Supplementary material for

Decomposing preferences into predispositions and evaluations

Nitisha Desai[1], Ian Krajbich[1,2]

[1]Department of Psychology, The Ohio State University

[2]Department of Economics, The Ohio State University

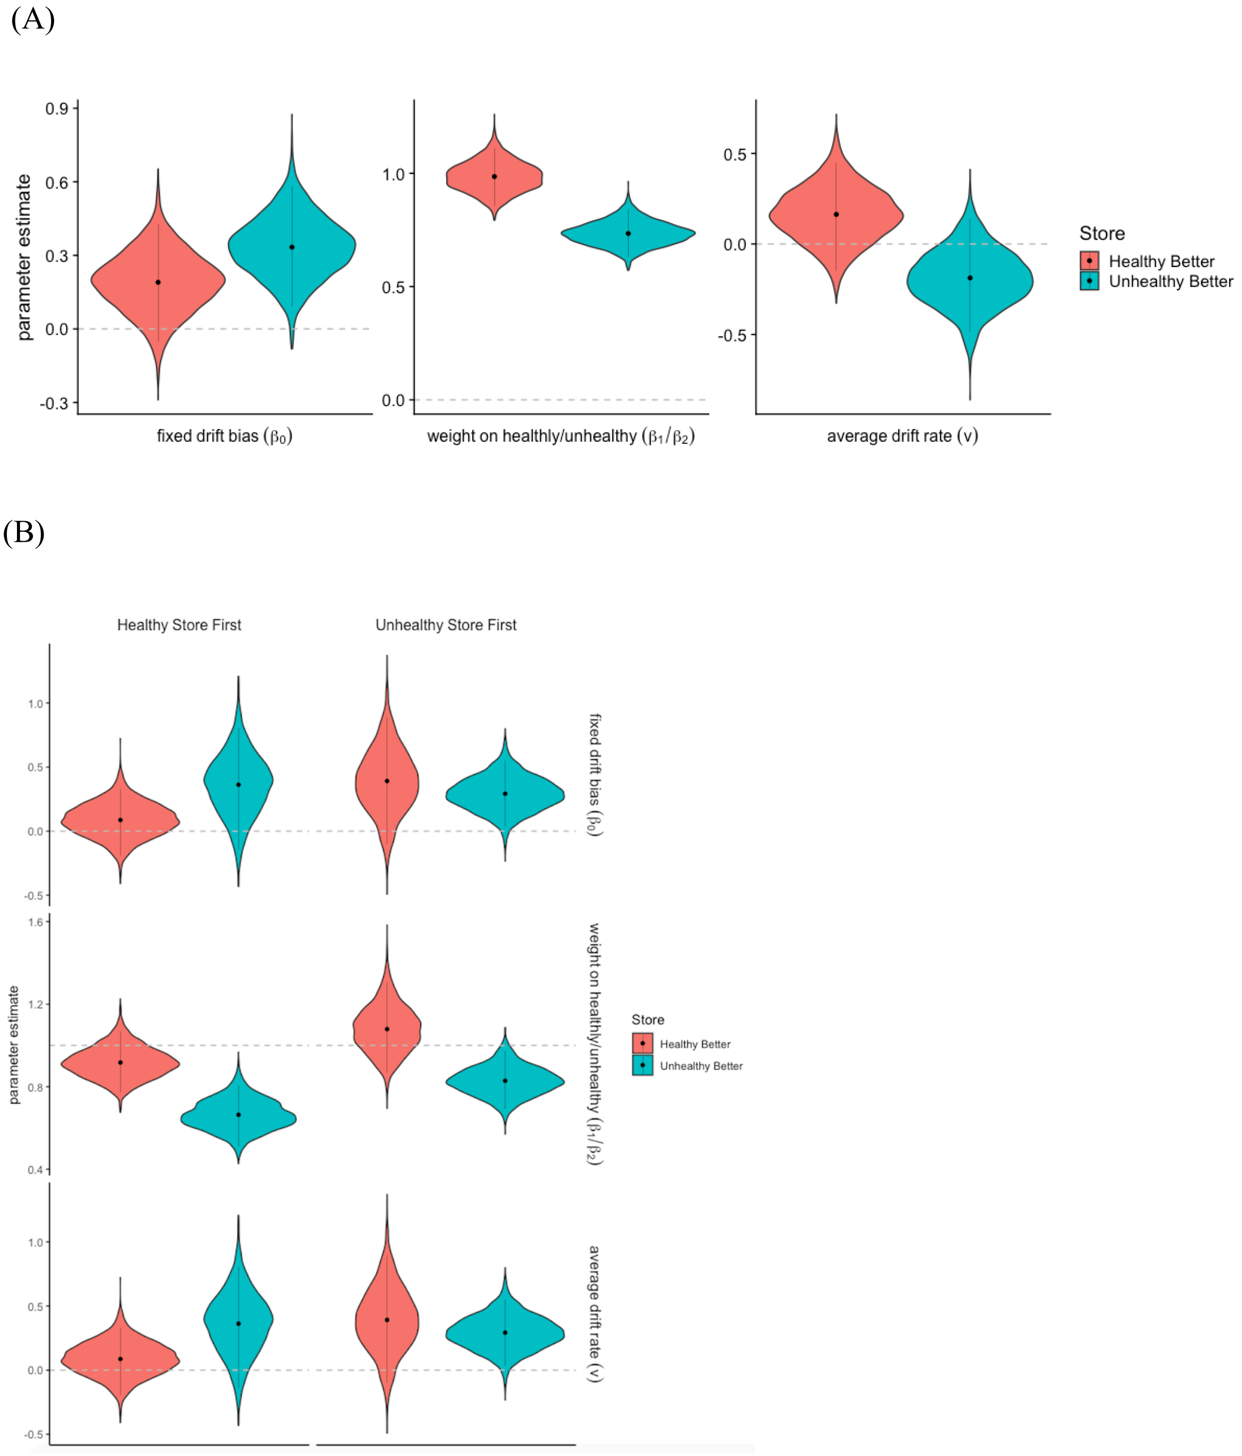

Figure S1. Experiment 1: Parameter estimates for all components of the drift regression (A) and order effects (B). The fixed drift rate (β_0), the weight on healthy relative to unhealthy ($\frac{\beta_1}{\beta_2}$), and the average drift rate (*v*) are shown. We calculated the average drift rate using values of 5 (the average rating for a food) for both items in the drift rate formula (Equation 2).

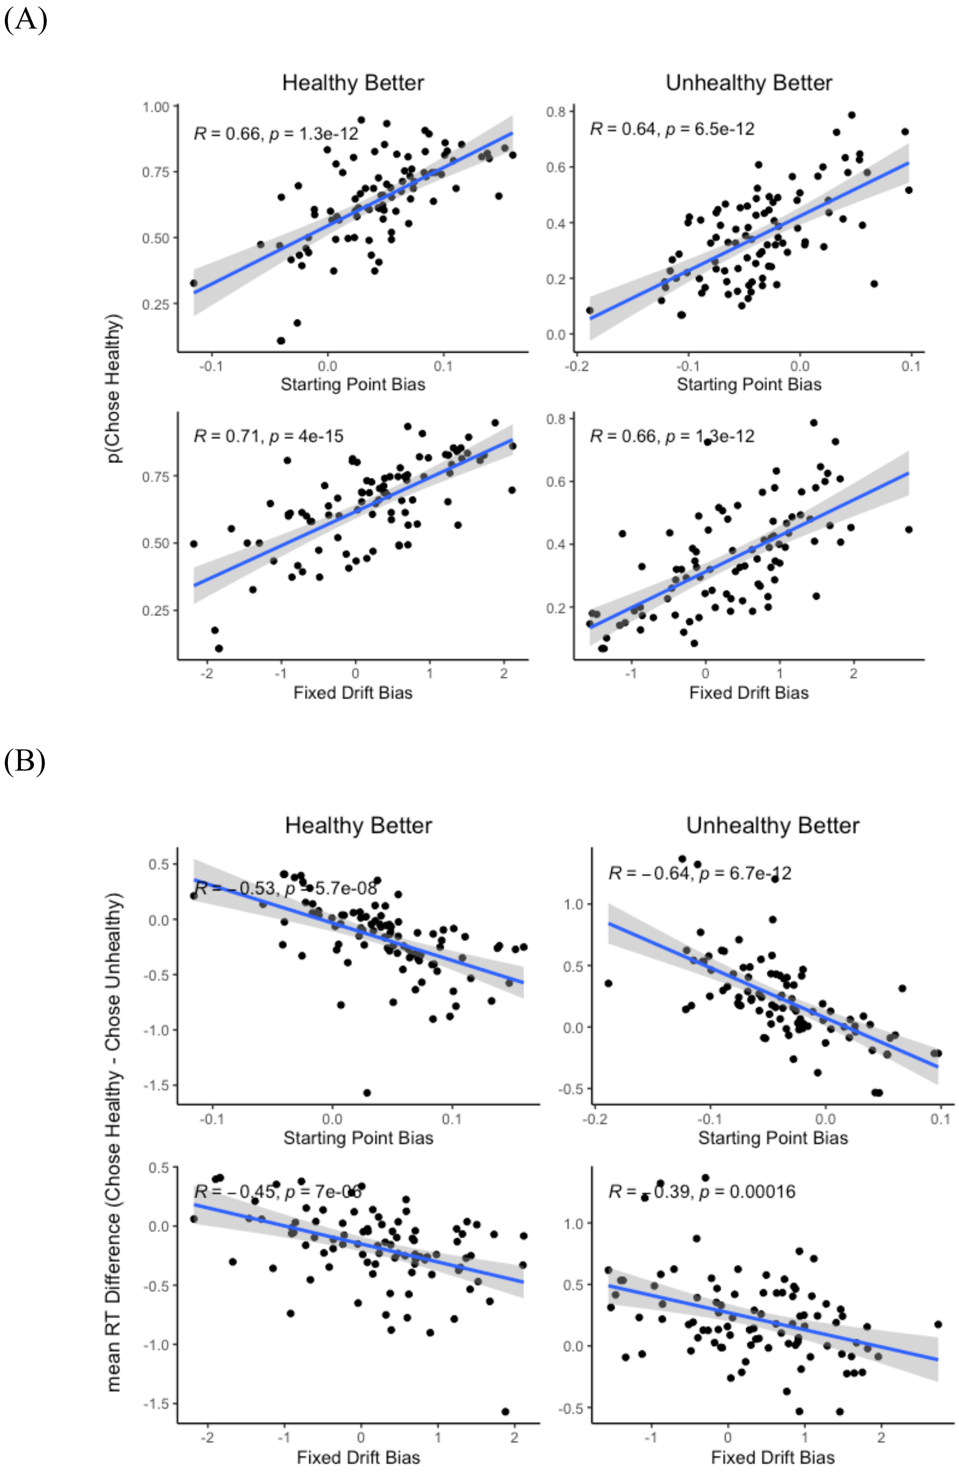

Figure S2. Experiment 1: Relationship between individual-level fixed drift bias and starting point bias vs. choice bias (A) and vs. mean RT difference (B). Choice bias is the probability of choosing the healthy option. Mean RT difference is the difference between a subject's mean RT for healthy choices and their mean RT for unhealthy choices. Both drift rates and starting points help to explain subjects' choices and RT.

