mean left/right outcome over 1000 simulated choices for each food pair) made by the original generating parameters and recovered standard or tDDM parameters. When the tDDM was the generating model, the recovered tDDM parameters had lower squared error (mean = 0.006) than the recovered standard DDM parameters (mean = 0.016). In contrast, when the standard DDM was the generating model, the recovered tDDM parameters had higher squared error (mean = 0.012) than the recovered standard DDM

parameters (mean = 0.008). Table S21 shows the results of a hierarchical Bayesian regression on the trial-wise improvement in prediction error for the tDDM over the standard DDM as a function of the taste and healthiness differences in each choice and the simulated agents' tDDM parameters. Note that in this regression on simulated data, the included tDDM parameters are the ones known to have generated the choices.

1.2 Results from control analyses in the tDCS study.

tDCS effects in other cognitive domains. We assessed potential differences in several cognitive domains under tDCS using the Bayesian regression model specified in Eq. S1 in the supplemental methods. We found no effects of tDCS stimulation over left dlPFC on working memory (Tables S9-10), response inhibition (Tables S11-12), or monetary inter-temporal choice (Tables S13-14). The paradigms used to measure these behaviours are described in the supplemental methods below. Note that although previous papers have reported effects of brain stimulation over the dlPFC in several of these domains, our stimulation target and electrode placement is different from those previous studies.

Pre-choice hunger levels. To assess the hunger level, we report percentages of maximum hunger level as participants indicated it on a visual analogue scale (Table S15). Self-reported hunger did not differ between the stimulation groups (for results of the Bayesian regression model specified in Eq. S2 see Table S16).

Stability of taste and health ratings before and after stimulation. In order to test the stability of taste and health ratings after stimulation, participants re-rated a random subset of 30 foods drawn from the original set of 180 foods after stimulation. We tested the stability of both rating types with Bayesian linear models specified as in Eq. S3 in the supplemental methods below. For the health ratings, neither stimulation session (mean beta estimate = 1.11 ± 0.75 SD, 95% credible interval: [-0.34; 2.58]), nor stimulation condition (estimate = -0.43 ± 0.82 , CI: [-2.05; 1.20]), nor their interaction (estimate = -0.13 ± 0.34 , CI: [-0.82; 0.54]) showed

significant differences at the population level (Table S22a). Taste ratings were slightly higher for all groups after stimulation (estimate = 3.70 ± 0.95 , CI: [1.86; 5.57]), and higher at baseline for the anodal (estimate = 4.92 ± 1.92 , CI: [1.14; 8.71]) and cathodal groups (estimate = 4.18 ± 1.92 ; CI: [0.42; 7.91]) compared to sham, but the critical interaction test showed that stimulation groups did not differ in their taste rating increase between pre- and post-stimulation measurements (estimate of the stimulation session X condition interaction = -0.16 ± 1.31 , CI: [-2.75; 2.41] for cathodal and for anodal estimate = -1.07 ± 1.31 , CI: [-3.69; 1.49]; Table S22b). Note that the ratings were scaled to fall between 0 and 100, and thus the coefficients can be interpreted as percentage of the visual analogue rating scale.

No relationship between RST and SSRT and WM in the baseline trials from the tDCS study. There was no correlation between the start time for health relative to taste (RST) estimated during baseline (i.e., without stimulation) choice trials and inhibitory control, measured as stop signal reaction time (mean $r = 0.05$, PP ($r > 0$) = 0.71, 95% HDI = [-0.12; 0.233]), or RST and working memory capacity, measured as forward digit span (mean $r = 0.02$, PP ($r > 0$) = 0.61, 95% HDI = [-0.125; 0.17]).

1.3 Testing the relationship between attribute rating response times and relative-start-times during choice

We reasoned that rating response times might be an index of how readily taste and healthiness attributes come to mind and play a role in determining their relative onset times. Therefore, we tested whether the time required to report taste and healthiness ratings during the separate rating sessions was associated with the relative-start-times for these attributes during the decision process (see Eq. S4 in supplementary methods). Using the data from the baseline session in our tDCS experiment (i.e., the largest sample in one homogenous experimental setup), we found that the time it took to report taste or healthiness ratings was not significantly associated with relative-start-times (taste regression coef. = 0.05 ± 11 sec, 95% HDI = [-0.17; 0.28], health regression coef = 0.07 ± 0.12 sec, 95% HDI = [-0.17; 0.31]; Table S23).

1.4 The influence of attribute consideration onset times on choice outcomes and response times

We also ran two hierarchical Bayesian regressions showing how relative-start-times (RSTs) relate to both response times and choices. The time-varying DDM with separate attribute consideration onset times predicts that the influence of RST on a given choice will depend on the individual's subjective weights for tastiness and healthiness. Furthermore, it also predicts that the influence of RST on choice outcomes will depend on whether or not the tastiness and healthiness attributes both favor the same option or they are in conflict (i.e. differs between health challenge and non-challenge trials). The regressions in Tables S24 and S25 show that these predicted interactions are present in our data.

2. Supplemental Discussion

2.1 DDM versus other sequential sampling models

We used the DDM as a starting point for our modelling analysis because this flavour of sequential sampling model is relatively simple, well established, and widely used to fit choice and response-time data across cognitive domains. However, a number of different sequential sampling model formulations exist, and in specific cases, these models make different predictions about choice and reaction-time distributions⁷⁷⁻⁸³. However, in our food choice datasets, most of these sequential sampling models will be nearly indistinguishable⁸⁴⁻⁸⁶; we therefore refer to our current model as a sequential sampling model to emphasize that we are adding a feature to one representative of this larger class of models. Adding this flexibility for attributes to enter into consideration at different points is also possible for many other sequential sampling models.

We also note that our results from the tDDM are consistent with theoretical and empirical work showing that sequential sampling models, including drift diffusion models and linear ballistic accumulators, can capture changes in perceptual decision processes that result from known changes in externally presented evidence over time^{82,87-90}. However, in contrast to previous work on perception, we tested for asynchronous attribute consideration onsets in value-based choices for which the

externally presented evidence is constant. In other words, we examined timing differences resulting from internal cognitive and neural processes instead of changes in the stimuli themselves. Regardless of the context (e.g. perceptual, value-based) or the exact form of sequential sampling, the key feature that allows these models to explain the data is the ability to account for variability in the strength of evidence over time.

Lastly, we note that an attribute-specific, independent race model is not a parsimonious explanation of our data. In an independent race model, the first competitor to cross the threshold determines how the decision is made⁹¹. Therefore, in a race model that included separate and independent accumulators for each attribute value or difference in attribute values, the response times will only be proportional to the value of the winning attribute (i.e. the one that crosses the threshold first). This prediction of race models does not hold in our data. Regardless of whether or not participants chose the healthier or tastier option, response times are significantly related to the difference in both attributes (Tables S26 and S27). Moreover, response times also depend on whether or not the subjective value derived from the initially considered attribute alone is compatible (i.e. points to the same choice) with a subjective value computed by integrating across both taste and healthiness attributes (see Table S28). Thus, our results are more consistent with hierarchical comparisons of attributes on one level and integrated option values at a second level, as in Hunt and colleagues' fMRI study of risky choices⁷⁷. There must be integration across attributes at some level in the evidence accumulation hierarchy or mutual inhibition between attribute-specific accumulators to explain the response time results we find.

3. Supplemental Methods

TDCS electrode montage

Figure S4. Left dlPFC target regions in the tDCS study. **(a)** The stimulation was centred over two previously identified dlPFC regions of interest (MNI peak coordinates = [-46 18 24] and [-30 42 24]) based on contrasts of health challenge success > failure in two previous fMRI studies (Hare, et al. ⁹² and Maier, et al. ⁹³). **(b)** The dlPFC and vertex coordinates were identified with neuronavigation based on anatomical brain scans for each participant. A 5x7 cm electrode (indicated by the white frame) was placed over left dlPFC to cover both stimulation targets, and a 10 x 10 cm reference electrode was placed over the vertex slightly offset to the right hemisphere (so that the centre of a comparable 5x7 cm area would be centred over the meeting point of the two central sulci, see Methods).

TDCS control task descriptions

Stop signal reaction time task (SSRT). We used a standard stop-signal-reaction time task ⁹⁴⁻⁹⁷ in order to assess whether stimulation changed inhibitory control.

Participants had to press a button as quickly as they could whenever a figure appeared on the screen (“go task”), but had to stop the initiated movement if another figure appeared above the first with a few milliseconds delay (“stop signal”). The initial delay between the stop signal and the go signal was 0.25 seconds. The task was adaptive and added 0.05 seconds delay to the next inhibition trial whenever the participant’s rate of successful movement inhibition was greater than 50% over all previous inhibition trials (adding up to a delay of maximum 0.95

seconds), and subtracted 0.05 seconds whenever the participant's success rate in inhibiting the button press fell below 50% over all inhibition trials completed up to this point. Stimuli were presented on the screen with a jittered duration between 0.5 and 1.25 seconds, and late responses that were given after the stimulus had disappeared from the screen were not counted as correct. Trials with (25) and without stop signal (75) were randomly mixed within the run.

The stop signal reaction times (calculated by subtracting the average presented delays from the average reaction times) were negative for a number of participants, indicating that they adopted a strategy of waiting for the stop signal to appear before initiating a response. In other words, they did not follow the task instructions to initiate a response as soon as the go signal appeared on the screen. We removed participants showing this pattern from all analyses using SSRTs. The remaining sample included 39 Anodal, 41 Cathodal and 42 Sham participants.

