

**Supplementary Material to accompany ‘Value-Based Decision Making: An Interactive  
Activation Perspective’ by Suri, Gross, and McClelland**

**Contents:**

1. MATLAB `getnet()` and `update()` routines that constitute the core of all simulations
2. Algorithmic details supplementary to Section 3.2
3. Simulation #4a: *The Anchoring Effect Depends on Attention*
4. Simulation #9: *Instructional Cue Effects*
5. Simulation #10: *Longer visual fixation increases choice rates*
6. Alternative network topologies for Simulation 8

## 1. MATLAB `getnet()` and `update()` routines that constitute the core of all simulations

The MATLAB routine `getnet`, computes the net input for each pool in the network. Specifically, for each pool, the `getnet` routine first accumulates the excitatory and inhibitory inputs from other units, then adds them to the scaled external input and scales them (by  $\alpha$  or  $\gamma$ ) to obtain the net input. Refer to Section 3.2 of the manuscript.

```

for i=1:numpools
    pool(i).inhibition = zeros(1, pool(i).nunits);
    pool(i).excitation = zeros(1, pool(i).nunits);

    for p = 1:length(pool(i).sendingpools)
        sp = pool(i).sendingpools(p);

        positive_acts_indices = find(pool(sp).activation > 0);

        if ~isempty(positive_acts_indices)
            for k = 1:length(positive_acts_indices)
                index = positive_acts_indices(k);
                wts = transpose(pool(i).proj(p).weight(:, index));

                pos = find(wts>0);
                neg = find(wts<0);

                pool(i).excitation(pos) = pool(i).excitation(pos) +
                    pool(sp).activation(index) * wts(wts>0);

                pool(i).inhibition(neg) = pool(i).inhibition(neg) +
                    pool(sp).activation(index) * wts(wts<0);

            end
        end
    end

    pool(i).excitation = pool(i).excitation * alpha;
    pool(i).inhibition = pool(i).inhibition * gamma;

    pool(i).netinput = pool(i).excitation + pool(i).inhibition + estr *
        pool(i).extinput;

end

```

The MATLAB routine `update`, updates the activation values for each unit. Refer to Section 3.2 of the manuscript

```

for i = 1:numpools

    pns = pool(i).netinput>0;

```

```

if (~isempty(pns))
    pool(i).activation(pns) = pool(i).activation(pns)...
        + (max- pool(i).activation(pns)).*pool(i).netinput(pns)...
        - decay*(pool(i).activation(pns) - rest);
end

nps=~pns;
if (~isempty(nps))
    pool(i).activation(nps) = pool(i).activation(nps)...
        + (pool(i).activation(nps) -min).*pool(i).netinput(nps)...
        - decay*(pool(i).activation(nps) - rest);
end

pool(i).activation(pool(i).activation > max) = max;
pool(i).activation(pool(i).activation < min) = min;

end

```

*Parameter values held constant for each simulation:*

```

max = 1.0;
min = -0.2;
rest = -0.1;
decay = 0.1;
estr = 0.4;
alpha = 0.1;
gamma = 0.1;

```

## 2. Algorithmic details supplementary to Section 3.2

We present a quantitative account of competition between units that inhibit each other (McClelland & Rumelhart, 1989). As discussed above the feature units of the input pool, the hidden units, and the action tendency units are inhibitory. We seek to demonstrate that if two units mutually inhibit one another, the one with even slightly greater net input will dominate the other over time.

Consider two mutually inhibiting units  $u_i$  and  $u_j$  with activations  $a_i$  and  $a_j$  greater than zero and  $a_i > a_j$ . Assume that the consolidated excitatory input into the two units is  $e_i$  and  $e_j$ , respectively. Recalling that gamma ( $\gamma$ ) is the scaling factor inhibition we can write

$$\text{net}_i = e_i - \gamma a_j$$

and

$$\text{net}_j = e_j - \gamma a_i$$

Since  $a_i > a_j$ ,  $\text{net}_i > \text{net}_j$ . This will cause activation of  $u_i$  to keep increasing relative to the activation of  $u_j$ . This has been referred to as the “rich get richer” effect (Grossberg, 1976).

### 3. Simulation #4a: *The Anchoring Effect Depends on Attention*

The experiment below (not in the value-based domain) shows that increasing levels of attention towards an incidental and irrelevant variable (a participant identification number) can result in increasing levels of influence of that variable. We used the same network structure and dynamics as described in Simulation 4 in the main text (which was used to simulate an anchoring empirical context in the value-based domain).