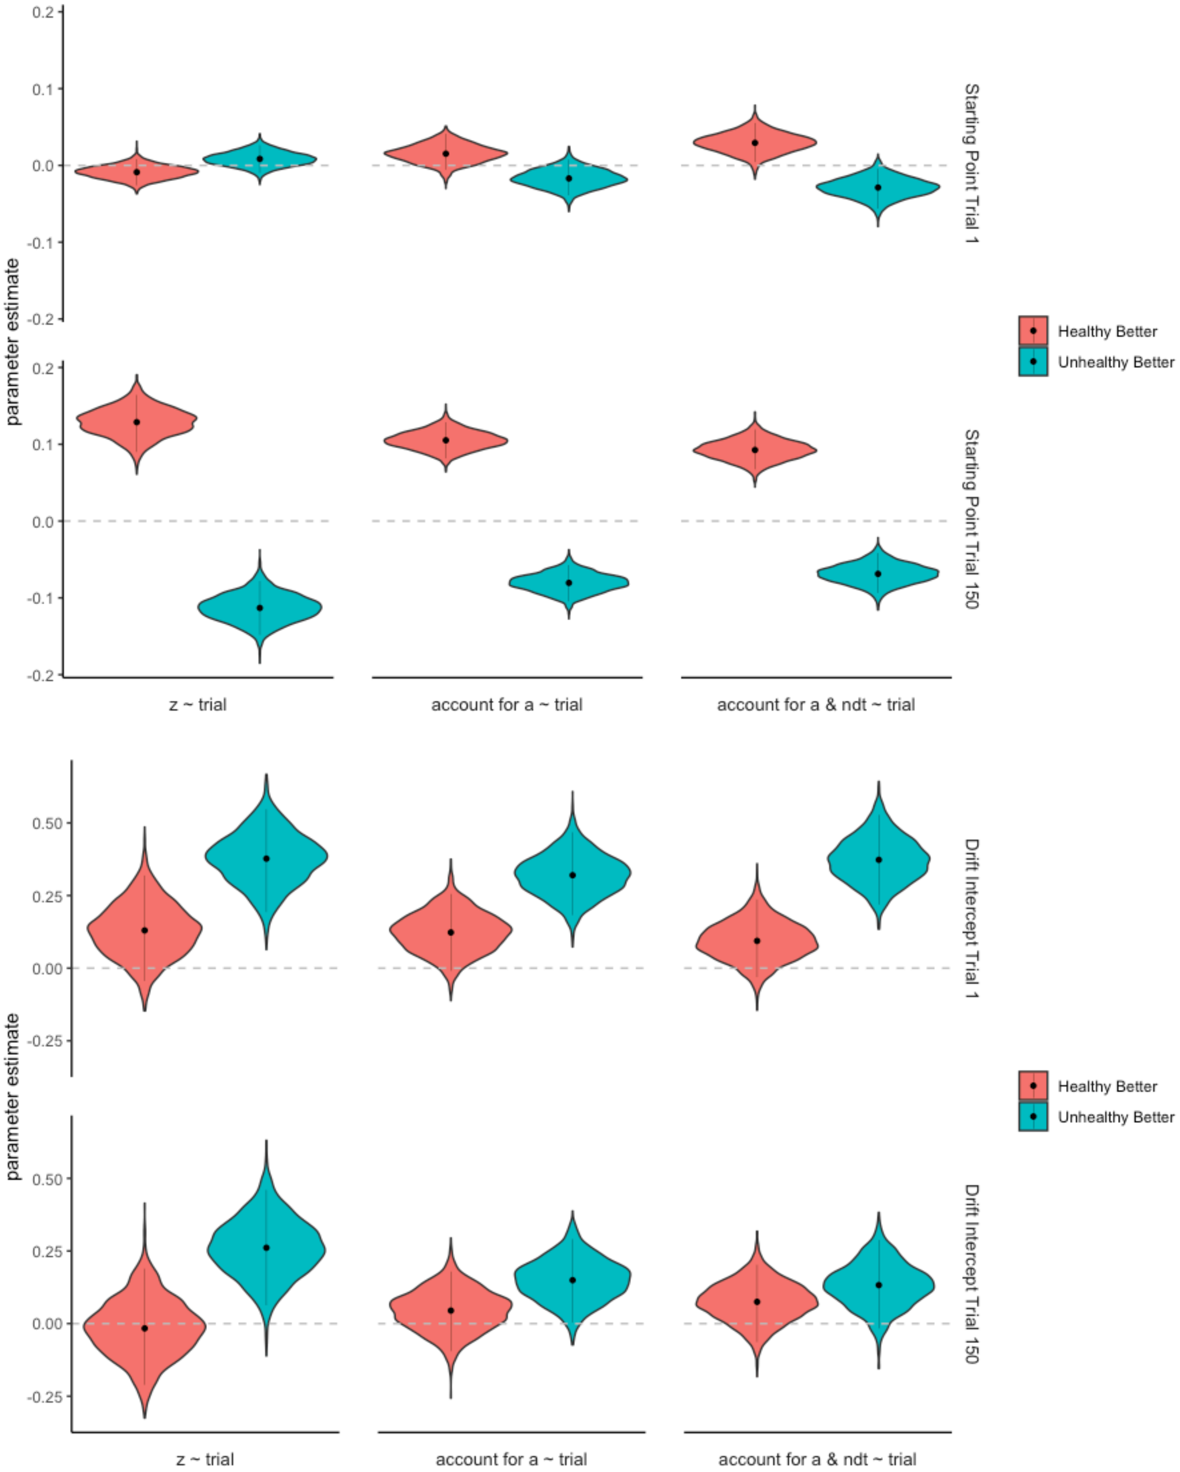

Figure S3. Experiment 1 time trends where starting point (*z*) (left); *z* and boundary separation (*a*) (center); and *z*, *a*, and non-decision time (*ndt*) (right) depend on trial. Even when accounting for how *a* and *ndt* change with trial number, the results remain essentially unchanged.

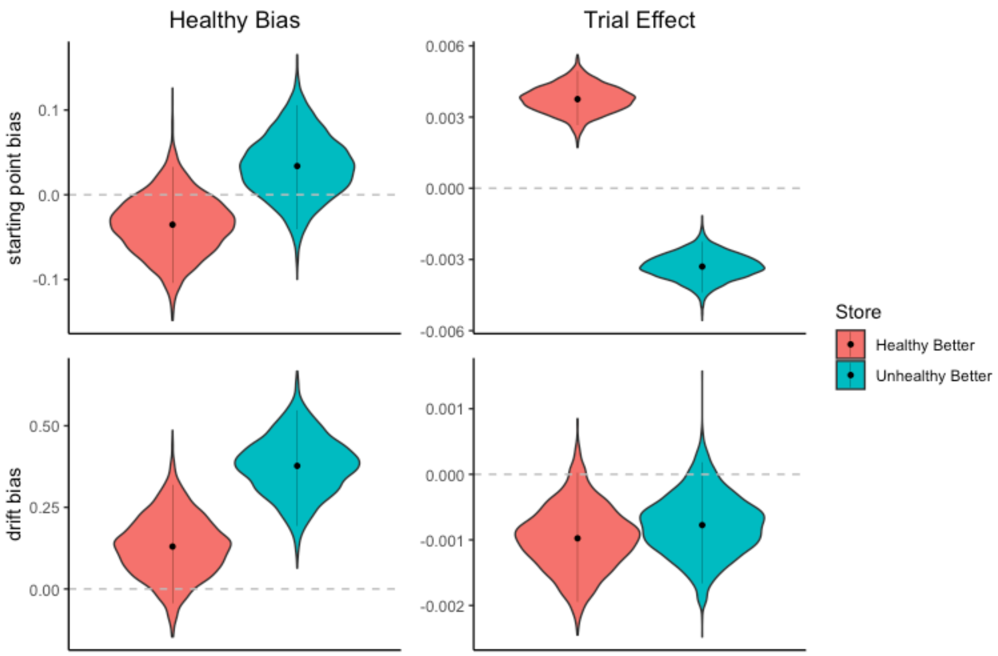

Figure S4. Experiment 1: Parameter estimates for the analysis of time trends in starting point. Starting point (top row) was regressed on trial with a logistic link function: $logit(z) = \gamma_0 + \gamma_1 \times Trial$. All intercepts ($\gamma_0$, healthy bias) were zero, indicating that subjects did not exhibit a predisposition. The coefficient on the trial regressor (γ_1) was either positive or negative depending on whether the store generally had better healthy items or unhealthy items. The drift rate function (bottom row) was extended to include *Trial* as a regressor: $v = \beta_0 + \beta_1 \times Healthy\ Rating - \beta_2 \times Unhealthy\ Rating + \beta_3 \times Trial$. The intercept ($\beta_0$) corresponds to an evaluation bias towards one of the categories (healthy or unhealthy) at the beginning of the block. The coefficients on the trial regressor (β_3) did not credibly differ from zero in either condition suggesting no change in the evaluation bias over the course of the block.

