

# Psychological Review

## Value-Based Decision Making: An Interactive Activation Perspective

Gaurav Suri, James J. Gross, and James L. McClelland

Online First Publication, September 16, 2019. <http://dx.doi.org/10.1037/rev0000164>

### CITATION

Suri, G., Gross, J. J., & McClelland, J. L. (2019, September 16). Value-Based Decision Making: An Interactive Activation Perspective. *Psychological Review*. Advance online publication. <http://dx.doi.org/10.1037/rev0000164>

# Value-Based Decision Making: An Interactive Activation Perspective

Gaurav Suri  
San Francisco State University

James J. Gross and James L. McClelland  
Stanford University

Prominent theories of value-based decision making have assumed that choices are made via the maximization of some objective function (e.g., expected value) and that the process of decision making is serial and unfolds across modular subprocesses (e.g., perception, valuation, and action selection). However, the influence of a large number of contextual variables that are not related to expected value in any direct way and the ubiquitous reciprocity among variables thought to belong to different subprocesses suggest that these assumptions may not always hold. Here, we propose an interactive activation framework for value-based decision making that does not assume that objective function maximization is the only consideration affecting choice or that processing is modular or serial. Our framework holds that processing takes place via the interactive propagation of activation in a set of simple, interconnected processing elements. We use our framework to simulate a broad range of well-known empirical phenomena—primarily focusing on decision contexts that feature nonoptimal decision making and/or interactive (i.e., not serial or modular) processing. Our approach is constrained at Marr's (1982) algorithmic and implementational levels rather than focusing strictly on considerations of optimality at the computational theory level. It invites consideration of the possibility that choice is emergent and that its computation is distributed.

**Keywords:** value-based decision making, interactive activation, neural networks, computer simulation, parallel distributed processing

**Supplemental materials:** <http://dx.doi.org/10.1037/rev0000164.supp>

What options do we prefer and how do we determine our preferences? One common answer is that we review the available options, calculate the expected value of each option, and then act to maximize our expected value (Atkinson, 1957; Boureau, Sokol-Hessner, & Daw, 2015; Eccles & Wigfield, 1995; Rangel, Camerer, & Montague, 2008; Simon, 1986; Slovic, 1995).

This way of thinking about value-based decision making is both intuitive and widespread. However, a growing body of empirical results is inconsistent with this view, and this leads us to propose an alternative. We refer to this as the interactive activation perspective. It is rooted in the parallel distributed processing (PDP) tradition (Rumelhart & McClelland & the PDP Research Group, 1986), and more specifically, in the assumptions of the interactive activation and competition (IAC) model introduced in McClelland and Rumelhart (1981).

## Summary of the Argument

There are two interrelated questions that are central to the value-based decision-making literature: *what* options do decision makers prefer, and *how* do they compute their preferences? A commonly accepted claim related to the “what” question is that decision makers prefer options that maximize expected value. A commonly accepted claim related to the “how” question is that for each decision it faces, the brain represents the available options, values them, and then acts to choose the option with the highest valuation. These subprocesses are thought to be largely serial and modular. We examine each of these claims in turn.

## Maximizing Expected Value

Researchers in different eras and in different domains have postulated that value-based decision making involves maximizing expected value (Huygens, 1657; Laplace, 1814; Samuelson, 1937; von Neumann & Morgenstern, 1944). The *expected value of an option* is defined as the product of the net benefits offered by that option (its value) and the probability of its realization (its expectation). It represents the average value yielded by an option if it is repeatedly encountered.

However, in many empirical studies, decision makers do not always act to maximize expected value (Bell, 1988). In response to these studies, researchers proposed that decision makers maximize subjective expected value. Subjective expected value allows for the subjective evaluation of the expectations and values of the available options. Subjective variables (i.e., expectation and value)

---

Gaurav Suri, Department of Psychology, San Francisco State University; James J. Gross and James L. McClelland, Department of Psychology, Stanford University.

We thank Ezequiel Morsella for his comments on a prior version of this article.