*Target Experiment(a)*: In the anchoring conditions, experimenters (Wilson, Houston, Etling, & Brekke, 1996) assigned participants an (allegedly) unique participant identity number and then asked them to estimate the number of physicians listed in the local yellow pages. Almost all participants reported having no prior information helpful in making such an estimate and the few that did were not included in the study.

The participant identity number was always between 1,928 and 1,935. It was not identical to prevent proximal participants from thinking they had the same number. No subsequent differences related to the small differences in identity numbers were noted. Using a cover story, one set of participants (color condition) was asked to identify the color of ink – red or blue – that the identity number was written in (it was written in blue for all participants). Another set of participants was asked to check whether the identity number had four digits (4-digit condition). A third set of participants were asked to identify whether the identity number was greater than 100 or not (GT-100 condition). Still other participants were asked to determine either if their assigned number was greater than 1,920 or if it was less than 1,940. No differences were found between these participants and they were both grouped into the same condition (GT-1920:1940 condition). A final condition asked participants to rewrite their identity number and compare it to the dependent variable – i.e. the number of physicians in the yellow pages (Write and Compare

or WC condition). There was also a control condition in which the participants did not have an identity number.

The experimenters designed these conditions (except for control), with a purpose of ensuring varying levels of attention towards the identity number. They assumed that the color condition, the 4-digit condition, and the GT-100 condition required low levels of attention. In the GT-1920:1940 condition, participants had to attend to their identity number more carefully, because 1920 and 1940 were closer in value to their identity numbers; this required medium levels of attention. Finally, participants in the write and compare (WC) condition had to pay a high level of attention to their identity number, because they had to write their number again and focus on whether it was less than or greater than the number of physicians in the phone book – a judgment that presumably took more thought.

*Empirical Results(a)*: Participant estimates of the number of physicians were unmistakably affected by the presence and type of anchors. In the control condition, the mean estimate was 221. In the low attention conditions mean estimates were slightly higher than control: in the color-condition the estimate was 343; in the 4-digit condition the estimate was 312; in the GT-100 condition, the estimate was 332. These estimates were not reliably different. In the medium attention GT-1920:1940 condition, the mean participant estimate was 527, and in the high attention comparison condition (WC) it was 755.

*Network Structure & Dynamics*: We used the identical network as described in Simulation 4 in the main text. The experience-based estimate ‘est’, (221 physicians) represented the central value of the range represented in the experience pool. Units in the estimate pool represented the following ranges based on ‘est’, the unanchored experienced-based estimate: (0, 0.5\*est), (0.5\*est, 0.9\*est), (0.9\*est, 1.1\*est), (1.1\*est, 1.5\*est), (1.5\*est, 3\*est). The last unit

had a larger range to include the possibility of modeling high estimates. The units in the estimate pool represented identical quantities to corresponding units in the experience pool. The identity number (between 1,920 and 1,940) was assumed to activate the 5th unit in the anchor pool, since it is a large number compared to commonly encountered numbers in similar contexts (Dehaene & Mehler, 1992).

The input activation in the 5th unit was assumed to be proportional to the level of attention elicited by the experimental manipulations: in the low attention conditions (color and GT-100), the input activation was 0.1; in the medium attention condition (GT-1920:1940), the input activation was 0.3; in the high attention condition (direct comparison to re-written identity number), the input activation was 1. In the control condition there was no activation to any unit other than the central estimate unit (in particular, no anchor units were activated).

*Simulation Results:* The winning unit (i.e. the unit in the estimate pool with the largest activation) was used to determine the range of the final estimate (e.g. if the second unit was the winner, then the final estimate was assumed to be between 0.5 and 0.9 times the original, experience-based estimate). To compute the precise estimate value within the range, we calculated the cumulative density function (cdf) of the activation value of the unit on a normal distribution (mean 0.2, standard deviation 0.05) kept fixed for all estimates in the present experiment and in the anchoring experiment described in the main text. For example, a cdf value of 0.5 in the winning second unit represented a value equal to 50% of the range represented by the second unit (i.e.  $0.7 * \text{est}$ ).