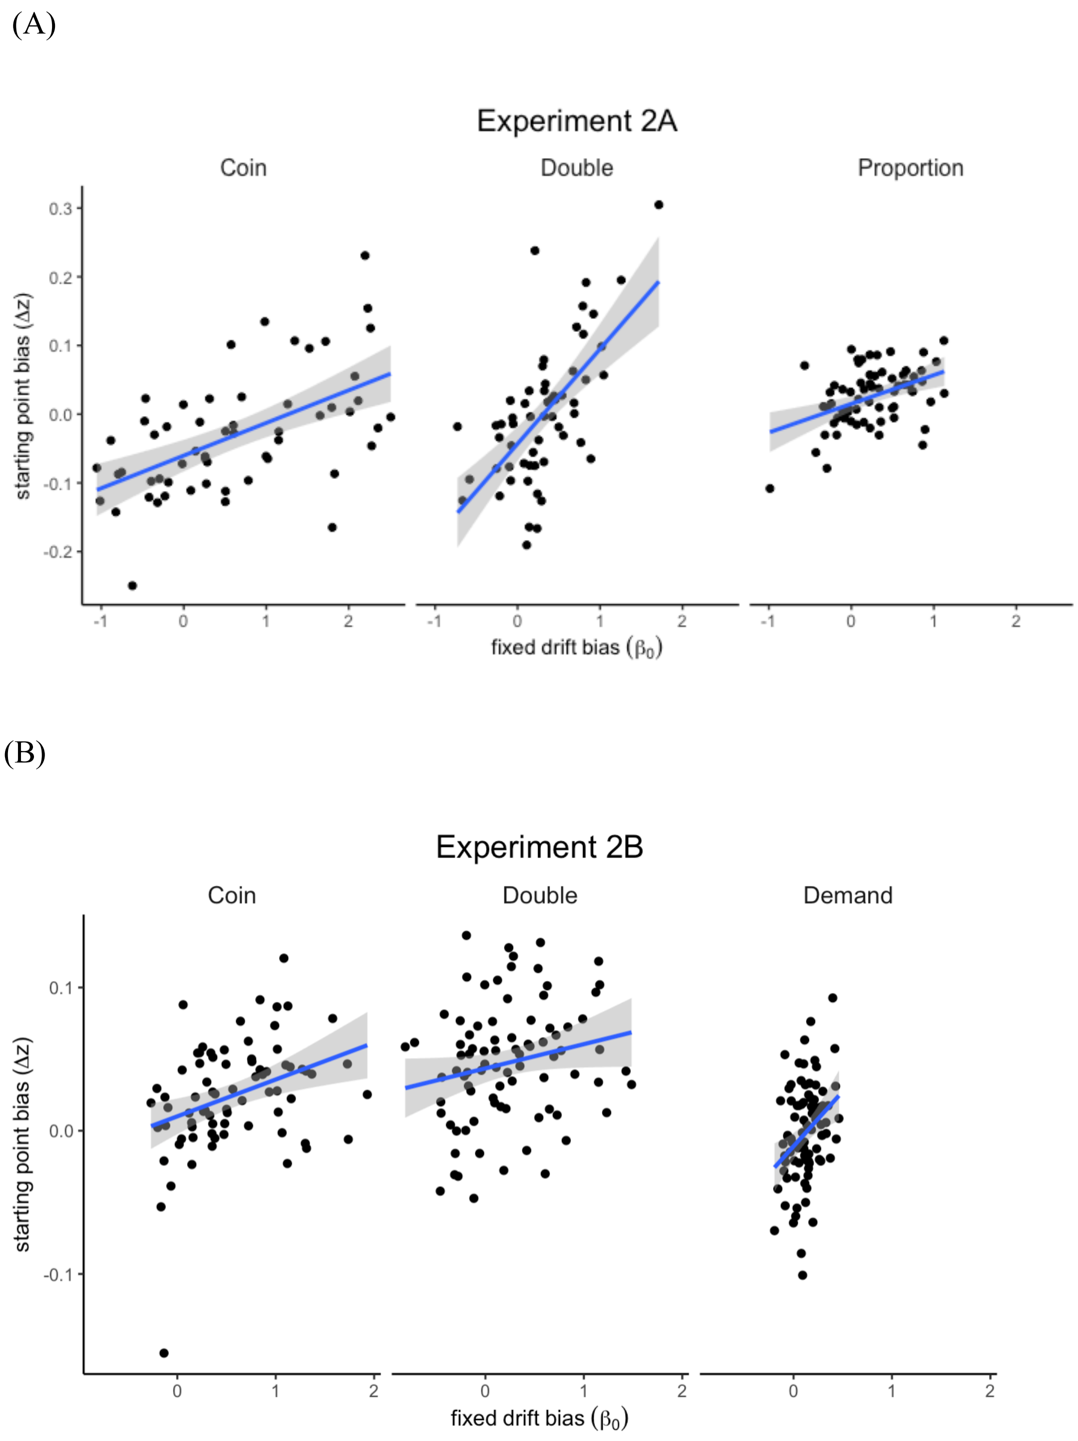

Figure S5. Relationship between individual-level starting point bias and fixed drift bias for Experiment 2A (A) and Experiment 2B (B). Data are presented separately for the Coin (N=57, 76), Double (N=58, 85), and Proportion/Demand (N=66, 89) conditions. In all conditions, fixed drift bias was positively correlated with starting-point bias, indicating that the stronger the average fixed drift bias toward the target options, the stronger the expectation that the target side was better.