Digit span task. In order to test and, if necessary, account for changes in working memory capacity under tDCS, participants completed a computerized digit span task according to the procedure of Wechsler⁹⁸. The screen first showed a series of 5 numbers, each for 1 second, and then prompted the participant to enter the numbers as she remembered them. If the participant entered a correct sequence in ascending ("forward") order two times in a row, the difficulty level increased by one digit (up to a maximum of 12 digits). If the participant failed two times in a row, or alternated between correct and incorrect answers more than seven times without reaching two sequential correct responses, the task stopped and prompted participants to enter the digits in reverse of the order they were originally shown ("backward") in the next round of trials. Within the backwards digit span trials, participants also needed two consecutive correct responses to reach the next level, stopped at 12 digits, or if they reached the failure criteria described above. We excluded data from two participants for the following reasons: For one participant in the Sham group, data for the post-stimulation control were lost due to a computer crash. One participant in the Anodal group was detected to cheat by writing down all sequences on the instruction sheet. We report data from 172 participants (57 Anodal, 57 Cathodal and 58 Sham). Note that due to an oversight in the task

programming, it was possible to cheat and solve the backward digit span by entering the numbers in the forward direction (i.e., the program did not force participants to enter responses in the requested order). We therefore only evaluated the forward digit span.

Inter-temporal choice task. To test and, if necessary, account for possible effects of tDCS on discounting behaviour, we ran an inter-temporal choice (ITC) task based on the paradigm of Cooper, Kable, Kim & Zauberman⁹⁹. Participants were instructed that they would earn part of their total payment in this task (60 CHF were paid as a baseline on the day of the study, and the present discounted value of 40 CHF from the ITC was paid at the time specified by the participant). They were told that we would randomly draw select one trial from either the baseline or stimulation session and realize (i.e., pay out) their decision on that trial. In this version of the ITC task, participants participated in a BDM auction¹⁰⁰ designed to elicit their indifference points between a payoff on the day of the experiment and 40 CHF at various points in the future. Participants were asked to specify a bid between 1 and 40 CHF for which they would be indifferent between receiving the bid amount today versus 40 CHF after the specified delay on that trial. Participants made bids for immediate payoffs versus 40 CHF in 14 linearly spaced delays ranging from 13 to 181 days from the day of the experiment.

The rules of the BDM auction were fully explained to participants and were as follows. If the participant's bid was more than a randomly determined counter offer that was uniformly drawn from the range 0:40 CHF, she would receive the delayed payment of 40 CHF in the indicated number of days. However, if she had bid less than the counter offer, she was paid the amount of the counter offer at the day of the study instead of 40 CHF in the future. If the bid was equal to the counter offer, then a coin flip decided which payment was made. This mechanism ensures that it is in the participants' best interest to bid their true value for the present equivalent of 40 CHF at the given delay.

We calculated a discounting score for each participant as the area under the curve for all bids, where higher bids indicate a greater willingness to wait for the delayed outcome. We excluded the data from several participants from analysis because

their bidding patterns indicated that they did not understand the task (e.g., they alternated between bidding high and low amounts with increasing delay, or bid lower amounts for short delays and increased their bids for longer delays, which is the opposite of the expected pattern). The ITC data analyses we report are based on N = 135 participants (45 Anodal, 45 Cathodal, and 45 from the Sham group) who showed a consistent pattern of temporal discounting across trials.

Regression models used in the tDCS study

All regression models were estimated using Bayesian Markov-chain Monte Carlo (MCMC) methods with R ("R Core Team," 2015), in combination STAN¹⁰¹ or JAGS¹⁰². We used the default, uninformative priors specified by the brms¹⁰³ or BEST¹⁰⁴ R-packages.

We tested working memory, monetary intertemporal choice and SSRT task performance (Tables S10, S12, and S14) as a function of stimulation session and stimulation condition using linear regressions taking the form of Eq. S1:

$$(\text{Eq. S1}) \text{ Score} = \beta_0 + \beta_1 \text{ stimulationON} + \beta_2 \text{ Cathodal Baseline} + \beta_3 \text{ Anodal Baseline} + \beta_4 \text{ Cathodal*stimulationON} + \beta_5 \text{ Cathodal*stimulationON} + e$$

where *Score* denoted the respective working memory, intertemporal choice or stop-signal reaction time performance, *stimulationON* was a binary dummy variable denoting stimulation session (0 = pre-stimulation, 1 = concurrent or immediately after stimulation), and Anodal and Cathodal were the active stimulation conditions. The Sham stimulation condition served as the baseline in this regression.

Baseline differences in hunger levels (Table S16) were modelled according to Eq. S2:

$$(\text{Eq. S2}) \text{ hunger level} = \beta_0 + \beta_1 \text{ Cathodal} + \beta_2 \text{ Anodal} + e$$

where *hunger level* was given in percent (as indicated on a visual analogue scale with 0% not at all and 100% = maximally hungry). The Sham stimulation condition served as the baseline in this regression.

The population-level regressors for the for taste and health ratings regressions (Table S22) are listed in equation S3 below. Note that the models also included intercepts for subjects and food items (i.e., each subject and food item was treated as a random effect) as well as subject- and food-specific specific slopes for the effect of stimulation.

$$\text{(Eq. S3) Rating} = \beta_0 + \beta_1 \text{ stimulationON} + \beta_2 \text{ Cathodal} + \beta_3 \text{ Anodal} + \beta_4 \text{ stimulationON} * \text{Cathodal} + \beta_5 \text{ stimulationON} * \text{Anodal} + e$$

where *Rating* was either health or taste ratings for the foods (tested in separate models), *stimulationON* was a factor for stimulation session (pre-stimulation, post-stimulation), and Anodal and Cathodal were the active stimulation conditions. The Sham condition was the baseline in this regression.

Health goal statement used in the tDCS study

The statement read: “In this study, we want to investigate how people make healthy food choices. Therefore, we ask you to maintain the goal of eating as healthy as possible during this study. Specifically, we ask you to try and choose the healthier of the two food options on each trial. However, these are real decisions, and you are required to eat the food that you chose in one randomly selected trial. We realize this may be more difficult for some people than others, and it is important for us to know whether you agree to this goal or not. Your participation and payment are not contingent on your response. However, this is important for the scientific validity of our study, so please mark your answer below honestly. Please mark “yes” if you agree to do your best to follow the health goal. Please mark “no” if you do not want to commit yourself to the health goal.” Participants made their selections, dated and signed the document, and handed it back to the experimenter.

Comparing response times during rating sessions to relative start times

We used response times during the rating sessions as an estimate of the participants' fluency in recalling or constructing taste and healthiness attributes (Table S23). To test whether the relative start time (RST) depended on the speed of ratings for either health or taste aspects, we estimated the following model for each participant:

$$(Eq. S4) RST = \beta_0 + \beta_1 mRT_{taste} + \beta_2 mRT_{health} + \beta_3 nDT + e$$

where mRT is the mean reaction time over all taste ratings or health ratings that the participant made at the beginning of the experiment, and nDT is the non-decision time estimated in the tDDM. We conducted this analysis using the data from the baseline session in our tDCS experiment (i.e., our largest set of data from a single choice paradigm/context).

Testing the influence of all tDDM parameters on the improvement in fit for the tDDM relative to the standard DDM

We tested how the improvement in fit for the tDDM relative to the standard DDM (Table S18) relates to all the parameters of the tDDM by estimating the following Bayesian linear model:

$$(Eq. S5) diffLL = \beta_0 + \beta_1 |\omega_{taste}| + \beta_2 |\omega_{health}| + \beta_3 |RST| + \beta_4 Thr + \beta_5 nDT + \beta_6 bias + e$$

where diffLL is the difference in Log-Likelihood between the tDDM and the standard DDM, $|\omega_{taste/health}|$ is the absolute value of taste or the health weights, $|RST|$ is the absolute value of relative-start-time (to accommodate changes in any direction), Thr

is the threshold, and nDT is the non-decision-time estimated by the tDDM. We conducted this analysis using models fit to the choice and reaction time distributions across all participants (N = 272). Note that similar results are obtained when using a logistic regression that treats the difference in Log-Likelihood as a binary outcome.

Improvement of prediction accuracy for the tDDM relative to standard DDM on choices made by a tDDM generating process and by human participants

We tested how much the out-of-sample prediction accuracy increased for a time-varying relative to a standard DDM on choices made by human participants (Table S19) as well as on choices made by simulated agents that used a time-varying DDM to make decisions (Table S21). We modeled this improvement as a function of trial and individual-specific variables using the following Bayesian hierarchical linear regression:

$$(Eq. S6) \text{ pERR} = \beta_0 + \beta_1 |TD| + \beta_2 |HD| + \beta_3 \omega_{\text{taste}} + \beta_4 \omega_{\text{health}} + \beta_5 |RST| + \beta_6 \text{Thr} + \beta_7 \text{nDT} + \beta_8 \text{bias} + e$$

where pERR indicates the percent error reduction of the tDDM over the standard DDM in predicting out-of-sample choices, $|TD|$ and $|HD|$ are the absolute trial-wise health and taste value differences, and $\omega_{\text{taste/health}}$, Thr, nDT, bias, and $|RST|$ indicate the 6 tDDM parameters: taste and health weighting, threshold, non-decision time, bias, and the absolute value of the relative start time for health.

Testing for differences in the strength of association between relative attribute weights and starting times in generating versus recovered parameters

We tested whether there was a difference in the strength of association between relative attribute weights and starting times in recovered relative to known generating parameters. The recovered parameters were estimated from choices generated by the best fitting tDDM parameters across 272 participants (Table S20). We conducted this analysis using the following Bayesian linear model:

$$(Eq. S7) RST = \beta_0 + \beta_1 (\omega_{\text{taste}} - \omega_{\text{health}}) + \beta_2 \text{RecParameters} + \beta_3 (\omega_{\text{taste}} - \omega_{\text{health}}) * \text{RecParameters} + e$$

where RST refers to the relative starting times across generating and recovered parameters, $(\omega_{\text{taste}} - \omega_{\text{health}})$ are the relative attribute weights, and RecParameters is a dummy variable where 1 or 0 indicates that the parameter was recovered or generated, respectively.