Correspondence concerning this article should be addressed to Gaurav Suri, Department of Psychology, San Francisco State University, 1600 Holloway Avenue, San Francisco, CA 94132. E-mail: [suri@sfsu.edu](mailto:suri@sfsu.edu)

are thought to be functions of the corresponding objective variables. This insight informed Prospect Theory (Tversky & Kahneman, 1981), which proposed that (a) subjective representations of probabilities were less extreme than corresponding objective probabilities and (b) representations of subjective value were nonlinear and asymmetric functions of objective value. Thus, Prospect Theory and related fields such as behavioral economics (Camerer, Loewenstein, & Rabin, 2011) did not contest the assumption of expected value maximization but did note that the representations of expected value that people actually used were functions of objective expected value.

Contemporary researchers have continued to embrace the notion that subjective expected value is the common currency of value-based decisions. Specifically, Bayesian decision-making theories (e.g., Peterson, 2017; Pooeh, Bernhardt, Guevara, Huys, & Smolka, 2018) have focused on the “expected” part of expected value calculations, and psychologists interested in motivation have focused on the role of needs and goals in shaping the “value” part of expected value calculations (e.g., Locke & Latham, 2006).

However, there is a growing class of apparent anomalies in which it appears doubtful whether decision makers are, in fact, maximizing subjective value (or any systematic function of expected value). For example, in the anchoring phenomenon (Ariely, Loewenstein, & Prelec, 2003; Orr & Guthrie, 2005), people are known to be systematically influenced by irrelevant anchors introduced prior to an estimate of the willingness to pay for a set of items. In one such study, participants were asked to attend to the last two digits of their social security number. Participants with higher numbers (e.g., 93) had a greater willingness to pay than participants with lower numbers (e.g., 15; Ariely et al., 2003; Orr & Guthrie, 2005). Another example of an apparent anomaly features the phenomenon of negative auto-shaping (Williams & Williams, 1969). In one such study, pigeons observed pairings between an illuminated key and the delivery of food. At first, they ignored the illuminated key and pecked at the delivered grain. Later in the experiment, the pigeons began to peck at the illuminated key even though such behavior, in the context of the experiment, prevented the delivery of food. In this article, we describe and simulate several such anomalies.

There are two reasonable responses to such apparent anomalies. The first is to propose that the above examples represent systematic—but as yet undiscovered—deviations in the computation of subjective expected value from objective expected value. Thus (subjective) expected value would remain the common currency of value-based decision making—albeit its computation would be even more complex than originally thought.

The second response is to consider the possibility that expected value computations are not the common currency of value-based decision making. Such a response would entail the development of a framework that cannot only explain the types of anomalies described above but can also include a viable explanation for the instances of decisions in which decision makers do appear to be maximizing expected value. In this article, our objective is to develop the second of these two responses.

## Serial and Modular Processing

Traditional models in economics limited their scope to explaining an individual’s observable choice. They did not seek to address

the difficult to observe brain processes that produced a particular choice. They allowed that in value-based decision making, humans behaved as if they had valued all available options but remained agnostic whether this “as if” computation actually occurred in the brain (Friedman & Friedman, 1953; Samuelson, 1937).

Modern neuro-economists and psychologists hypothesized that the brain does indeed make subjective expected value calculations (Kable & Glimcher, 2009). They proposed that this calculation is made via a set of serially unfolding, modular processes (Opris & Bruce, 2005; Rangel et al., 2008). The brain is thought to first represent the variables that are relevant to the decision at hand. Next, it is thought to compute value signals related to the variables being considered, and then it is thought to act to select the action possibility with the strongest value signal. Each subprocess is thought to complete itself (i.e., it is modular) and then pass on information to the next subprocess (i.e., it is serial). Although some theorists do not insist on rigidity in serial and modular processing (e.g., Rangel et al., 2008), serial/modular processing is often thought to be a core feature underlying value-based decision making.

The assumption of serial and modular processing has proved conceptually useful. However, there are several anomalies in which the brain does not appear to be following a serial and modular process. For example, participants perceived an ambiguous figure as either a horse or seal depending on whether a farm animal or a sea creature led to more rewards (Balcetis & Dunning, 2006). Similar results were seen when participants viewed a figure that could be perceived either as the letter “B” or as the number “13.” In both these cases, the value of future outcomes appeared to influence what they perceived—suggesting that the processing of visual stimuli interacted with the rewards associated with those stimuli (which were presumably processed in the valuation stage). In this article, we describe and simulate this and other such anomalies.