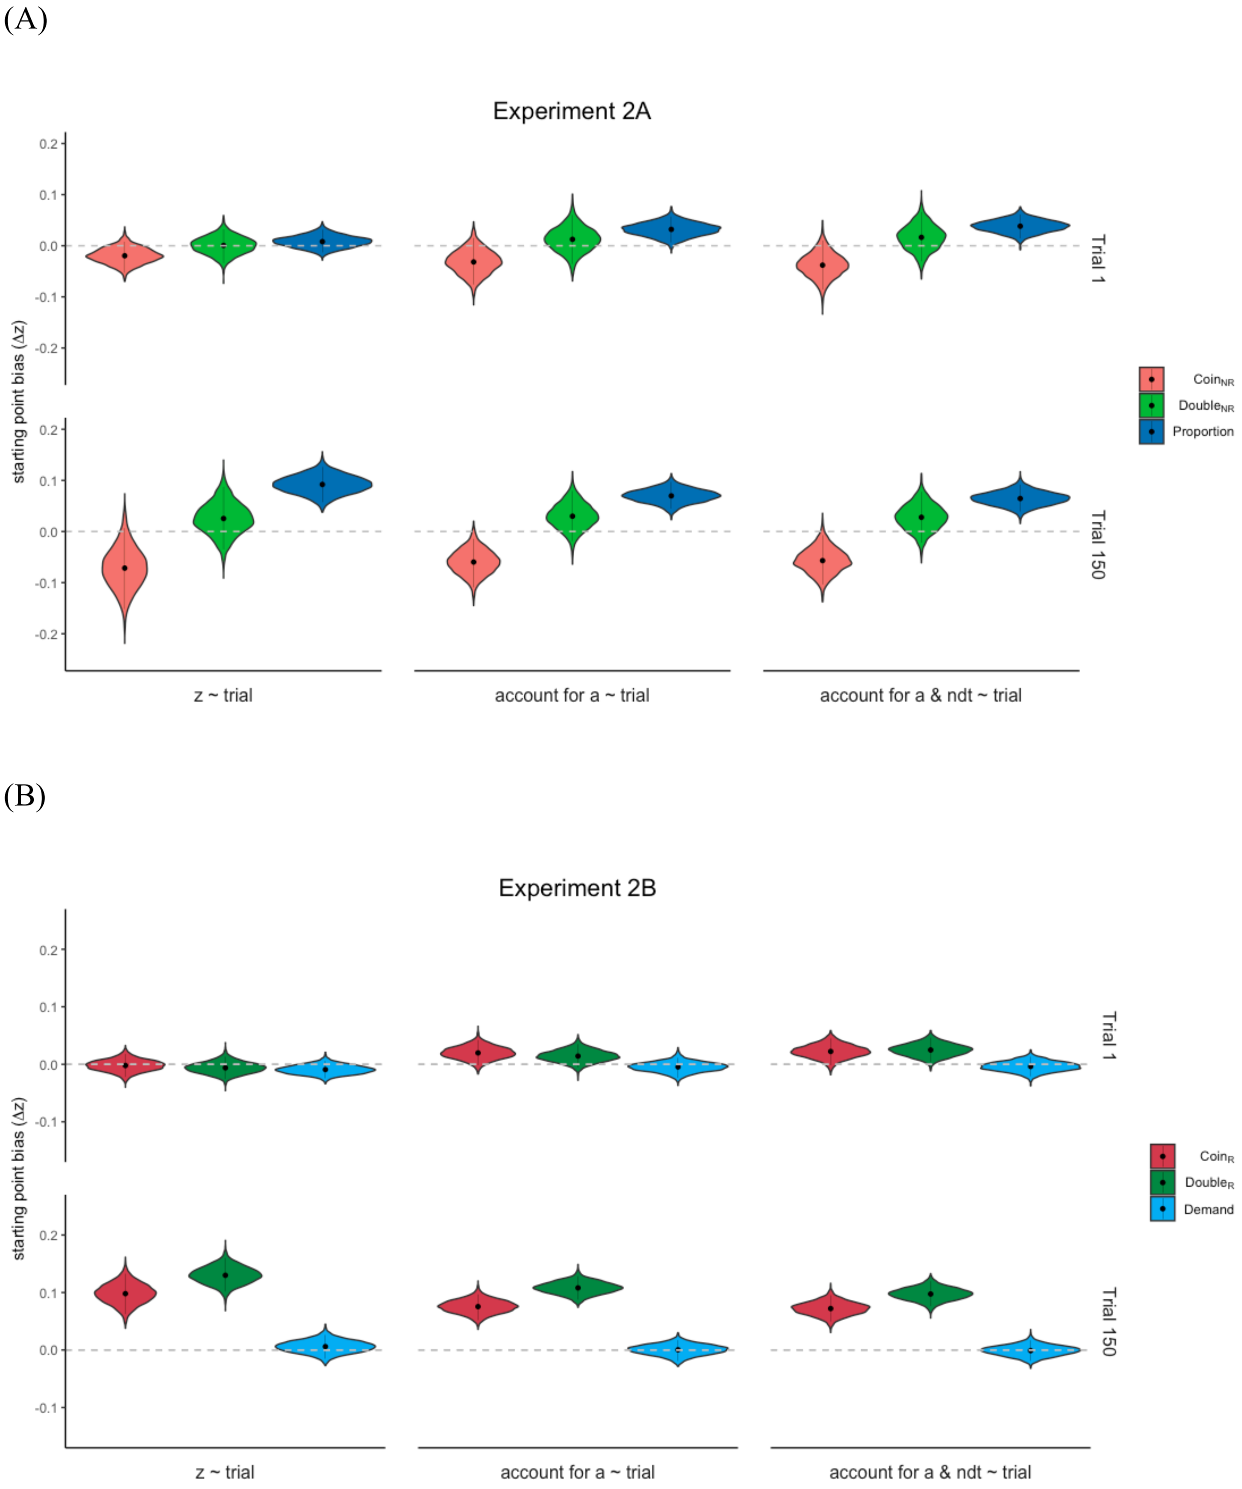

Figure S6. Experiment 2A (A) and 2B (B) time trends, where starting point (*z*) (left); *z* and boundary separation (*a*) (center); and *z*, *a*, and non-decision time (*ndt*) (right) depend on trial. Even when accounting for how *a* and *ndt* change with trial number, the results remain essentially unchanged.

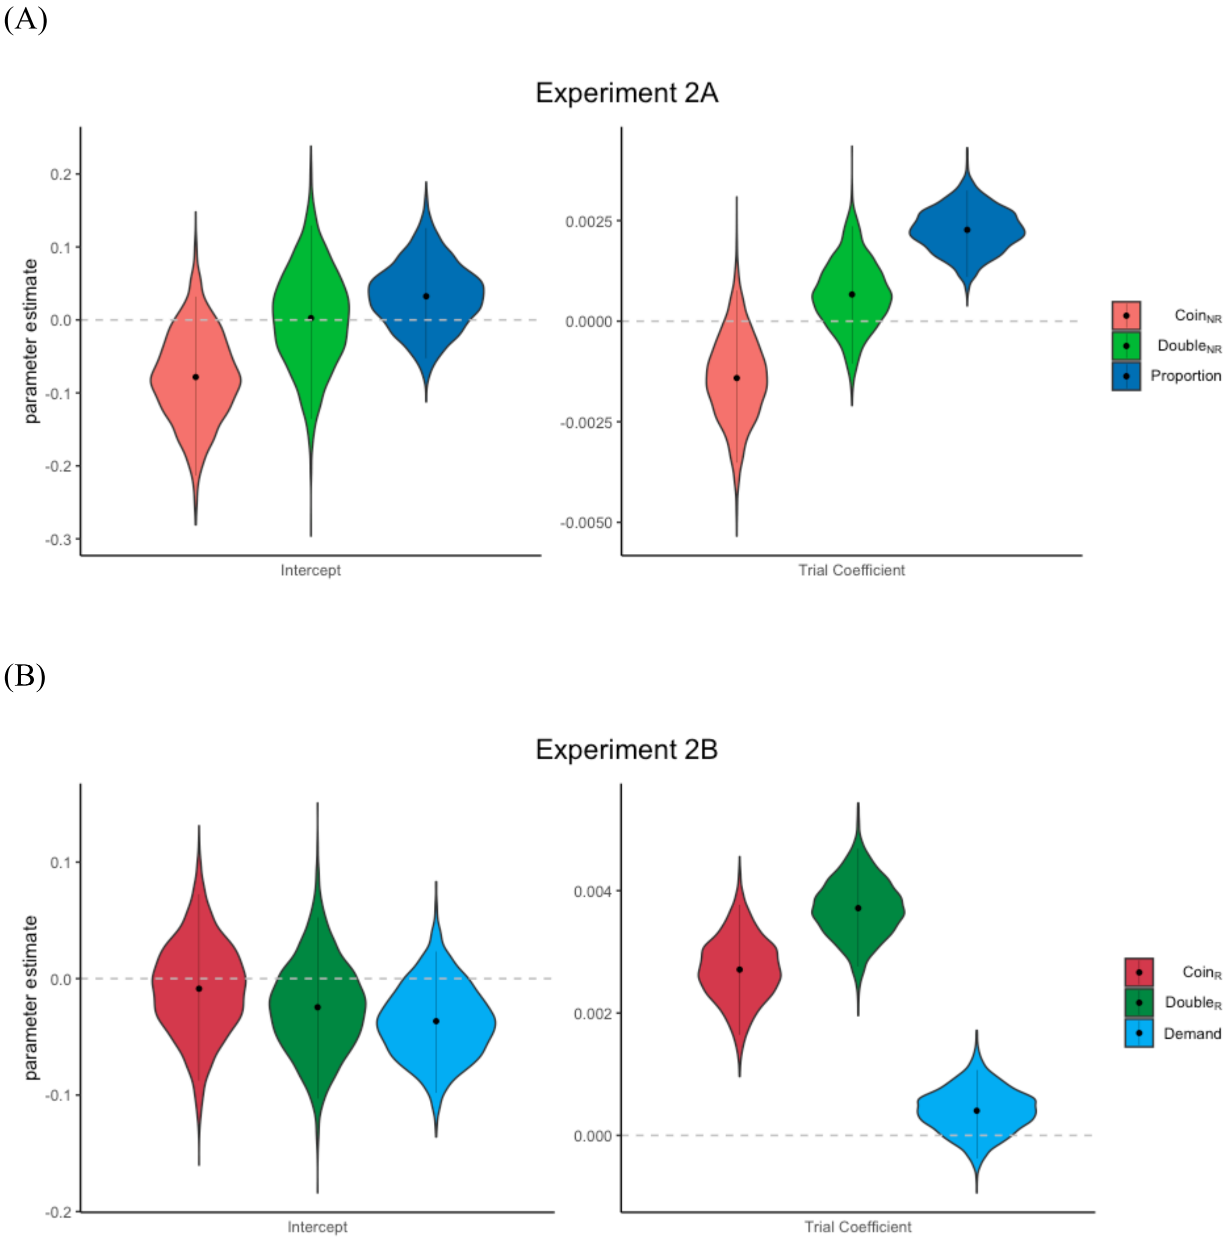

Figure S7. Parameter estimates for the analysis of time trends in starting point for Experiment 2A (A) and 2B (B). Starting point was regressed on trial with a logistic link function: $logit(z) = \gamma_0 + \gamma_1 \times Trial$. All intercepts ($\gamma_0$) were zero, indicating that subjects did not initially exhibit a predisposition. The coefficient on the trial regressor (γ_1) was positive in the Proportion, Coin_B, and Double_B conditions where predispositions were reinforced, and zero in the Demand, Coin_A, and Double_A conditions where predispositions were not reinforced.

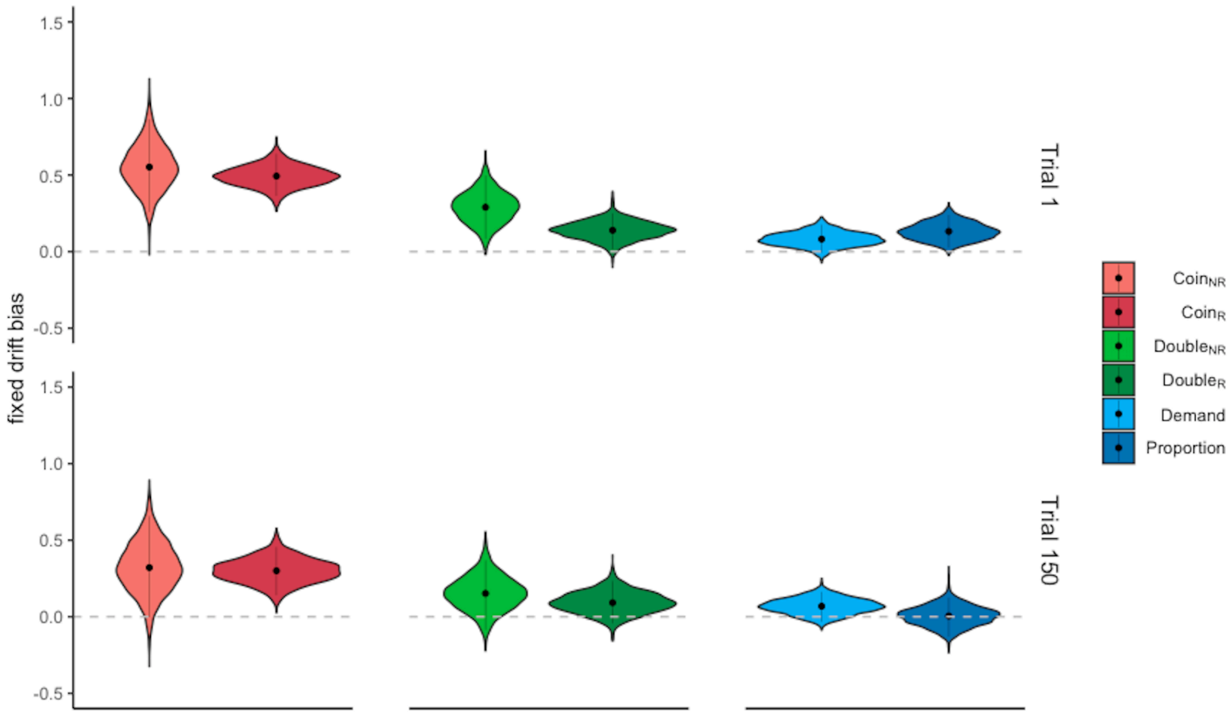

Figure S8. Experiment 2 evaluation bias. Drift biases were estimated as a function of trial number. Based on this model, we estimate the drift biases at the beginning (Trial 1) and end (Trial 150) of each condition. These results suggest that the evaluation biases did not credibly change from the beginning to the end of the block.