Testing healthy food choices as a function of attribute weighting and timing parameters

We modeled healthy food choices as a function of the estimated parameters of the tDDM, the difference in taste and healthiness ratings, and whether the participants faced a challenge on the respective trial by fitting a Bayesian hierarchical logistic regression (Table S24). The population-level regressors are given below:

$$(Eq. S8) Hc = \beta_0 + \beta_1 \omega_{\text{taste}} + \beta_2 \omega_{\text{health}} + \beta_3 RST + \beta_4 \text{Chall} + \beta_5 TD + \beta_6 HD + \beta_7 \omega_{\text{taste}} * RST + \beta_8 \omega_{\text{health}} * RST + \beta_9 \omega_{\text{taste}} * \text{Chall} + \beta_{10} \omega_{\text{health}} * \text{Chall} + \beta_{11} RST * \text{Chall} + \beta_{12} \omega_{\text{taste}} * RST * \text{Chall} + \beta_{13} \omega_{\text{health}} * RST * \text{Chall} + e$$

where Hc is a binary variable indicating whether a healthy food item was chosen, $\omega_{\text{taste/health}}$ are the estimated taste or health weighting parameters of the tDDM, RST is the relative start time for health, Chall is a dummy variable indicating whether the participants faced a challenge on the respective trial, and TD and HD are the differences in taste and healthiness ratings between the two options, respectively. The model included subject-specific intercepts as well as subject-specific slopes for TD and HD to capture individual differences in sensitivity to these two aspects when making healthy choices.

Testing response times in health challenge and non-challenge trials as a function of the tDDM parameters

We used a hierarchical Bayesian linear regression to model response times in health challenge and non-challenge trials as a function of the tDDM parameters pooled

across the baseline conditions of all 4 studies (Table S25). The population-level regressors are indicated below:

$$\begin{aligned} \text{(Eq. S9)} \log(\text{RT}) = & \beta_0 + \beta_1 \text{Chall} + \beta_2 \omega_{\text{taste}} + \beta_3 \omega_{\text{health}} + \beta_4 \text{RST} + \beta_5 \text{HC} + \beta_6 |\text{TD}| + \beta_7 \\ & |\text{HD}| + \beta_8 \omega_{\text{taste}} * \text{RST} + \beta_9 \omega_{\text{health}} * \text{RST} + \beta_{10} \omega_{\text{taste}} * \text{HC} + \beta_{11} \omega_{\text{health}} * \text{HC} + \beta_{12} \text{RST} * \text{HC} + \beta_{13} \\ & \text{Chall} * \omega_{\text{taste}} + \beta_{14} \text{Chall} * \omega_{\text{health}} + \beta_{15} \text{Chall} * \text{RST} + \beta_{16} \text{Chall} * \text{HC} + \beta_{17} \text{Chall} * |\text{TD}| + \beta_{18} \\ & \text{Chall} * |\text{HD}| + \beta_{19} \omega_{\text{taste}} * \text{RST} * \text{HC} + \beta_{20} \omega_{\text{health}} * \text{RST} * \text{HC} + \beta_{21} \text{Chall} * \omega_{\text{taste}} * \text{RST} + \beta_{22} \\ & \text{Chall} * \omega_{\text{health}} * \text{RST} + \beta_{23} \text{Chall} * \omega_{\text{taste}} * \text{HC} + \beta_{24} \text{Chall} * \omega_{\text{health}} * \text{HC} + \beta_{25} \text{Chall} * \text{RST} * \text{HC} + \beta_{26} \\ & \text{Chall} * \omega_{\text{taste}} * \text{RST} * \text{HC} + \beta_{27} \text{Chall} * \omega_{\text{health}} * \text{RST} * \text{HC} + e \end{aligned}$$

where $\log(\text{RT})$ are log-transformed reaction times, Chall is a dummy variable indicating whether the participants faced a challenge on the respective trial, $\omega_{\text{taste/health}}$ are the estimated taste or health weighting parameters of the tDDM, HC is a binary variable indicating whether a healthy food item was chosen, and $|\text{TD}|$ and $|\text{HD}|$ are the absolute value differences in taste and healthiness ratings between the two options, respectively. The model included subject-specific intercepts and subject-specific slopes for the $|\text{TD}|$ and $|\text{HD}|$ regressors as well as the interactions of these regressors with the challenge trial and healthy choice dummy variables to capture individual differences in sensitivity to taste and health aspects when making healthy or unhealthy choices in challenging or non-challenging settings.

Modeling response times in cases where the healthier or tastier option was chosen

We tested how response times were affected by the difference in attribute ratings and whether participants faced a challenge on the respective trial separately for cases where the healthier (Table S26) or tastier option was chosen (Table S27) with the following regression:

$$\begin{aligned} \text{(Eq. S10)} \log(\text{RT}) = & \beta_0 + \beta_1 |\text{TD}| + \beta_2 |\text{HD}| + \beta_3 \text{Chall} + \beta_4 |\text{TD}| * \text{Chall} + \beta_5 |\text{HD}| * \text{Chall} \\ & + e \end{aligned}$$

where $\log(\text{RT})$ are log-transformed reaction times, $|\text{TD}|$ and $|\text{HD}|$ are the absolute value differences in taste and healthiness ratings between the two options,

respectively, and Chall is a dummy variable indicating whether the participants faced a challenge on the respective trial. The model included subject-specific intercepts and subject-specific slopes for the |TD| and |HD| regressors and their interaction with the challenge trial dummy variable to capture individual differences in sensitivity to taste and health aspects in challenging or non-challenging settings.

Modelling response times as a function of subjective value when one or two attributes are considered

To test whether response times also depend on whether or not the subjective value derived from the initially considered attribute alone is compatible (i.e. points to the same choice) with a subjective value computed by integrating across both taste and healthiness attributes (Table S28), we modeled response times with the following regression:

$$(Eq. S11) \log(RT) = \beta_0 + \beta_1 \text{Comp} + \beta_2 |TD| + \beta_3 |HD| + e$$

where log(RT) are log-transformed reaction times, Comp is a dummy variable indicating whether the subjective value on the respective trial for the initially considered attribute was compatible with the subjective value of both health and taste attributes combined, and |TD| and |HD| are the absolute value differences in taste and healthiness ratings between the two options, respectively. The model included subject-specific intercepts as well as subject-specific slopes for |TD| and |HD| to capture individual differences in sensitivity to these two attributes.

Statistics packages used

All analyses presented in this paper were performed with the R (“R Core Team,” 2015), STAN¹⁰¹ and JAGS¹⁰² statistical software packages. The sequential sampling model was fit using the Rcpp toolbox¹⁰⁵. Bayesian regression models were run using the brms package¹⁰³ which is an interface between R and STAN. All correlations and t-tests were computed using Bayesian MCMC sampling methods using R in combination with JAGS^{104,106}. The plots in Figures 2,3 and 5 were created using the yarr package¹⁰⁷. The packages pracma¹⁰⁸ and plyr¹⁰⁹ were used for data handling and restructuring.

4. Supplemental Tables

Table S1. Bayesian information criterion (BIC) values for alternative specifications of the DDM.

DDM type	Squared Error	Squared Error Difference	BIC	BIC difference
Different Onsets	0.204	-	280632	-
Standard DDM	0.218	0.014	286493	5861
Different Offsets	0.215	0.011	318724	38092
Different Onsets & Offsets	0.219	0.015	289060	8428
Starting-point bias for tastier	0.281	0.077	281909	1277

This table lists the squared error in out-of-sample predictions and the Bayesian information criterion (BIC) values for the time-varying DDM used throughout the main text, the standard DDM, and three other alternative DDM formulations. Column 1 lists the squared error in predictions for even-numbered trial outcomes generated from parameter fits to participants odd-numbered choices. Column 2 lists the increase in error for a given model relative to the tDDM with separate consideration onset times for taste and healthiness attributes. Thus, higher values indicate better performance for our primary tDDM. Column 3 lists the BIC values for each model across all trials from all 272 participants in the four data sets we examine here. Column 4 lists the differences between each alternative and our primary tDDM specification (i.e. alternative BIC – primary BIC). Here again, higher values indicate better performance for our primary tDDM.

Different Onsets = a time-varying DDM that includes a free parameter allowing for the consideration onset time of each attribute to differ. This is the model we use throughout the rest of the paper.

Different Offsets = a DDM that instead of different onsets allows for the time at which each attribute ceases to be considered to differ. In other words, one attribute could stop adding to the evidence accumulation process before a threshold is reached.

Different Onsets & Offsets = a DDM including both separate onset and offset times for each attribute.

Starting-point-bias for tastier = a standard DDM with simultaneous attribute on and offset times. However, in this case the data have been transformed such that tastier outcomes are always on the upper boundary of the DDM. This allows the starting point bias that formally measured left/right or eat/don't eat biases in our other formulations of the DDM to capture a bias in favour of the tastier option.

Standard = a standard DDM with boundaries and starting-point biases for left/right or eat/don't outcomes.

In every case, we fit the DDMs to either all or the odd-numbered choices in one experimental condition from all 272 participants. In other words, there are no repeated within-subjects measures in these comparisons. All DDMs were fit using the same procedures as those described in the main-text Methods section.