In the control condition, the simulation estimate was 218 (vs. 221 in the empirical data). In the low attention condition, the simulation estimate was 346 (vs. an average of 330 in the empirical data). In the medium attention condition, the simulation estimate was 564 (vs. an

empirical estimate of 527), and in the high attention condition, the simulation estimate was 663 (vs. an empirical estimate of 775).

*Significance:* This simulation captured an important element of the anchoring phenomenon: increased attention toward an anchor increases its effect on an estimate. This was a consequence of increased input activation. This element was demonstrated using the same network used in the main text to simulate an occurrence of anchoring in the value-based domain.



#### 4. Simulation #9: *Instructional Cue Effects*

*Target Experiment:* In a behavioral task, experimenters manipulated taste and health cues related to food items (Hare, Malmud, & Rangel, 2011). Participants rated and made decisions on 180 different food items, including junk foods (e.g., chips and candy bars) and healthier snacks (e.g., apples and broccoli). On every trial, subjects were shown a picture of one of the food items and were given up to three seconds to indicate whether they wanted to eat that food at the end of the experiment. Participants indicated their responses using a four-point scale: Strong No, No, Yes, and Strong Yes. At the end of the experiment, choices from a random trial were used to determine whether the participant was required to eat the food item or not.

Participants were asked to make choices in trials preceded by three visual-cue conditions. In the health condition (HC), they were asked to consider the healthiness of the food before choosing. In the taste condition (TC), they were asked to consider the taste of the food before choosing. In the control or natural condition (NC), they were asked to consider whatever features of the food came naturally to mind. Importantly, the instructions emphasized that subjects should always make the decision that they prefer, regardless of the condition.

*Empirical Results:* Affirmative responses (strong yes or yes) for each item/cue combination are shown in Figure SM1: first, without cues (i.e. in the NC condition), un-tasty items were valued considerably less than tasty items regardless of their healthiness. Second, the effect of the health cue increased selection rates for healthy-untasty items and decreased selection rates for unhealthy-tasty and unhealthy-untasty items. No effect of the health cue was observed for the healthy-tasty condition. Third, there were no significant effects of the taste-cue.

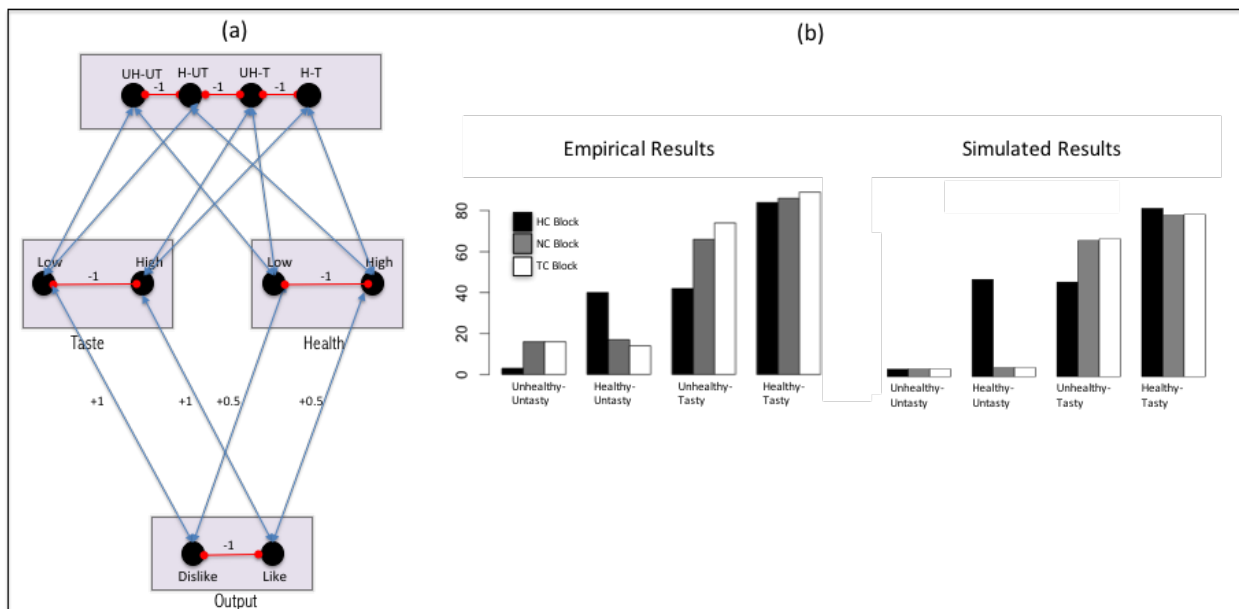
*Network Structure:* We used the following network structure specified in Figure SM1a. The hidden pool had four units corresponding to items that are unhealthy-untasty (UH-UT),