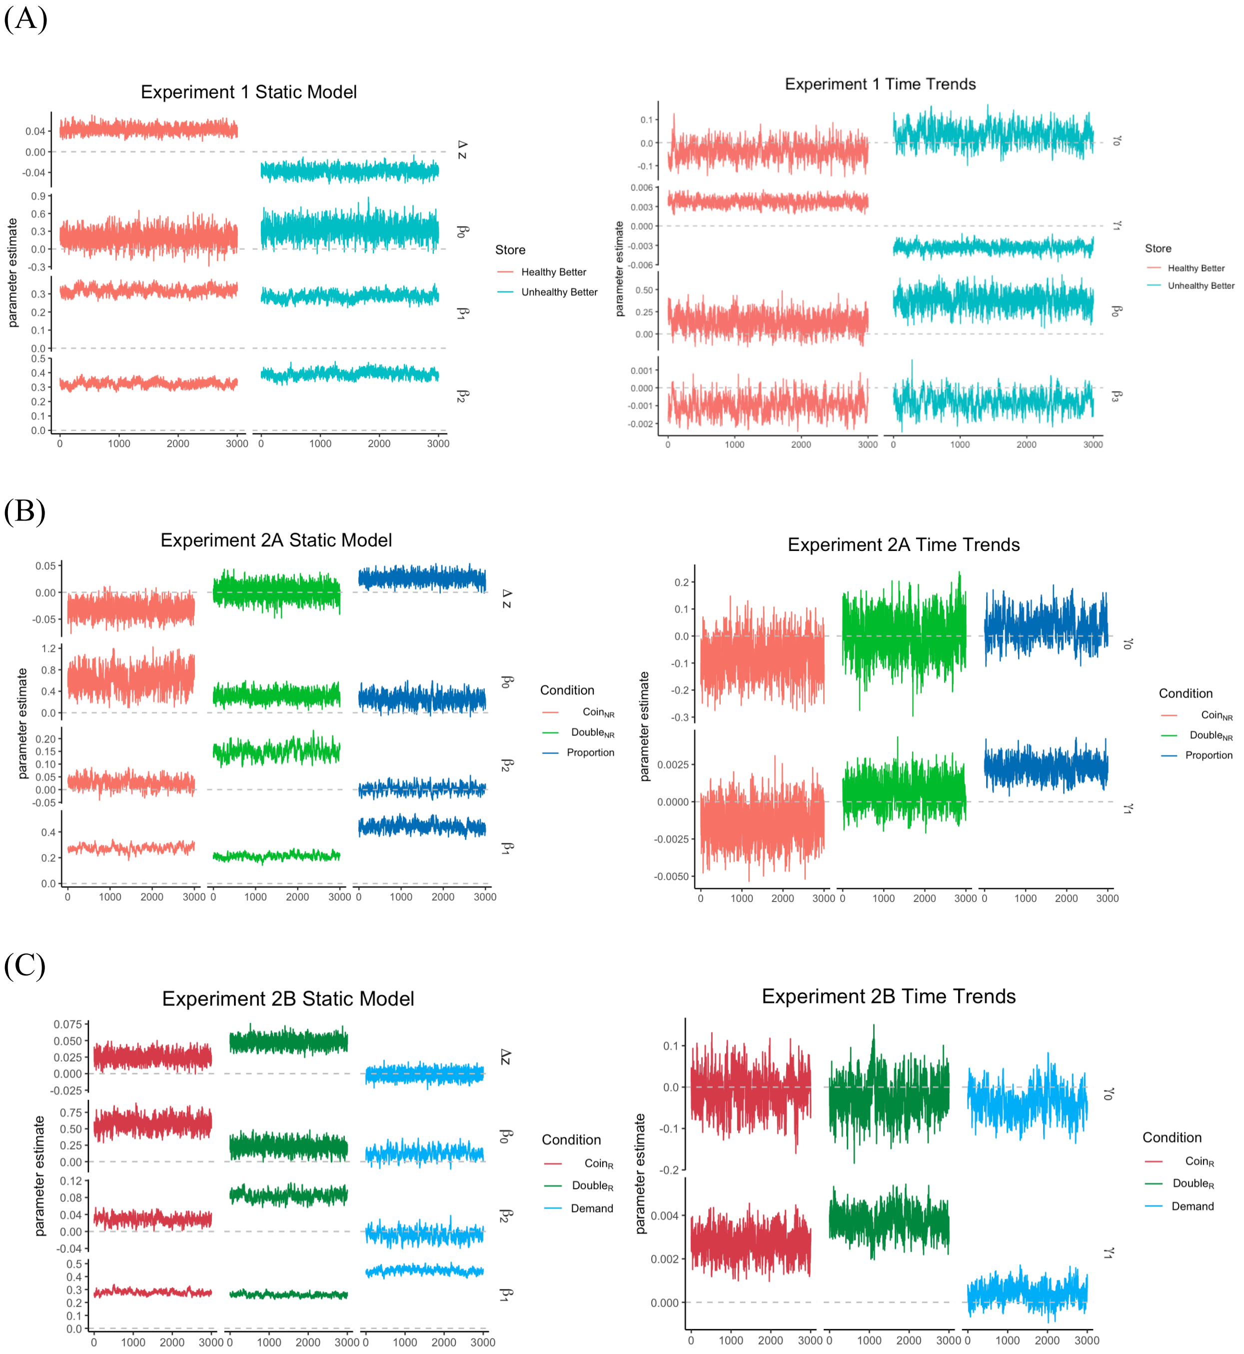

Figure S9. Trace plots. A visual inspection of all traces indicates convergence. For the Experiment 1 models (A) and the Experiment 2 time trends analyses (B and C, right), 3 chains of 2000 samples with 1000 burn-in were run. Concatenated sample from the 3 chains are displayed. All R-hats < 1.1. For the Experiment 2 static models (B and C, left), 1 chain of 6000 samples with 3000 burn-in was run.