Table S2. The influence of tastiness relative to healthiness before and after the average initial advantage for the attribute considered first has elapsed.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	-0.04 $\pm$ 0.05	[-0.14; 0.06]
HFirst	0.02 $\pm$ 0.06	[-0.09; 0.13]
InitAdv	-0.09 $\pm$ 0.13	[-0.34; 0.17]
TD	1.34 $\pm$ 0.06	[1.21; 1.47]
HD	0.65 $\pm$ 0.07	[0.52; 0.79]
HFirst * InitAdv	-0.35 $\pm$ 0.17	[-0.69; -0.04]
HFirst * TD	-0.57 $\pm$ 0.07	[-0.7; -0.43]
HFirst * HD	0.61 $\pm$ 0.07	[0.47; 0.75]
InitAdv * TD	0.78 $\pm$ 0.14	[0.51; 1.05]
InitAdv * HD	0.33 $\pm$ 0.12	[0.1; 0.57]
HFirst * InitAdv * TD	-0.52 $\pm$ 0.15	[-0.81; -0.25]
HFirst * InitAdv * HD	0.81 $\pm$ 0.15	[0.52; 1.12]
Bayesian R ²	0.4	

This table reports the group-level results from the Bayesian hierarchical logistic regression described in Eq. 2 for all 272 participants. The outcome variable in this equation was a binary indicator of the choice outcome (left/right, eat/do not eat). HFirst is a dummy variable (1 = healthiness, 0 = taste) indicating which attribute is considered first – as determined by the tDDM. InitAdv is a dummy variable (1 = before, 0 = after) indicating whether the response was made before the median value of the sum of relative-starting-time difference plus non-decision time across participants had elapsed. This sum was equal to 1 second. The abbreviations TD and HD stand for the differences in tastiness and healthiness, respectively, on each trial. These differences were computed as left item – right item, or item to potentially eat – 0. Subject-specific coefficients were estimated for all regressors except HFirst, because each participant had only one level of that regressor in his or her in baseline condition. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval.

Table S3. Model comparisons between a Bayesian hierarchical logistic regression including response time and consideration order effects and two simplified models without those features.

Model	LOOIC	SE
Full Model	22873.33	161.16
Reduced Model 1	23168.74	161.25
Reduced Model 2	23342.44	159.83
Full – Reduced 1	-295.41	33.46
Full – Reduced 2	-469.12	43.48

This table shows the expected out-of-sample prediction error (LOOIC) for three Bayesian hierarchical logistic regressions that measure the relative influence of health and tastiness attributes on choice outcomes. The LOOIC values were computed using Pareto smoothed importance-sampling leave-one-out cross-validation^{110,111}. Smaller values indicate better expected out-of-sample performance. The standard errors (SE) for the LOOIC or differences in LOOIC are shown in the last column.

The full model is given by Eq. (2) that we repeat here for convenience:

$$\text{Left} = \beta_0 + \beta_1 \text{HFirst} + \beta_2 \text{InitAdv} + \beta_3 \text{TD} + \beta_4 \text{HD} + \beta_5 \text{HFirst} * \text{InitAdv} + \beta_6 \text{HFirst} * \text{TD} + \beta_7 \text{HFirst} * \text{HD} + \beta_8 \text{InitAdv} * \text{TD} + \beta_9 \text{InitAdv} * \text{HD} + \beta_{10} \text{HFirst} * \text{InitAdv} * \text{TD} + \beta_{11} \text{HFirst} * \text{InitAdv} * \text{HD} + e$$

In this equation, Left is a binary indicator of the choice outcome. HFirst is a dummy variable (1 = healthiness, 0 = taste) indicating which attribute is considered first – as determined by the tDDM. InitAdv is a dummy variable (1 = before, 0 = after) indicating whether the response was made before the median value of the sum of relative-starting-time difference plus non-decision time across participants had elapsed. This sum was equal to 1 second. The abbreviations TD and HD stand for the differences in tastiness and healthiness, respectively, on each trial.

We compared this full model to two simpler models that omitted either, 1) the dependency on response time (i.e. InitAdv dummy variable) or 2) both the dependency on response time and the indicator for which attribute a participant considered first (i.e. InitAdv and HFirst dummy variables). The full model is better than both reduced models and the results from it are given in Table S2.

Table S4. The influence of taste relative to healthiness before and after the average initial advantage for the attribute considered first has elapsed in simulated choices.

	tDDM simulations		Empirical data		DDM simulations	
	Mean $\pm$ SD	95% CI	Mean $\pm$ SD	95% CI	Mean $\pm$ SD	95% CI
(Intercept)	0.07 $\pm$ 0.10	[-0.11; 0.27]	0.12 $\pm$ 0.12	[-0.12; 0.36]	0.12 $\pm$ 0.09	[-0.06; 0.31]
HFirst	-0.22 $\pm$ 0.13	[-0.47; 0.04]	-0.29 $\pm$ 0.16	[-0.61; 0.02]	-0.24 $\pm$ 0.13	[-0.49; 0.00]
InitAdv	1.67 $\pm$ 0.77	[0.21; 3.24]	0.87 $\pm$ 0.31	[0.29; 1.51]	1.03 $\pm$ 0.97	[-0.78; 3.07]
TD	1.75 $\pm$ 0.19	[1.38; 2.12]	1.33 $\pm$ 0.14	[1.06; 1.60]	2.41 $\pm$ 0.27	[1.87; 2.94]
HD	0.45 $\pm$ 0.25	[-0.03; 0.93]	0.29 $\pm$ 0.17	[-0.05; 0.63]	0.43 $\pm$ 0.30	[-0.15; 1.03]
HFirst * InitAdv	-3.62 $\pm$ 1.01	[-5.81; -1.81]	-1.73 $\pm$ 0.42	[-2.59; -0.95]	-2.59 $\pm$ 1.21	[-5.22; -0.39]
HFirst * TD	-0.77 $\pm$ 0.24	[-1.24; -0.29]	-0.50 $\pm$ 0.18	[-0.86; -0.15]	-1.10 $\pm$ 0.34	[-1.76; -0.44]
HFirst * HD	1.28 $\pm$ 0.33	[0.65; 1.95]	0.94 $\pm$ 0.23	[0.50; 1.38]	2.04 $\pm$ 0.41	[1.27; 2.84]
InitAdv * TD	0.37 $\pm$ 0.36	[-0.29; 1.13]	1.09 $\pm$ 0.30	[0.56; 1.73]	0.69 $\pm$ 0.67	[-0.39; 2.24]
InitAdv * HD	-0.62 $\pm$ 0.31	[-1.29; -0.08]	-0.06 $\pm$ 0.24	[-0.53; 0.42]	-0.25 $\pm$ 0.47	[-1.20; 0.65]
HFirst * InitAdv * TD	-1.39 $\pm$ 0.48	[-2.38; -0.51]	-1.31 $\pm$ 0.33	[-2.01; -0.70]	-0.34 $\pm$ 0.68	[-1.87; 0.83]
HFirst * InitAdv * HD	1.31 $\pm$ 0.50	[0.43; 2.41]	1.04 $\pm$ 0.39	[0.37; 1.85]	0.74 $\pm$ 0.78	[-0.58; 2.52]

Table S4 reports the results of the same Bayesian hierarchical logistic regression used in Table S3 (Eq. 2), but fit to simulated choices derived from the best-fitting tDDM (left columns) or standard DDM (right columns) parameters. The simulations were based on parameters fit to choices from the subset of participants whose |RST| duration fell in the 3rd quantile across all four studies (approximately 0.3-0.5 s). Thus, this subset includes participants whose RST is large enough to produce clear effects on simulated choice outcomes, but not at the extreme ends of the distribution. The regression coefficients from the hierarchical logistic regression on the actual choices made by this subset of participants are shown in the centre columns.

Table S5. The variability in relative-start-times is only modestly related to other tDDM parameters.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.19 $\pm$ 0.24	[-0.29; 0.65]
ω_{taste}	0.15 $\pm$ 0.04	[0.07; 0.22]
ω_{health}	0.01 $\pm$ 0.04	[-0.07; 0.08]
nDT	0 $\pm$ 0.19	[-0.37; 0.38]
Thr	-0.04 $\pm$ 0.11	[-0.26; 0.18]
Bias	-0.27 $\pm$ 0.82	[-1.90; 1.29]
Study IAC	0.07 $\pm$ 0.10	[-0.12; 0.27]
Study MRT	-0.10 $\pm$ 0.11	[-0.32; 0.11]
Study TDCS	-0.32 $\pm$ 0.09	[-0.49; -0.15]
Bias X Study IAC	0.39 $\pm$ 0.85	[-1.22; 2.07]
Bias X Study MRT	-0.18 $\pm$ 0.92	[-1.92; 1.65]
Bias X Study TDCS	0.74 $\pm$ 0.83	[-0.86; 2.38]
Bayesian R ²	0.3	

This table reports the results from the Bayesian linear regression described in Eq. 3. We modelled the relative-start-time as a function of all other parameters in the tDDM fits to the “baseline” choice condition in each study. These “baseline” condition fits correspond to those shown in Figure 3 and include data from the mouse response trials in Sullivan et al 2015, data from the natural choice condition in Hare et al 2011, data from the first session/day of the gambles and food choices study, and data from the pre-stimulation baseline choices in our tDCS experiment. Overall, the other tDDM parameters explained only 30% of the variability in relative-start-times across participants. To account for the fact that the bias measured different answers across studies, and thus, carries a different meaning (e.g., left/right, eat/do not eat), we included an interaction between the *Bias* parameter from the tDDM with a dummy variable for each *Study* to account for the different task designs. The significant coefficient for the Study TDCS dummy regressor indicates that relative-start-times favoured healthiness more in that study relative to the GFC study that served as the baseline in this regression. This fact can be seen in Figure 3 as well.