healthy-untasty (H-UT), unhealthy-tasty (UH-T) and healthy-tasty (H-T) (A1). The feature pools corresponding to high/low health and high/low taste were connected to their corresponding hidden units with excitatory (+1) weights (A5). Weights from the features to the action tendency units were excitatory for high taste and high health, and inhibitory for low taste and low health (A6). Taste weights (+1) were set to be higher than health weights (+0.5) (corresponding to the finding that many taste preferences are innate, but health goals are learned and more variable (Ventura & Mennella, 2011)). All weights were reciprocal and were kept identical through-out all conditions.

*Network Dynamics:* The perception of an item involved external input into the network. The perception of the tastiness and healthiness of an item without cues was assumed to be higher for taste features (equal to 0.5 input) than for health features (equal to 0.1 input). We made this assumption based on results showing that taste related features are processed more quickly than health related features and therefore attract greater attention (Sullivan, Hutcherson, Harris & Rangel, 2015). We also assumed that the health cue increased input into units in the health pool from 0.1 to 1, whereas the taste cue increased input into units in the taste pool more modestly – from 0.5 to 0.6 (A7). We made this assumption based on results showing that taste features are generally processed more automatically compared to health related features; the latter often require attention to be processed and may thus be more influenced by external cues (Marteau, Hollands, & Fletcher, 2012). Thus, a taste cue may not have a large incremental effect, but a health cue might.

The model ran through 300 cycles, before which all unit activations had converged (A8). We used output unit (i.e. ‘like’ output unit) activation levels to compute the probability of an item being rated as liked.

*Simulation Results:* In each condition-cue combination we treated activation of the ‘like’ output unit as a point on a normal distribution (mean = 0.4; standard deviation = 0.3), and used the point’s associated probability density functions to compute the probability of being liked. (Figure SM1b) (A9). The simulation captured the main features of the empirical results: items that were both healthy and tasty were the most likely to be liked, and items that were neither were least likely to be liked. Untasty-healthy items were liked in the presence of the health cue, but not otherwise; tasty unhealthy items were liked – but the liking level decreased in the presence of the health cue.



**Figure SM1:** The network and simulation results for Simulation 9: Attention Towards Features Increases their Impact. Panel (a) depicts the network structure in which a taste feature pool and a health feature pool are connected to a hidden pool (with uniform weights) and an output pool with respective weights of +1 and +0.5. All connections within the pool were inhibitory (not fully shown in the hidden pool to promote readability). Panel (b) depicts the percentage probability that a food item would be liked. The same model was able to simulate all four conditions of the experiment – Unhealthy-Untasty, Healthy-Untasty, Unhealthy-Tasty, and Healthy Tasty. The effect of the health and taste cue was modeled as an increase in external activation into the health and taste pools.

*Significance:* In this experiment, participants were encouraged to select items according to their preferences – and only their preferences. Yet, health cues unambiguously affected decisions. This simulation shows a mechanism by which activation in input units (caused by the

presence of cues), can change patterns of choice without any changes to the action-tendency related knowledge represented by the connections between feature units and output units.

## 5. Simulation 10: Longer visual fixation increases choice rates

*Target Experiment(s):* Experiments that exogenously and explicitly manipulated fixation durations during the process of choice have shown that the willingness-to-pay for items increases significantly with increased fixation. In a typical such experiment (Armel, Beaumel & Rangel, 2008, Study 3), participants were asked to provide liking ratings for posters (on a 100 point scale). A computer algorithm then chose two items that had liking ratings within five points of each other. The pictures of the two items then alternated, one at a time, with one presented for 300ms, and the other for 900ms on each alternation. Alternations were repeated six times. The participant was then asked to choose one of the two items. The participant was awarded her choice from a randomly chosen trial.

*Empirical Results:* Liking differences varied between -5 and +5. In particular, the item shown for 900ms could be more liked (by up to +5 points) or less liked (by up to -5 points) than the item shown for 300ms, and vice versa. Across all liking levels, items that were viewed for a longer time, were more likely to be chosen than items that were viewed for a shorter time (Table SM1). For example, when an item was disliked by -5 points, it was chosen in 25% of trials when it was viewed for 900ms, but only in 17% of trials when it was viewed for 300ms. When items were neither liked nor disliked relative to their choice, items shown for 900ms were chosen in 55% of trials, but items shown for 300ms were chosen in 45% of trials. When liked by +5 points, the respective choice rates were 65% to 47%.