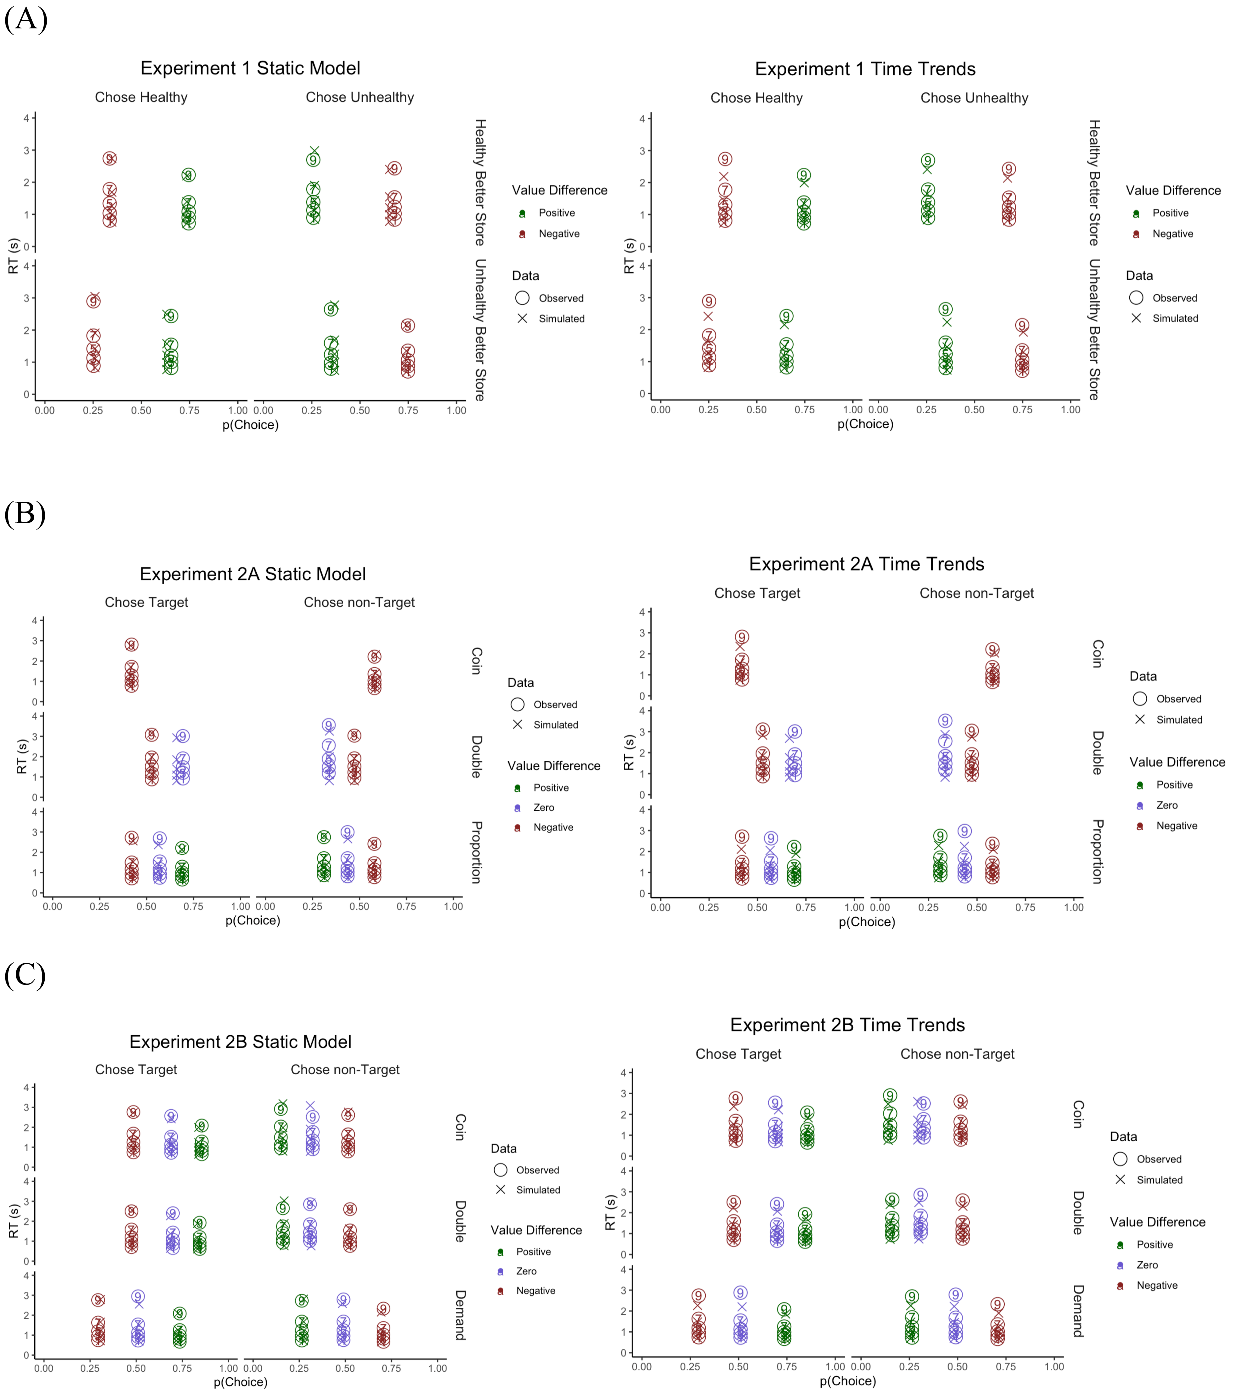

Figure S10. Quantile-Probability Plots. For each model, we generated 500 datasets using the estimated parameter values. The data is binned by choice (chose Healthy/Target vs. chose Unhealthy/non-Target) and by value difference (positive value difference, equally valued, negative value difference). The y-axis shows the mean RT at the 0.1, 0.3, 0.5, 0.7, and 0.9 quantiles for observed vs. simulated data. The x-axis shows the probability of choosing either Target/Healthy or non-Target/Unhealthy. The static models are on the left and the time trend models are on the right.

(A)

Experiment 2A			
Third of Trials	Coin	Double	Proportion
First Third	0.443 (0.031)	0.557 (0.026)	0.644 (0.015)
Middle Third	0.417 (0.033)	0.552 (0.026)	0.645 (0.014)
Last Third	0.403 (0.035)	0.508 (0.032)	0.644 (0.016)

(B)

Experiment 2B			
Third of Trials	Coin	Double	Demand
First Third	0.677 (0.016)	0.650 (0.015)	0.506 (0.010)
Middle Third	0.677 (0.015)	0.659 (0.014)	0.507 (0.008)
Last Third	0.664 (0.018)	0.697 (0.014)	0.524 (0.011)

Table S1. Probability of selecting the target option for trials binned by the first third, middle third, and last third of trials for Experiment 2A (A) and Experiment 2B (B). Each bin has approximately 50 trials.

Behavioral Analyses

In Experiment 1, we modeled the behavioral data with following mixed-effects logistic regression allowing for subject-level random effects for the slopes and intercept:

$$logit(Choose\ Healthy)$$
$$= \beta_0 + \beta_1 \times (Healthy\ Rating - Unhealthy\ Rating) + \beta_2 \times Better\ Categroy$$

In Experiment 2, we modeled the behavioral data with following mixed-effects logistic regression allowing for subject-level random effects for the slope and intercept:

$$logit(Choose\ Target) = \beta_0 + \beta_1 \times Value\ Difference$$

In Experiment 1, we modeled the RT data with the following mixed-effects regression allowing for subject-level random effects for the slopes and intercept:

$$log(RT) = \beta_0 + \beta_1 \times |Healthy\ Rating - Unhealthy\ Rating|$$
$$+ \beta_2 \times (Target\ Category\ Is\ Higher\ Valued)$$

In Experiment 2, we modeled the RT data with two mixed-effects regression allowing for subject-level random effects for the slopes and intercept. The first was identical to the RT regression Experiment 1:

$$log(RT) = \beta_0 + \beta_1 \times |Value\ Difference| + \beta_2 \times (Target\ Is\ Higher\ Valued)$$

And we conducted a second analysis using a quadratic regression:

$$log(RT) = \beta_0 + \beta_1 \times Value\ Difference + \beta_2 \times Value\ Difference^2$$

Experiment 2 RT Analyses

We examined the inverted U-shape in the RT curves. We expected the peaks to coincide with the indifference points from the choice data (Figure 7A). In other words, we expected to find peaks in the RT curves when the targets were rated lower than the non-targets. To test these two hypotheses, we ran mixed-effects regressions with *log(RT)* as a function of *Value Difference* and *Value Difference*^{2}, with random effects of subject on the intercept and slope.

In Experiment 2A, we found significant quadratic effects in the Double$_A$ ($\beta_2 = -0.708$, $t = -2.60$, $p = 0.01$) and Proportion ($\beta_2 = -1.47$, $t = -4.15$, $p = 10^{-5}$) conditions, but not in the Coin$_A$ condition ($\beta_2 = 0.0001$, $t = 0.00$, $p = 1$). From the linear coefficients of the model we can derive the peaks of the RT curves, which were estimated to be –at 0.382 in the Proportion condition and –0.473 in the Double$_A$ condition. Unfortunately, in the Coin$_A$ condition, the range of value differences that we selected for the experiment (Target – Non-target: [-6, -2]) was not ideal for testing these hypotheses because subjects were on average indifferent at –2. Therefore, in this condition, we were not set up to detect both sides of the inverted-U, only one side of it, which is indistinguishable from a linear increase.

In Experiment 2B, the inverted-U-shaped RT curves were also biased, with peaks to the left of value difference = 0 (Figure 7B), mirroring the choice data. The mixed-effects regressions indicated quadratic effects in the Coin$_B$ ($\beta_2 = –0.432$, $t = -6.46$, $p = 10^{-8}$), Double$_B$ ($\beta_2 = –0.443$, $t = -8.92$, $p = 10^{-13}$), and Demand ($\beta_2 = –0.547$, $t = -10.2$, $p = 10^{-16}$) conditions. From the linear coefficients of the model we can derive the peaks of the RT curves, which were estimated to be at –2.44 in the Double$_B$ condition, –2.50 in the Coin$_B$ condition, and –0.301 in the Demand condition.

Thus, in the majority of the conditions, we verify the inverted-U shape and that the peaks of the RT curves were shifted in the expected direction.

Priors

Using HDDM, we fit hierarchical Bayesian models. With these hierarchical models, we can estimate individual-level and group-level parameters concurrently. Individual-level parameters are assumed to come from group-level distributions, defined by the group-level parameters. For instance, a subject's starting point bias (*z*) follows a group-level normal distribution parameterized by a mean (μ_z) and standard deviation (σ_z). The priors on the parameters are specified below:

Static Model:
$\mu_a \sim \text{Gamma}(1.5, 0.75)$
$\mu_t \sim \text{Normal}(0.4, 0.3)$
$\mu_z \sim \text{Normal}(0.5, 0.5)$
$\mu_{\beta 0} \sim \text{Normal}(2, 3)$
$\mu_{\beta 1\text{-}2} \sim \text{Normal}(0, 15)$

$\sigma_a \sim \text{Half-Normal}(\text{var} = 0.1)$
$\sigma_t \sim \text{Half-Normal}(1)$
$\sigma_z \sim \text{Half-Normal}(0.05)$
$\sigma_{\beta 0} \sim \text{Half-Normal}(2)$
$\sigma_{\beta 1\text{-}2} \sim \text{Uniform}(\text{lower} = 10^{-10}, \text{upper} = 100)$

$sv \sim \text{Half-Normal}(2)$
$st \sim \text{Half-Normal}\ (0.3)$
$sz \sim \text{Beta}(\alpha = 1, \beta = 3)$

$a_{\text{subject}} \sim \text{Gamma}(\mu_a, \sigma_a^2)$
$t_{\text{subject}} \sim \text{Gamma}(\mu_t, \sigma_t^2)$
$z_{\text{subject}} \sim \text{invlogit}(\ \text{Normal}(\mu_z, \sigma_z^2)\ )$
$\beta_{0,\text{subject}} \sim \text{Normal}(\mu_{\beta 0}, \sigma_{\beta 0}^2)$
$\beta_{1\text{-}2,\text{subject}} \sim \text{Normal}(\mu_{\beta 1\text{-}2}, \sigma_{\beta 1-2}^2\ )$

Time Trends Model:
$\mu_a \sim \text{Gamma}(1.5, 0.75)$
$\mu_t \sim \text{Normal}(0.4, 0.3)$
$\mu_z \sim \text{Normal}(0.5, 0.5)$
$\mu_{\beta 0} \sim \text{Normal}(2, 3)$
$\mu_{\beta 1\text{-}3} \sim \text{Normal}(0, 15)$
$\mu_{\gamma 0} \sim \text{Normal}(0, 15)$
$\mu_{\gamma 1} \sim \text{Normal}(0, 15)$

$\sigma_a \sim$ Half-Normal(var = 0.1)
$\sigma_t \sim$ Half-Normal(1)
$\sigma_z \sim$ Half-Normal(0.05)
$\sigma_{\beta 0} \sim$ Half-Normal(2)
$\sigma_{\beta 1\text{-}3} \sim$ Uniform(lower = 10^{-10}, upper = 100)
$\sigma_{\gamma 0} \sim$ Uniform(lower = 10^{-10}, upper = 100)
$\sigma_{\gamma 1} \sim$ Uniform(lower = 10^{-10}, upper = 100)

$sv \sim$ Half-Normal(2)
$st \sim$ Half-Normal(0.3)
$sz \sim$ Beta(1, 3)

$a_{\text{subject}} \sim \text{Gamma}(\mu_a, \sigma_a^2)$
$t_{\text{subject}} \sim \text{Gamma}(\mu_t, \sigma_t^2)$
$\beta_{0,\text{subject}} \sim \text{Normal}(\mu_{\beta 0}, \sigma_{\beta 0}^2)$
$\beta_{1\text{-}3,\text{subject}} \sim \text{Normal}(\mu_{\beta 1\text{-}3}, \sigma_{\beta 1-3}^2)$
$\gamma_{0,\text{subject}} \sim \text{Normal}(\mu_{\gamma 0}, \sigma_{\gamma 0}^2)$
$\gamma_{1,\text{subject}} \sim \text{Normal}(\mu_{\gamma 1}, \sigma_{\gamma 1}^2)$

Modifications to the HDDM code

We modified the HDDM code because there is a problem with the prior on the intercept in the starting point regressions when one is using the inverse logit link function. The HDDM code is set up so that the intercept and other regression coefficients have different prior distributions. The intercept's distribution is based on the distribution of the parameter in the standard (non-regression-based) model, while the other regression coefficients have Normal distributions. In most cases this is fine.