Table S6. Shared variance between relative start-times and drift weights.

	ω_{taste}	ω_{health}	$(\omega_{\text{taste}} - \omega_{\text{health}})$	$(\omega_{\text{taste}} + \omega_{\text{health}})$
RST versus:	12 % (0.351)	12 % (-0.348)	20 % (0.448)	< 1% (-0.030)

The values shown in this table represent the amount of shared variance (Pearson's $r^2 * 100$) between the relative-start-time parameter (RST) and the separate, relative ($\omega_{\text{taste}} - \omega_{\text{health}}$) or summed ($\omega_{\text{taste}} + \omega_{\text{health}}$) drift rates across all participants and conditions ($n = 549$ parameter sets). The signed Pearson's r -values are shown in parenthesis to indicate the direction of each correlation. The estimated ω_{taste} and ω_{health} parameters share 6% of their variance between themselves ($r = -0.213$).

Table S7. Logistic regression testing whether stimulation over left DIPFC causes changes in health challenge success.

Model parameter	Mean beta $\pm$ SD	95% Credible Interval
<i>Main effect on health challenge success</i>		
(Intercept)	-0.04 $\pm$ 0.12	[-0.28; 0.21]
TD	-0.61 $\pm$ 0.11	[-0.83; -0.40]
HD	0.72 $\pm$ 0.11	[0.51; 0.94]
stimulationON	0.03 $\pm$ 0.11	[-0.18; 0.24]
Cathodal	-0.10 $\pm$ 0.17	[-0.44; 0.24]
Anodal	-0.02 $\pm$ 0.18	[-0.37; 0.32]
TD X stimulationON	-0.02 $\pm$ 0.11	[-0.23; 0.20]
HD X stimulationON	0.24 $\pm$ 0.11	[0.04; 0.45]
TD X Cathodal	0.08 $\pm$ 0.16	[-0.22; 0.39]
TD X Anodal	0.10 $\pm$ 0.16	[-0.20; 0.41]
HD X Cathodal	-0.01 $\pm$ 0.14	[-0.28; 0.27]
HD X Anodal	0.14 $\pm$ 0.14	[-0.15; 0.42]
stimulationON X Cathodal	-0.32 $\pm$ 0.15	[-0.63; -0.03]
stimulationON X Anodal	-0.03 $\pm$ 0.15	[-0.33; 0.28]
TD X stimulationON X Cathodal	-0.18 $\pm$ 0.15	[-0.49; 0.12]
TD X stimulationON X Anodal	-0.06 $\pm$ 0.15	[-0.35; 0.23]
HD X stimulationON X Cathodal	0.11 $\pm$ 0.13	[-0.15; 0.39]
HD X stimulationON X Anodal	0.00 $\pm$ 0.14	[-0.26; 0.27]
Bayesian R ²	0.18	

This table reports the results from a Bayesian hierarchical logistic regression model (Eq. 4) explaining trial-wise health challenge success as a function of the categorical factor tDCS polarity (Anodal, Cathodal, and Sham), a dummy variable for session (stimulation=1, baseline=0), and continuous regressors for the standardized (z-score) absolute value of the healthiness difference (HD) and taste difference (TD) on each trial. The stimulationON regressor itself measures the difference between the initial baseline choice session and the second round of choices for the sham stimulation group (i.e. it captures potential effects of time and/or experience). Thus, the HD x stimulationON interaction indicates that the influence of healthiness on choice outcomes changes over time and/or experience with the food choice task. The effects of active tDCS are captured by the interaction coefficients, “stimulationON X Cathodal” and “stimulation X Anodal”. The model included subject-specific intercepts as well as subject-specific slopes for the effects of stimulation versus baseline (stimulationON), HD, and TD. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval.

Table S8. Demographics of tDCS stimulation groups.

Group	Sham	Cathodal	Anodal
BMI	22.37 ± 2.75	22.30 ± 2.13	22.74 ± 2.92
Age (in years)	23.54 ± 2.90	23.75 ± 3.25	23.48 ± 23.04
TFEQ – Restrained Eating	9.94 ± 2.21	9.34 ± 2.46	9.59 ± 2.62
TFEQ - Disinhibition	7.52 ± 2.04	8 ± 2.17	7.70 ± 2.35
TFEQ – Hunger susceptibility	6.22 ± 2.11	6.43 ± 2.20	6.05 ± 2.26

This table reports descriptive statistics for Age, Body Mass Index (BMI) and scores on each subscale of the Three Factor Eating Questionnaire (TFEQ, German validated version by Pudel and Westenhöfer¹¹²). All values are given as mean and SD. Linear regression models testing for differences in each of these measures across stimulation groups did not reveal significant differences for any measure.

Table S9. Mean and standard deviation for working memory by tDCS stimulation group (pre and post stimulation).

Group	Mean pre	SD pre	Mean post	SD post
Anodal	7.39	1.45	7.84	1.53
Cathodal	7.18	1.23	7.61	1.29
Sham	7.12	1.29	7.38	1.31

The values in this table represent the forward digit span that we measured before and after stimulation using a computerized task that asked participants to remember sequences of 5 to 12 digits. Sequences were pseudo-randomized so that unique sequences would be presented before and after stimulation.

Table S10. Comparison of working memory at baseline and after tDCS stimulation across stimulation groups.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	7.12 $\pm$ 0.18	[6.77; 7.47]
Main Effect of stimulationON	0.26 $\pm$ 0.25	[-0.23; 0.73]
Cathodal Baseline	0.05 $\pm$ 0.25	[-0.45; 0.55]
Anodal Baseline	0.27 $\pm$ 0.25	[-0.23; 0.77]
Cathodal X stimulationON	0.18 $\pm$ 0.36	[-0.52; 0.86]
Anodal X stimulationON	0.19 $\pm$ 0.36	[-0.50; 0.91]

This table reports the results of a Bayesian regression model (given in Eq. S1) testing for differences in working memory (forward digit span; measured as explained in caption S10) as a function of tDCS polarity and stimulation session (0 = pre-stimulation and 1 = concurrent or immediately after stimulation). The Sham condition served as the baseline for this regression. There were no significant main effects or interactions. All regression coefficients are reported as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval.

Table S11. Mean and standard deviation for stop-signal-reaction-time (SSRT) in seconds by stimulation group (pre and post stimulation).

Group	Mean pre	SD pre	Mean post	SD post
Anodal	0.20	0.09	0.19	0.08
Cathodal	0.20	0.09	0.19	0.08
Sham	0.21	0.08	0.19	0.09

The values in this table represent the stop-signal-reaction-time (SSRT) in seconds that we measured before and after stimulation. A computerized task asked participants to respond as quickly as possible once a stimulus appeared on the screen but to withhold their response movement if a second “stop” stimulus appeared as well. A higher SSRT indicates better ability to stop initiated movements even at a later stage.

Table S12. Between-stimulation group comparison of stop signal reaction time (SSRT) at baseline and after stimulation.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.21 $\pm$ 0.01	[0.18; 0.24]
Main Effect of stimulationON	-0.02 $\pm$ 0.02	[-0.06; 0.02]
Cathodal Baseline	-0.01 $\pm$ 0.02	[-0.05; 0.03]
Anodal Baseline	-0.01 $\pm$ 0.02	[-0.05; 0.03]
Cathodal X stimulationON	0.01 $\pm$ 0.03	[-0.04; 0.06]
Anodal X stimulationON	0 $\pm$ 0.03	[-0.05; 0.06]

This table reports the results of a Bayesian regression model (given in Eq. S1) testing for differences in SSRT (measured as explained in the caption of S12) as a function of tDCS polarity and stimulation session (0 = pre-stimulation and 1 = concurrent or immediately after stimulation). The Sham condition served as the baseline for this regression. There were no significant main effects or interactions. All regression coefficients are reported as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval.

Table S13. Mean and standard deviation of intertemporal preference measures by stimulation group (pre and post stimulation).

Group	Mean pre	SD pre	Mean post	SD post
Anodal	5123	1361	5214	1372
Cathodal	4628	1285	4583	1414
Sham	5144	1488	5123	1523

The values in this table represent the area under the curve (AUC) for bids for all 14 delays in an intertemporal choice task that we measured before and after stimulation. A computerized task asked participants to bid in a Becker-DeGroot-Marschak (BDM) auction their true current value of 40 CHF in X days (with X being 14 linearly increasing delays from 12 to 183 days). A higher AUC indicates less temporal discounting because bids in the BDM auction represented the present discounted value of the delayed rewards. Table S12 presents the results of a regression testing for significant differences in temporal discounting between stimulation groups and choice sessions (i.e., pre and post stimulation).

Table S14. Between-stimulation group comparison of monetary intertemporal choice at baseline and after stimulation.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	5149.06 $\pm$ 206.47	[4747.41; 5563.95]
Main Effect of stimulationON	-25.94 $\pm$ 295.23	[-620.27; 551.19]
Cathodal Baseline	-519.78 $\pm$ 292.98	[-1088.85; 54.17]
Anodal Baseline	-25.23 $\pm$ 297.86	[-610.24; 554.62]
Cathodal X stimulationON	-17.86 $\pm$ 417.47	[-847.88; 815.47]
Anodal X stimulationON	114.77 $\pm$ 420.16	[-710.31; 937.52]

This table reports the results of a Bayesian regression model (given in Eq. S1) testing for differences in temporal discounting (measured as AUC for all delays, as explained in the caption of table S11) as a function of tDCS polarity and stimulation session (0 = pre-stimulation and 1 = concurrent or immediately after stimulation). The Sham condition served as the baseline for this regression. There were no significant main effects or interactions. All regression coefficients are reported as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval.