*Network Structure:* We created a network with a hidden pool, a feature pool, and an output pool. The hidden pool consisted of 5 inhibiting units that represented items that were very liked, liked, neutral, disliked, and very disliked. The feature pools consisted of 2 sub-pools. One pool consisted of 5 features that were more or less liked. The first of these units was connected to

the ‘very liked’ hidden unit, the second to the ‘liked’ hidden unit, the third with the ‘neutral’ hidden unit, the fourth with the ‘disliked’ hidden unit, and the fifth with the ‘very disliked’ hidden unit. Each of these weights was set at +1. These feature units were also connected with a single approach output unit with the following respective weights: 0.5, 0.4, 0.3, 0.2 and 0.1 (least liked). These weights represented differences in liking ratings between items.

The second feature sub-pool consisted of a single unit that was connected to all hidden units representing the common features across all items (weight +1). It was connected with the output unit with a weight of 0.5.

*Network Dynamics:* The short presentation level (300ms) was simulated by providing feature input for 15 cycles, whereas the long presentation level (900ms) was simulated by providing feature input for 45 cycles. For items seen for 900ms, increased input, caused by longer presentation times, increased the activation build-up in the approach unit resulting in an overall preference for these items over items seen for 300ms at the same liking level.

<i>Liking</i>	-5		-2		0		+2		+5	
	<i>Empirical</i>	<i>Sim.</i>	<i>Empirical</i>	<i>Sim.</i>	<i>Empirical</i>	<i>Sim.</i>	<i>Empirical</i>	<i>Sim.</i>	<i>Empirical</i>	<i>Sim.</i>
	%	%	%	%	%	%	%	%	%	%
<i>900ms</i>	30	31	47	44	58	57	60	69	65	80.1
<i>300ms</i>	18	20	38	30	47	43	50	56	49	69

*Simulation Results:* We ran parallel copies of networks at different levels of liking. For example, corresponding to a liking disadvantage of -5 points for the longer-viewed items, we ran one network with input into the least liked feature unit, and the common unit for 45 cycles, and another network with input into the most liked feature unit, and the common unit for 15 cycles.

We then compared the difference in activation in the respective approach units. For other liking differences we averaged across all possible combinations corresponding to that particular difference. Simulation results matched the pattern seen in the empirical data. Table SM1 summarizes the empirical results and simulation data for liking differences of +5, +2, 0, -2, and -5.

*Significance:* The simulation provides an activation-based mechanism that is consistent with the empirical observation that increased viewing time increases choice rates at all levels of liking. Increased viewing time was simulated by increased activation.

## 6. Alternative network topologies for Simulation 8

The IAC model specifies separate categories for high and low quantities of features, rather than a single unit that activates more or less strongly for a given dimension. Here we briefly describe the reason a single unit architecture is not a viable model in the context of the IAC framework.

A core principle of all IAC networks is that a feature unit can receive activation from other feature units (via hidden units) that it frequently co-occurs with (Assumption 1). For example, a ‘high health’ unit may get indirect activation from the ‘low sugar’ feature unit and a ‘vegetable logo’, because those properties may frequently travel together. In this example, all three feature units (high health, low sugar, and vege logo) would be connected to a common hidden unit.

This principle makes an architecture that relies on different activation levels in a single unit incompatible with the principles of interactive activation. If there were a single unit for health, for example, it could not co-activate with units representing co-occurring features. In particular, assume that such a unit is connected with the ‘vege logo’ unit but not the ‘cola logo’ unit. In this case, evidence for the drink being a cola would not provide evidence for the property of low health. Conversely, if such a unit was connected with the ‘cola logo’ unit but not the ‘vege logo’ unit, the sight of a logo featuring vegetables would not provide evidence for the property of high health.

This feature of the IAC framework may explain why single-cell recordings in putative value circuits often show co-mingling of neurons with "on" responses to value (i.e. higher response to increasing value) and neurons with "off" responses to value (higher responses to decreasing value) (Morrison & Salzman, 2009). It is possible that such neurons could be



representing two ends of a continuum in order to enable feature associations that are context sensitive.