However, there is a problem with starting points. The problem is that starting points range from 0 to 1. Therefore, in the tutorial/documentation, users are instructed to use an inverse-logit link function, which “transforms values between plus and minus infinity into values ranging from (just above) 0 to (nearly) 1.” If a user follows these instructions, they are running the following regression:

$$z = \frac{1}{1 + \exp\left(-\beta_0 + \beta_1 \times X_1\right)}$$

However, because z has a [0,1] distribution, the HDDM code gives β_0 that prior distribution as well. The result is that β_0 is constrained to be in [0,1], when in reality β_0 should be “between plus and minus infinity” with mean = 0, not 0.5. In other words, the intercept in the argument to the inverse logit function is constrained, when instead, it is the output of the inverse logit that should be constrained to [0,1]. The result is that the “baseline” starting points (i.e., the inverse logit of the intercept), will always be above 0.5.

There are multiple solutions to this problem. One of which is to remove the inverse logit prior on the starting point intercept and replace it with a Normal. This way, the untransformed beta will not be constrained from 0 to 1. A second possibility is keeping the inverse logit prior on the starting point intercept but use a linear link function. We chose the first solution, as shown in the description of our priors.

Experiment 1 – Drift bias components (Figure S1)

We conducted an analysis, which was not pre-registered, to test whether the store influenced the weight on the healthy food relative to the unhealthy food ($\frac{\beta_1}{\beta_2}$) in Equation 2. We did find a difference in the relative weight between stores (Healthy Better store: $\frac{\beta_1}{\beta_2}$ = 0.986 (0.855, 0.111); Unhealthy Better store: $\frac{\beta_1}{\beta_2}$ = 0.735 (0.632, 0.839)). But because the relative weight ($\frac{\beta_1}{\beta_2}$) and the intercept (β_0) pointed in opposite directions, we conducted an additional analysis, combining the relative weight and the intercept to test whether the average total drift rate differed between the Healthy Better and Unhealthy Better stores. To do so, we calculated the total drift rate for average-valued healthy and unhealthy foods. In other words, we used values of 5 (the average rating for a food) for both items in the drift rate formula (Equation 2). Here, HDIs for the average total drift rates heavily overlapped (Healthy Better store: v = 0.163 (-0.150, 0.446); Unhealthy Better store: v = -0.187 (-0.486, 0.139)), suggesting that the evaluation process did not credibly differ between the Healthy Better and Unhealthy Better stores (see Figure S1A). Moreover, the difference between the distributions contains 0, further indicating no credible difference in the average drift rates (Δv = 0.351 (-0.087, 0.783)).

Next, we examined the components of the drift rate for both sets of subjects (see Figure S1B). The weight on healthy relative to unhealthy remained the same for subjects who entered the Healthy Better store first (Healthy Store: $\frac{\beta_1}{\beta_2}$ = 0.923 (0.762, 1.07); Unhealthy Store: $\frac{\beta_1}{\beta_2}$ = 0.665 (0.520, 0.815)) and for those who entered the Unhealthy Better store first (Unhealthy Store: $\frac{\beta_1}{\beta_2}$ = 0.832 (0.698, 0.978); Healthy Store: $\frac{\beta_1}{\beta_2}$ = 1.07 (0.841, 1.31)).

Then, we confirmed that the average drift rate did not change from Phase 1 to Phase 2, irrespective of the store the subject entered first. For subjects who entered the Healthy Better store in Phase 1, their average drift did not change across stores (Healthy Store: v = 0.113 (-0.137, 0.351); Unhealthy Store: v = 0.459 (-0.0345, 0.927)), as well as for subjects who entered the Unhealthy Better store first (Unhealthy Store: v = 0.287 (0.041, 0.557); Healthy Store: v = 0.385 (-0.104, 0.858)).

Experiment 2 - Drift sensitivity

We explored how the conditions influenced drift sensitivity, i.e. the coefficient on *Value Difference* in Equation 5 (β_1). The coefficient on *Value Difference* was greater in the Proportion ($\beta_1 = 0.442$ (0.385, 0.503)) and Demand conditions ($\beta_1 = 0.445$ (0.41, 0.481)) relative to the Coin (Coin_A: $\beta_1 = 0.27$ (0.233, 0.318); Coin_B: $\beta_1 = 0.278$ (0.25, 0.307)) and Double (Double_A: $\beta_1 = 0.211$ (0.177, 0.246); Double_B: $\beta_1 = 0.26$ (0.235, 0.282)) conditions. This indicates that in the Coin and Double conditions, the uncorrected value difference did not play as great of a role in determining drift. The largest coefficients in the Proportion and Demand conditions were as hypothesized, while the lack of difference between the Coin and Double conditions differed from what we expected based on the logistic regression results.

Time Trend Analyses

We modelled z as a function of trial and used the inverse logit link function to restrict z to values between 0 and 1:

$$z = \frac{1}{1 + \exp\left(-\gamma_0 + \ \gamma_1 \times Trial\right)}$$

Here, a non-zero γ_0 reflects a bias towards the target (positive) or non-target (negative) at the start of the condition. A positive γ_1 indicates that the starting point bias is increasing with time, and a negative γ_1 indicates that it is decreasing with time.

Experiment 1 Time Trends Model Comparison

We conducted a model comparison exercise between the following 4 models, specified below:

1. A linear relationship between predisposition and trial
2. A logarithmic relationship between predisposition and trial
3. A quadratic relationship between predisposition and trial (in addition to the linear relationship)
4. A logarithmic relationship between predisposition and trial (in addition to the linear relationship)

Model 1

$$v = \beta_0 + \beta_1 \times Healthy\ Rating - \beta_2 \times Unhealthy\ Rating + \beta_3 \times Trial$$

$$logit(z) = \gamma_0 + \gamma_1 \times Trial$$

Model 2

$$v = \beta_0 + \beta_1 \times Healthy\ Rating - \beta_2 \times Unhealthy\ Rating + \beta_3 \times \log(Trial)$$

$$logit(z) = \gamma_0 + \gamma_1 \times \log(Trial)$$

Model 3

$$v = \beta_0 + \beta_1 \times Healthy\ Rating - \beta_2 \times Unhealthy\ Rating + \beta_3 \times Trial + \beta_4 \times Trial^2$$

$$logit(z) = \gamma_0 + \gamma_1 \times Trial + \gamma_2 \times Trial^2$$

Model 4

$$v = \beta_0 + \beta_1 \times Healthy\ Rating - \beta_2 \times Unhealthy\ Rating + \beta_3 \times Trial + \beta_4 \times \log(Trial)$$

$$logit(z) = \gamma_0 + \gamma_1 \times Trial + \gamma_2 \times \log(Trial)$$

(A)

Experiment 1 DIC		
Model	Healthy Better Store	Unhealthy Better Store
1	29811.7	29907.7
2	29801.4	29913.0
3	**29793.5**	29906.2
4	29795.8	**29899.4**

(B)

Experiment 1 Parameter Estimates for Best Fitting Model (95% HDI)		
Parameter	Healthy Better Store (Model 3)	Unhealthy Better Store (Model 4)
γ_0	-0.143 (-0.263, -0.028)	0.190 (0.010, 0.363)
γ_1	0.009 (0.006, 0.012)	-0.002 (-0.004, -0.001)
γ_2	-0.00003 (-0.00005, -0.00002)	-0.060 (-0.120, -0.003)

Table S2. Experiment 1 model comparison results (A) and parameter values for the best fitting models (B). A lower DIC indicates a better fit. We find that for the Healthy Better Store, Model 3 fits the data best and that for the Unhealthy Better Store, Model 4 fits the data best. Overall, Model 4 is the best model, though by a small margin. None of the models substantially outperform the simple linear model. In the Healthy Better store, the coefficient on the linear trial term is positive, indicating that the starting point bias increases towards the healthy option with trial number. However, the coefficient on the quadratic term is negative. This indicates a flattening of the curve later in the block, as can be seen in Figure S11. In the Unhealthy Better store, the coefficient on the linear term is negative, indicating that the starting point bias increases towards the unhealthy option with trial number. The coefficient on the logarithmic term is negative, again indicating a flattening of the curve later in the block, as can be seen Figure S11.