Table S15. Hunger levels by tDCS stimulation group.

Group	Mean	SD
Anodal	74 %	20 %
Cathodal	76 %	22 %
Sham	75 %	20 %

The values in this table represent the subjective hunger level that we measured before stimulation. Participants reported their current feelings of hunger on a computerized visual analogue scale with anchors from -5 ("not at all") to +5 ("extremely"). Values were transformed to percent of the total scale length.

Table S16. Baseline hunger level comparison between tDCS stimulation groups.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	74.6 $\pm$ 2.73	[69.17; 80.01]
Cathodal	1.05 $\pm$ 3.90	[-6.86; 8.57]
Anodal	-0.56 $\pm$ 3.93	[-8.29; 7.1]

This table reports the results of a Bayesian regression model (given in Eq. S2) testing for differences in baseline hunger levels (measured as explained in the caption of table S13) as a function of tDCS group. There were no significant main effects. The Sham condition served as the baseline for this regression. All regression coefficients are reported as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval. Hunger level is reported as a percentage (100% = maximally hungry).

Table S17. Upper and lower bounds for each tDDM parameter during model fitting.

Parameter	Lower bound	Upper bound
ω_{taste}	-2	2
ω_{health}	-2	2
Thr	0.6	3
nDT	0.01	1
RST	-1	1
bias	-1	1

ω_{taste} : weighting factor determining how much taste contributes to the evidence accumulation rate.

ω_{health} : weighting factor determining how much healthiness contributes to the evidence accumulation rate.

Thr: evidence threshold for responding (symmetric around zero)

nDT: non-decision time

RST: relative-start-time for health (positive values mean that health enters the process after taste, negative values mean health enters before taste)

Bias: starting point bias for the evidence accumulation process

Table S18. The improvement in the fit for the tDDM relative to the standard DDM is proportional to increases in weighting parameters and relative-starting-times.

Model parameter	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	1.02 $\pm$ 0.59	[-0.12; 2.14]
$ \omega_{\text{taste}} $	-0.45 $\pm$ 0.15	[-0.74; -0.17]
$ \omega_{\text{health}} $	-0.62 $\pm$ 0.13	[-0.87; -0.36]
 RST 	-1.09 $\pm$ 0.20	[-1.49; -0.69]
Thr	-0.27 $\pm$ 0.30	[-0.85; 0.31]
nDT	0.02 $\pm$ 0.48	[-0.94; 0.97]
bias	-0.47 $\pm$ 0.28	[-1.02; 0.06]

Results from a Bayesian linear model (given in Eq. S5) explaining the difference in Log-Likelihood between the tDDM and the standard DDM by all parameters of the tDDM. In addition to threshold (Thr), non-decision time (nDT) and bias, absolute values of taste and health weights as well as relative-start-time (RST) were used in order to accommodate changes of these parameters in either direction. The standard DDM was fit using the same procedures as the tDDM, except that the equation omitted the relative-start-time parameter, which meant that both tastiness and healthiness entered into consideration at the same time. Note that similar results are obtained when using a logistic regression that treats the difference in Log-Likelihood as a binary outcome.

Table S19. Improvement in out-of-sample prediction accuracy for the time-varying relative to the standard DDM on choices made by human participants.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.86 $\pm$ 0.42	[0.01; 1.65]
TD	-0.32 $\pm$ 0.18	[-0.66; 0.04]
 HD 	-0.87 $\pm$ 0.17	[-1.22; -0.53]
ω_{taste}	-0.07 $\pm$ 0.61	[-1.26; 1.1]
ω_{health}	-0.69 $\pm$ 0.6	[-1.92; 0.48]
RST	0.85 $\pm$ 1.34	[-1.75; 3.43]
Thr	-8.17 $\pm$ 1.98	[-12.07; -4.27]
nDT	-5.18 $\pm$ 2.97	[-11.05; 0.78]
bias	0.61 $\pm$ 1.84	[-2.87; 4.3]

This table reports the results of a Bayesian hierarchical linear regression (Eq. S6) testing how trial (Taste and Healthiness Difference) and individual-specific (tDDM parameters) features relate to the relative accuracy of out-of-sample predictions made by the standard versus time-varying DDM on our 272 human participants even-numbered choices. The two trial-specific features and the six tDDM parameters (Taste weight, Healthiness weight, Relative-start-time (RST), threshold (Thr), non-decision time (nDT) and starting point bias) were mean centered before being entered into the regression. The taste and healthiness differences were included as absolute rather than signed values because choice difficulty was determined by the absolute difference in those attributes. Similarly, the regression used the absolute value of the RST parameter because the tDDM differs more strongly from the standard DDM the further this parameter value is from zero. The positive intercept coefficient indicates that the tDDM made more accurate out-of-sample choice predictions than the standard DDM on average. The relative accuracy of the tDDM when there was a larger difference in healthiness ratings between the two food items or participants were estimated to have a higher decision threshold. The difference in prediction accuracy between the two diffusion models was computed as the square-root of the squared error for the standard minus the tDDM on each trial multiplied by 100 to put it in terms of percent error. Thus, positive values indicate higher prediction accuracy for the tDDM. The set of generating-parameter agents used the best-fitting tDDM parameters for each of the 272 human participants to make food choices. The standard and tDDM parameters were fit to the odd-numbered trials for all participants and then those fits were used to make predictions about the most probable outcomes on even-numbered trials (see Methods for details). This regression included subject-specific intercepts.

Abbreviations:

ω_{taste} : weighting factor determining how much taste contributes to the evidence accumulation rate.

ω_{health} : weighting factor determining how much healthiness contributes to the evidence accumulation rate.

RST: relative-start-time for health

Table S20. A Bayesian linear regression (Eq. S7) testing for potential differences in the strength of the association between the relative attribute weights and starting times between generating and recovered parameters.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.02 $\pm$ 0.03	[-0.04; 0.07]
Relative weight	0.15 $\pm$ 0.02	[0.11; 0.19]
Recovered parameters	-0.03 $\pm$ 0.04	[-0.10; 0.04]
Rel. weight* Rec. parameters	-0.02 $\pm$ 0.03	[-0.07; 0.04]
Bayesian R ²	0.17	

This table shows that there is no significant difference in the strength of the association between relative attribute weights ($\omega_{\text{taste}} - \omega_{\text{health}}$) and the relative starting times in recovered parameters estimated from choices generated by the best fitting tDDM parameters. In other words, the coefficient for the Relative weight* Recovered parameters interaction term is not different from zero. These best fitting parameters were initially estimated for each of our 272 participants based on the choices they made during the experiments. We used these estimated parameters to generate choices and response times for the same decisions each participant originally faced. We then fit the tDDM to these simulated choices to see how well we could recover the known generating parameters (see Fig. S2). The Bayesian linear regression coefficient estimates (betas) are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval. There is a significant main effect for relative weights, indicating a positive relationship between relative weights and start times. Importantly, neither the main effect, Recovered parameters, nor the interaction, Relative weights * Recovered parameters, are significant. These results indicate that the recovered values for the relative-start-time (RST) parameter and the correlation between RST and relative attribute weights do not significantly differ in the recovered compared to the known generating parameters.

Table S21. Improvement in out-of-sample prediction accuracy for the time-varying relative to the standard DDM on choices made by simulated agents using a time-varying DDM to make decisions.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	1.92 $\pm$ 0.18	[1.57; 2.25]
TD	-0.04 $\pm$ 0.02	[-0.08; 0.01]
HD	-0.05 $\pm$ 0.03	[-0.1; 0]
ω_{taste}	0.06 $\pm$ 0.26	[-0.47; 0.55]
ω_{health}	-0.09 $\pm$ 0.24	[-0.58; 0.38]
 RST 	2.64 $\pm$ 0.51	[1.62; 3.62]
Threshold	-0.57 $\pm$ 0.76	[-2.03; 0.93]
Non-decision Time	-0.57 $\pm$ 1.32	[-3.03; 2.09]
Starting Point Bias	1.55 $\pm$ 0.76	[0.14; 3.08]

This table reports the results of a Bayesian hierarchical linear regression (Eq. S6) testing how trial (Taste and Healthiness Difference) and individual-specific (tDDM parameters) features relate to the relative accuracy of out-of-sample predictions made by the standard versus time-varying DDM when the true choice generating model was one that included separate consideration onset times for taste and healthiness attributes (tDDM). The two trial-specific features and the six tDDM parameters (Taste weight, Healthiness weight, Relative-start-time (RST), Threshold, Non-decision time, and Starting point bias) were mean centered before being entered into the regression. The taste (TD) and healthiness (HD) differences were included as absolute rather than signed values because choice difficulty was determined by the absolute difference in those attributes. Similarly, the regression used the absolute value of the RST parameter because the tDDM differs more strongly from the standard DDM the further this parameter value is from zero. The positive intercept coefficient indicates that the tDDM made more accurate out-of-sample choice predictions than the standard DDM on average. The relative accuracy of the tDDM was also greater when the generating-parameter agent had a bigger difference in consideration start times for tastiness versus healthiness (i.e. larger absolute RST). The difference in prediction accuracy between the two diffusion models was computed as the square-root of the squared error for the standard minus the tDDM on each trial multiplied by 100 to put it in terms of percent error. Thus, positive values indicate higher prediction accuracy for the tDDM. The set of generating-parameter agents used the best-fitting tDDM parameters for each of the 272 human participants to make food choices. The standard and tDDM recovered-parameter agents used the best parameters found when each model was fit to a separate set of choices made by the generating-parameter agents (see Methods for details). This regression included agent/subject-specific intercepts.

Abbreviations:

TD: taste difference

HD: healthiness difference

ω_{taste} : weighting factor determining how much taste contributes to the evidence accumulation rate.

ω_{health} : weighting factor determining how much healthiness contributes to the evidence accumulation rate.

RST: relative-start-time for health

Table S22. Tests of differences in taste or health ratings across stimulation groups at baseline or following tDCS

	Mean beta $\pm$ SD	95% Credible Interval
<i>A. Dependent variable = Health ratings</i>		
(Intercept)	49.73 $\pm$ 2.03	[45.61; 53.68]
stimulationON	1.03 $\pm$ 0.49	[0.07; 2.00]
Cathodal	0.74 $\pm$ 1.29	[-1.77; 3.16]
Anodal	-1.18 $\pm$ 1.33	[-3.65; 1.36]
Cathodal X stimulationON	-0.21 $\pm$ 0.69	[-1.58; 1.13]
Anodal X stimulationON	-0.29 $\pm$ 0.69	[-1.67; 1.11]
<i>B. Dependent variable = Taste ratings</i>		
(Intercept)	52.85 $\pm$ 1.51	[49.93; 55.80]
stimulationON	3.70 $\pm$ 0.95	[1.86; 5.57]
Cathodal	4.18 $\pm$ 1.92	[0.42; 7.91]
Anodal	4.92 $\pm$ 1.92	[1.14; 8.71]
Cathodal X stimulationON	-0.16 $\pm$ 1.31	[-2.75; 2.41]
Anodal X stimulationON	-1.07 $\pm$ 1.31	[-3.69; 1.49]

This table reports the results from two hierarchical Bayesian regressions (given in Eq. S3) testing for differences in taste or health ratings across stimulation groups at baseline or following tDCS. Taste and health ratings were coded as percent maximum tastiness or healthiness. Both regressions included regressors for the main effect of stimulation group (sham, anodal, cathodal), stimulation session (baseline, post-stimulation), and their interaction as population-level effects. The models also included intercepts for subjects and food items (i.e., each subject and food item was treated as a random effect) as well as subject-specific and food-specific slopes for the effect of stimulation versus baseline. **A)** Health ratings increased at the second rating during or after stimulation, but no group differences were observed at baseline or following stimulation. **B)** Taste ratings were higher at baseline for the anodal and cathodal groups compared to the sham group and increased for all groups at the second rating. Critically, there was no stimulation group X stimulation session interaction, indicating that stimulation condition did not differentially influence the ratings at the second stimulation session.

Table S23. Effects of the average taste and health rating response times on the relative-start-time for health.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	-0.63 $\pm$ 0.25	[-1.12; -0.13]
Mean RT health	0.05 $\pm$ 0.11	[-0.17; 0.28]
Mean RT taste	0.07 $\pm$ 0.12	[-0.17; 0.31]
nDT	0.33 $\pm$ 0.24	[-0.13; 0.80]

This table reports the results from a Bayesian regression (see Eq. S4) of the tDDM-estimated relative-start-time for health in baseline choices on the rating speed for taste and health, in the sample of participants from the TDCS study. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval. There was no significant relation between the mean reaction time (RT) for health and taste ratings (in seconds) and the relative-start-time for health in the individual choice process.

Table S24. Bayesian logistic regression for healthy choice outcomes as a function of attribute weighting and timing parameters.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	-0.17 $\pm$ 0.26	[-0.68; 0.35]
ω_{taste}	0.52 $\pm$ 0.17	[0.18; 0.85]
ω_{health}	1.43 $\pm$ 0.21	[1.02; 1.87]
RST	0.72 $\pm$ 0.36	[0.01; 1.43]
Challenge	-0.73 $\pm$ 0.25	[-1.22; -0.25]
Taste Difference (TD)	0.19 $\pm$ 0.03	[0.14; 0.24]
Health Difference (HD)	0.69 $\pm$ 0.08	[0.54; 0.84]
ω_{taste} X RST	-0.29 $\pm$ 0.27	[-0.81; 0.23]
ω_{health} X RST	-1.34 $\pm$ 0.26	[-1.86; -0.85]
ω_{taste} X Challenge	-1.60 $\pm$ 0.18	[-1.95; -1.26]
ω_{health} X Challenge	-0.07 $\pm$ 0.21	[-0.50; 0.34]
RST X Challenge	-0.23 $\pm$ 0.38	[-0.96; 0.49]
ω_{taste} X RST X Challenge	-0.28 $\pm$ 0.27	[-0.83; 0.24]
ω_{health} X RST X Challenge	0.22 $\pm$ 0.27	[-0.26; 0.78]

This table reports results from a Bayesian logistic regression (given in Eq. S8) modelling healthy food choices as a function of the estimated taste and health weighting parameters of the tDDM as well as the relative start time for health (RST). The model also included the signed difference in taste (TD) and healthiness (HD) ratings between the healthier and less healthy option (i.e. healthier – less healthy), and a dummy variable (Challenge) indicating whether participants faced a challenge on the respective trial. The model included subject-specific intercepts as well as subject-specific slopes for the TD, HD and Challenge regressors. Data were pooled across the baseline conditions of all 4 studies. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval. Abbreviations:

ω_{taste} : weighting factor determining how much taste contributes to the evidence accumulation rate.

ω_{health} : weighting factor determining how much healthiness contributes to the evidence accumulation rate.

RST: relative-start-time for health

Table S25. Bayesian linear regression for response times in health challenge and non-challenge trials as a function of tDDM parameters.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.55 $\pm$ 0.04	[0.47; 0.63]
Challenge	-0.12 $\pm$ 0.04	[-0.19; -0.05]
ω_{taste}	-0.04 $\pm$ 0.03	[-0.11; 0.02]
ω_{health}	-0.09 $\pm$ 0.04	[-0.16; -0.02]
RST	-0.07 $\pm$ 0.07	[-0.21; 0.06]
Healthy choice (HC)	-0.07 $\pm$ 0.03	[-0.14; -0.01]
 Taste Difference (TD) 	-0.06 $\pm$ 0.01	[-0.07; -0.04]
 Health Difference (HD) 	-0.02 $\pm$ 0.01	[-0.03; -0.01]
ω_{taste} X RST	0.08 $\pm$ 0.05	[-0.03; 0.18]
ω_{health} X RST	0.07 $\pm$ 0.05	[-0.03; 0.15]
ω_{taste} X HC	0.02 $\pm$ 0.03	[-0.03; 0.08]
ω_{health} X HC	-0.01 $\pm$ 0.03	[-0.07; 0.05]
RST X HC	0.03 $\pm$ 0.05	[-0.07; 0.13]
Challenge X ω_{taste}	0.03 $\pm$ 0.03	[-0.03; 0.09]
Challenge X ω_{health}	0.08 $\pm$ 0.03	[0.01; 0.14]
Challenge X RST	0.06 $\pm$ 0.06	[-0.06; 0.18]
Challenge X HC	0.07 $\pm$ 0.04	[0; 0.13]
Challenge X TD 	0.05 $\pm$ 0.01	[0.03; 0.06]
Challenge X HD	0 $\pm$ 0.01	[-0.01; 0.02]
ω_{taste} X RST X HC	-0.07 $\pm$ 0.04	[-0.16; 0.01]
ω_{health} X RST X HC	0.05 $\pm$ 0.04	[-0.03; 0.12]
Challenge X ω_{taste} X RST	-0.11 $\pm$ 0.05	[-0.21; -0.02]
Challenge X ω_{health} X RST	-0.06 $\pm$ 0.04	[-0.14; 0.02]
Challenge X ω_{taste} X HC	0 $\pm$ 0.03	[-0.06; 0.07]
Challenge X ω_{health} X HC	-0.04 $\pm$ 0.03	[-0.11; 0.02]
Challenge X RST X HC	-0.04 $\pm$ 0.06	[-0.15; 0.08]
Challenge X ω_{taste} X RST X HC	0.17 $\pm$ 0.05	[0.07; 0.26]
Challenge X ω_{health} X RST X HC	0.04 $\pm$ 0.04	[-0.05; 0.12]

This table reports results from a Bayesian hierarchical linear regression (given in Eq. S9) modelling response times in food choices as a function of the estimated taste and health weighting parameters of the tDDM as well as the relative start time for health (RST), and the absolute distance in taste (|TD|) and health (|HD|) ratings between the two options. Dummy variables indicated whether participants faced a challenge on the respective trial and whether they chose the healthier food (HC). The model included subject-specific intercepts and subject-specific slopes for the |TD| and |HD| regressors as well as the interactions of these regressors with the challenge trial and healthy choice dummy variables to capture individual differences in sensitivity to taste and health aspects when making healthy or unhealthy choices in challenging or non-challenging settings. Data were pooled across the baseline conditions of all 4 studies. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval.

Abbreviations:

ω_{taste} : weighting factor determining how much taste contributes to the evidence accumulation rate.

ω_{health} : weighting factor determining how much healthiness contributes to the evidence accumulation rate.

RST: relative-start-time for health

Table S26. Bayesian linear regression for response times in cases where the healthier option was chosen.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.41 $\pm$ 0.02	[0.37; 0.46]
 Taste Difference (TD) 	-0.09 $\pm$ 0.01	[-0.11; -0.06]
 Health Difference (TD) 	-0.02 $\pm$ 0.01	[-0.04; 0.00]
Challenge	-0.03 $\pm$ 0.02	[-0.07; 0.02]
 TD X Challenge	0.12 $\pm$ 0.01	[0.10; 0.15]
 HD X Challenge	-0.04 $\pm$ 0.01	[-0.07; -0.02]

This table reports results from a Bayesian linear regression (given in Eq. S10) modelling response times in trials where participants selected the healthier option as a function the absolute difference in taste (|TD|) and health (|HD|) ratings between the two options. A dummy variable (challenge) indicated whether participants faced a health challenge on the respective trial because the healthier option was not the more tasty option. The model included subject-specific intercepts and subject-specific slopes for the |TD| and |HD| regressors and their interaction with the challenge trial dummy variable to capture individual differences in sensitivity to taste and health aspects in challenging or non-challenging settings. Data were pooled across the baseline conditions of all 4 studies. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval. Participants' response times were associated with the differences in taste between the two options when they chose the healthier option in both health challenge and non-challenge trials.