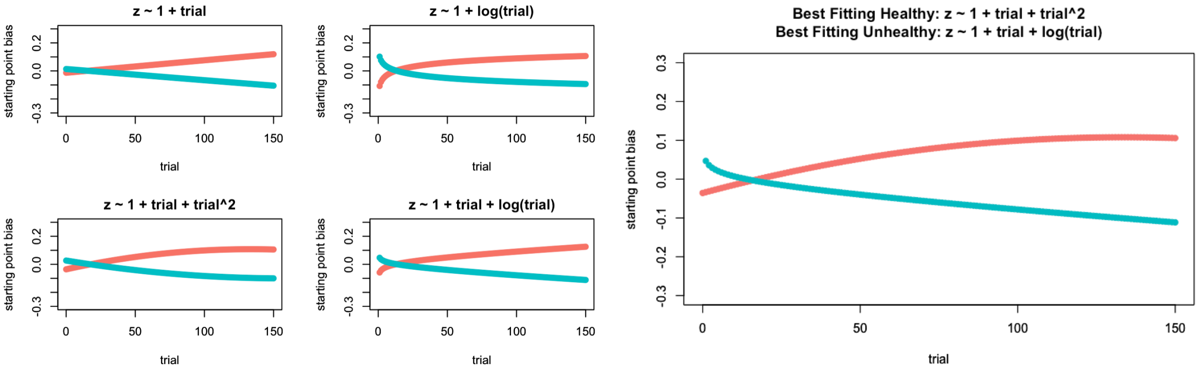

Figure S11. Experiment 1 graphs for all starting point models (left) and the best fitting starting point models (right). Red is Healthy Better, and blue is Unhealthy Better.

Experiment 2 Time Trends Model Comparison

We conducted a model comparison exercise between the following 4 models, specified below:

1. A linear relationship between predisposition and trial
2. A logarithmic relationship between predisposition and trial
3. A quadratic relationship between predisposition and trial (in addition to the linear relationship)
4. A logarithmic relationship between predisposition and trial (in addition to the linear relationship)

Model 1

$$v = \beta_0 + \beta_1 \times Value\ Difference + \beta_2 \times Target\ Value + \beta_3 \times Trial$$

$$logit(z) = \gamma_0 + \gamma_1 \times Trial$$

Model 2

$$v = \beta_0 + \beta_1 \times Value\ Difference + \beta_2 \times Target\ Value + \beta_3 \times \log(Trial)$$

$$logit(z) = \gamma_0 + \gamma_1 \times \log(Trial)$$

Model 3

$$v = \beta_0 + \beta_1 \times Value\ Difference + \beta_2 \times Target\ Value + \beta_3 \times Trial + \beta_4 \times Trial^2$$

$$logit(z) = \gamma_0 + \gamma_1 \times Trial + \gamma_2 \times Trial^2$$

Model 4

$$v = \beta_0 + \beta_1 \times Value\ Difference + \beta_2 \times Target\ Value + \beta_3 \times Trial + \beta_4 \times \log(Trial)$$

$$logit(z) = \gamma_0 + \gamma_1 \times Trial + \gamma_2 \times \log(Trial)$$

(A)

Experiment 2A DIC			
Model	Coin_A	Double_A	Proportion
1	19279.1	**23953.1**	**24254.8**
2	**19263.8**	23955.1	24258.9
3	19268.9	23954.9	24257.4
4	19266.4	23957.5	24256.4

(B)

Experiment 2B DIC			
Model	Coin_B	Double_B	Demand
1	27444.7	**28900.6**	32423.3
2	**27438.9**	28924.8	**32423.2**
3	27441.9	28904.2	32423.4
4	27440.7	28902.3	32427.5

(C)

Experiment 2A Parameter Estimates for Best Fitting Model (95% HDI)			
Parameter	Coin_A (Model 2)	Double_A (Model 1)	Proportion (Model 1)
γ_0	0.205 (-0.022, 0.436)	-0.015 (-0.092, 0.068)	0.009 (-0.061, 0.088)
γ_1	-0.099 (-0.138, -0.062)	0.0006 (-0.0003, 0.002)	0.002 (0.0014, 0.003)

(D)

Experiment 2B Parameter Estimates for Best Fitting Model (95% HDI)			
Parameter	Coin$_B$ (Model 2)	Double$_B$ (Model 1)	Demand (Model 2)
γ_0	-0.350 (-0.495, -0.213)	-0.062 (-0.124, -0.008)	-0.060 (-0.190, 0.074)
γ_1	0.132 (0.103,0.164)	0.0036 (0.0030, 0.0043)	0.014 (-0.017, 0.044)

Table S3. Experiment 2A model comparison results (A) and parameter values for the best fitting models (C). Experiment 2B model comparison results (B) and parameter values for the best fitting models (D). A lower DIC indicates a better fit. In the Coin conditions, the logarithmic models (Model 2) fit the best. In the Coin$_A$ condition, earlier trials more strongly influence the starting point away from the target. And in the Coin$_B$ condition, earlier trials more strongly influence the starting point towards the target. In the Double conditions, the linear models (Model 1) fit the best. For Double$_A$, there is no change in the starting point with trial number. And for Double$_B$, the starting point shifts toward the target linearly with trial number. In the Proportion condition, the best-fitting model is the linear model, and the starting point shifts toward the target linearly with trial number. In the Demand condition, the logarithmic model fits slightly better than the linear model, but the parameter values suggest there is no change in starting point with trial number. Again, the improvements of the non-linear models over the linear ones are typically quite small, indicating that the linear model is a reasonable approximation (see Figure S12).

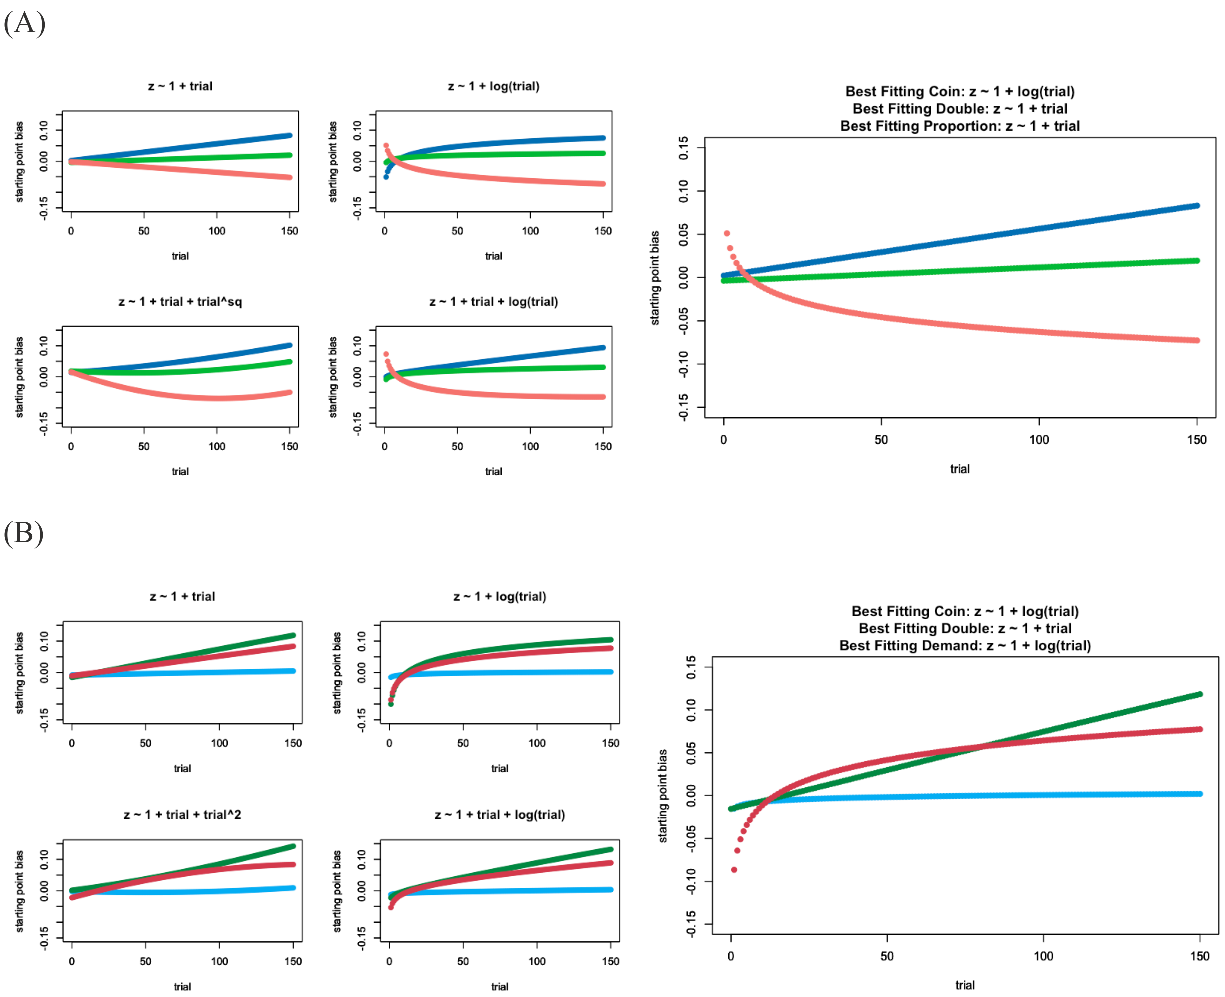

Figure S12. Experiment 2A (A) and 2B (B) graphs for the starting point models (left) and the best fitting starting point models (right). Red is Coin, green is Double and blue is Proportion (A)/Demand (B).

Deviations from Pre-Registrations:

Experiment 1:

There were no deviations from the pre-registration (https://aspredicted.org/blind.php?x=h8er63).

Experiment 2:

We made two deviations from our pre-registered exclusion criteria (https://aspredicted.org/blind.php?x=vv39xx):

1. We pre-registered: "For each individual, we will eliminate trials in which the RT is three standard deviations above the individual mean RT."

   Instead, we removed RTs > 10 seconds and < 1 second. We chose a more lenient threshold during data preparation because we used a feature in HDDM to account for outliers. This feature fits a mixture model where 5% of the data do not inform the parameter estimates.

2. We also removed subjects that chose the same option > 95% of the time. This exclusion criterion was not pre-registered. We applied this criterion because, in these cases, there were too few non-target choices to estimate the model.