Table S27. Bayesian linear regression for response times in cases where the tastier option was chosen.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.45 $\pm$ 0.02	[0.41; 0.49]
 Taste Difference (TD) 	-0.08 $\pm$ 0.01	[-0.10; -0.06]
 Health Difference (TD) 	-0.02 $\pm$ 0.01	[-0.03; -0.00]
Challenge	-0.06 $\pm$ 0.02	[-0.09; -0.02]
 TD X Challenge	0.04 $\pm$ 0.01	[0.02; 0.06]
 HD X Challenge	0.05 $\pm$ 0.01	[0.03; 0.07]

This table reports results from a Bayesian linear regression (given in Eq. S10) modelling response times in trials where participants selected the tastier option as a function the absolute difference in taste (|TD|) and health (|HD|) ratings between the two options. A dummy variable (challenge) indicated whether participants faced a health challenge on the respective trial because the healthier option was not the more tasty option. The model included subject-specific intercepts and subject-specific slopes for the |TD| and |HD| regressors and their interaction with the challenge trial dummy variable to capture individual differences in sensitivity to taste and health aspects in challenging or non-challenging settings. Data were pooled across the baseline conditions of all 4 studies. Model fits are given as the population level mean of the posterior distribution $\pm$ standard deviation (SD) and the 95% credible interval. Participants' response times were associated with the differences in healthiness between the two options when they chose the tastier option in both health challenge and non-challenge trials.

Table S28. Response times as a function of whether a subjective value based on the initially considered attribute alone favours the same choice as a subjective value derived from both tastiness and healthiness.

	Mean beta $\pm$ SD	95% Credible Interval
(Intercept)	0.36 $\pm$ 0.01	[0.34; 0.39]
Compatible Trials	-0.03 $\pm$ 0.01	[-0.05; -0.02]
 TD 	-0.02 $\pm$ 0	[-0.03; -0.01]
 HD 	-0.02 $\pm$ 0	[-0.03; -0.02]

This table reports the results of a Bayesian hierarchical linear regression (given in Eq. S11) testing for differences in response times on compatible and incompatible trials. Compatible trials were defined as those in which a subjective value based on the initially considered attribute alone favours the same choice as a subjective value derived from both tastiness and healthiness. Incompatible trials were those in which a subjective value based on the initially considered attribute alone favours the opposite choice as a subjective value derived from both tastiness and healthiness. Consistent with the predictions of a tDDM with separate consideration onset times for taste and healthiness, compatible trials were faster than incompatible trials. The model included subject-specific intercepts as well as subject-specific slopes for |TD| and |HD| to capture individual differences in sensitivity to these two attributes. The abbreviations |TD| and |HD| stand for the absolute value of the taste and healthiness differences, respectively.

5. References

- 77 Hunt, L. T., Dolan, R. J. & Behrens, T. E. Hierarchical competitions subserving multi-attribute choice. *Nat. Neurosci.* **17**, 1613-1622, doi:10.1038/nn.3836 (2014).
- 78 Brown, S. D. & Heathcote, A. The simplest complete model of choice response time: linear ballistic accumulation. *Cogn Psychol* **57**, 153-178, doi:10.1016/j.cogpsych.2007.12.002 (2008).
- 79 Stewart, N., Chater, N. & Brown, G. D. Decision by sampling. *Cognit. Psychol.* **53**, 1-26 (2006).
- 80 Roe, R. M., Busemeyer, J. R. & Townsend, J. T. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol. Rev.* **108**, 370-392 (2001).
- 81 Busemeyer, J. R. & Townsend, J. T. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432-459 (1993).
- 82 Cisek, P., Puskas, G. A. & El-Murr, S. Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* **29**, 11560-11571, doi:10.1523/JNEUROSCI.1844-09.2009 (2009).
- 83 Usher, M. & McClelland, J. L. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550-592 (2001).
- 84 Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700-765, doi:10.1037/0033-295X.113.4.700 (2006).
- 85 Thura, D. & Cisek, P. On the difference between evidence accumulator models and the urgency gating model. *J. Neurophysiol.* **115**, 622-623, doi:10.1152/jn.01050.2015 (2016).
- 86 Hawkins, G. E., Wagenmakers, E. J., Ratcliff, R. & Brown, S. D. Discriminating evidence accumulation from urgency signals in speeded decision making. *J. Neurophysiol.* **114**, 40-47, doi:10.1152/jn.00088.2015 (2015).
- 87 Diederich, A. & Oswald, P. Multi-stage sequential sampling models with finite or infinite time horizon and variable boundaries. *J. Math. Psychol.* **74**, 128-145 (2016).
- 88 Holmes, W. R. & Trueblood, J. S. Bayesian analysis of the piecewise diffusion decision model. *Behavior research methods*, 1-14 (2017).
- 89 Holmes, W. R., Trueblood, J. S. & Heathcote, A. A new framework for modeling decisions about changing information: The Piecewise Linear Ballistic Accumulator model. *Cogn Psychol* **85**, 1-29, doi:10.1016/j.cogpsych.2015.11.002 (2016).
- 90 Ratcliff, R. A note on modeling accumulation of information when the rate of accumulation changes over time. *J. Math. Psychol.* **21**, 178-184 (1980).
- 91 Turner, B. M., Schley, D. R., Muller, C. & Tsetsos, K. Competing theories of multialternative, multiattribute preferential choice. *Psychol. Rev.* **125**, 329 (2018).

- 92 Hare, T. A., Camerer, C. F. & Rangel, A. Self-control in decision-making involves modulation of the vmPFC valuation system. *Science (New York, N.Y.)* **324**, 646-648, doi:10.1126/science.1168450 (2009).
- 93 Maier, S. U., Makwana, A. B. & Hare, T. A. Acute Stress Impairs Self-Control in Goal-Directed Choice by Altering Multiple Functional Connections within the Brain's Decision Circuits. *Neuron* **87**, 621-631, doi:10.1016/j.neuron.2015.07.005 (2015).
- 94 Logan, G. D. & Cowan, W. B. On the ability to inhibit thought and action—a theory of an act of control. *Psychol. Rev.* **91**, 295-327 (1984).
- 95 Logan, G. D., Cowan, W. B. & Davis, K. A. On the Ability to Inhibit Simple and Choice Reaction-Time Responses - a Model and a Method. *J. Exp. Psychol. Hum. Percept. Perform.* **10**, 276-291 (1984).
- 96 Cubillo, A. *et al.* Reduced activation and inter-regional functional connectivity of fronto-striatal networks in adults with childhood Attention-Deficit Hyperactivity Disorder (ADHD) and persisting symptoms during tasks of motor inhibition and cognitive switching. *J. Psychiatr. Res.* **44**, 629-639 (2010).
- 97 Cubillo, A. *et al.* Shared and Drug-Specific Effects of Atomoxetine and Methylphenidate on Inhibitory Brain Dysfunction in Medication-Naive ADHD Boys. *Cereb. Cortex* **24**, 174-185 (2014).
- 98 Wechsler, D. *Wechsler Adult Intelligence Scale - III*. (The Psychological Corporation, 1997).
- 99 Cooper, N., Kable, J. W., Kim, B. K. & Zauberman, G. Brain activity in valuation regions while thinking about the future predicts individual discount rates. *J. Neurosci.* **33**, 13150-13156 (2013).
- 100 Becker, G. M., Degroot, M. H. & Marschak, J. Measuring utility by a single-response sequential method. *Behav. Sci.* **9**, 226-232 (1964).
- 101 Carpenter, B. *et al.* Stan: A probabilistic programming language. *Journal of Statistical Software* **20**, 1-37 (2016).
- 102 Plummer, M. in *Proceedings of the 3rd international workshop on distributed statistical computing*.
- 103 Bürkner, P.-C. brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software* **80**, 1-28, doi:10.18637/jss.v080.i01 (2017).
- 104 Kruschke, J. K. Bayesian estimation supersedes the t test. *J. Exp. Psychol. Gen.* **142**, 573-603, doi:10.1037/a0029146 (2013).
- 105 Eddelbuettel, D. *et al.* Rcpp: Seamless R and C++ integration. *Journal of Statistical Software* **40**, 1-18 (2011).
- 106 Kruschke, J. K. *Bayesian estimation of correlations and differences of correlations with a multivariate normal* <<http://doingbayesiandataanalysis.blogspot.com/2017/06/bayesian-estimation-of-correlations-and.html>> (2017).
- 107 yarr: A Companion to the e-Book "YaRrr!: The Pirate's Guide to R" v. R package version 0.1.5 (2017).
- 108 pracma: Practical Numerical Math Functions v. R package version 2.1.4 (2018).
- 109 Wickham, H. The split-apply-combine strategy for data analysis. *Journal of Statistical Software* **40**, 1-29 (2011).

- 110 loo: Efficient leave-one-out cross-validation and WAIC for Bayesian models. R package version 0.1. 6 (2016).
- 111 Vehtari, A., Gelman, A. & Gabry, J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and computing* **27**, 1413-1432 (2017).
- 112 Pudiel, V. & Westenhöfer, S. Fragebogen zum Ernährungsverhalten. (1989).