There are two reasonable responses to this and other such apparent anomalies. The first is to propose that the brain makes decisions in a serial and modular fashion as a fundamental design principle—but does so imperfectly in practice. Any examples featuring the dynamic interaction between putatively modular subprocesses is an outcome of imperfect segregation or leakage among the components.

The second response is to consider the possibility that value-based decision making is an inherently interactive process whose putative components can operate concurrently and influence each other as a fundamental characteristic of their operation. In this article, our objective is to develop the second of these two responses.

## The Interactive Activation Approach

To accomplish our objectives, we will draw on the IAC framework (McClelland & Rumelhart, 1981; McClelland, 1981) that embodies a specific set of assumptions within the broader PDP tradition (Rumelhart et al., 1986). The framework specifies that the neuron-like processing units that constitute a network are organized into pools such that units within a pool are mutually excitatory and connections between pools are mutually inhibitory. The framework adheres to the assumption that propagation of excitatory and inhibitory influences is fundamentally bidirectional; if A

excites or inhibits B, then B excites or inhibits A to the same extent (i.e., units are connected bidirectionally and have symmetrical weights). An important principle of the network is that units representing attributes of items and situations do not excite each other directly; instead, the framework relies upon conjunction units that bind together all of the units that represent the features of an item or situation and mediates their influence on each other.

The principles of interactive activation have been applied in a wide range of psychological phenomena (see the *Situating the Present Work in Prior Literature* section). However, prior to this work, they have not been applied to the development of an integrated alternative to the view that value-based decision making involves the maximization of expected value, and that it occurs via a serial and modular process.

In our application of the framework to value-based decision making, we retain the features of the framework described above. Two extensions are required to address value-based decision making. First, drawing on an extension to the framework introduced in Kumaran and McClelland (2012) we provide for experience-dependent strengthening of the bidirectional excitatory connections between the features of an item or situation and the conjunction unit for it. Second, to address value-based behavior itself, we allow some (but not all) of the units corresponding to features of a situation to be connected to output units representing approach or avoid action tendencies. Similar units have been previously used in the personality domain (e.g., Read et al., 2010).

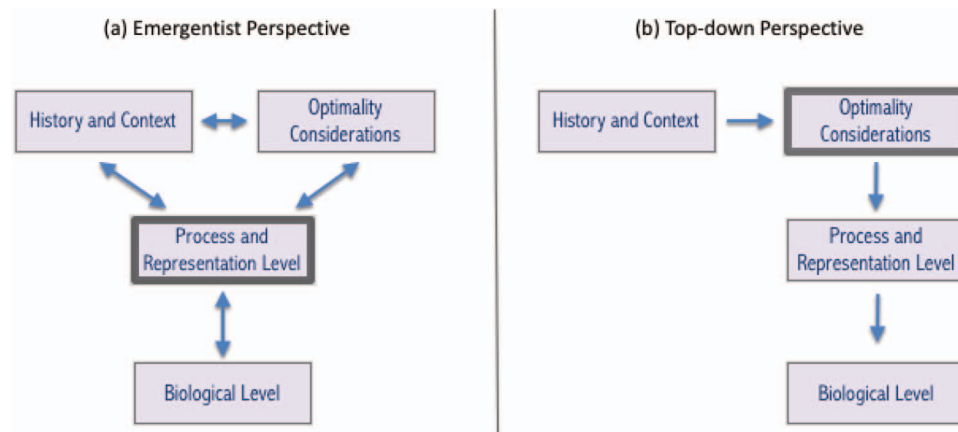
These elements play a crucial role in enabling our model to simulate situations where choice patterns are consistent with value maximization, as well as situations in which choice patterns do not appear to be consistent with value maximization. As an example of the latter, we consider negative auto-shaping described above as a pattern inconsistent with value maximization. We simulated this phenomenon by assuming that repeated pairing between the illuminated key and the delivery of food caused the units representing these events to be connected to each other via a conjunction unit for the key-food pairing situation. After this connection was established, activation in the illuminated key unit caused activation

to flow to the food unit which, in turn, activated the approach output unit (the key unit was not connected to the approach unit, but the food unit was). This caused the pigeon to approach and peck at the illuminated key—even though doing so ensured the nondelivery of food (see Simulation 5).

The IAC framework instantiates an emergentist vision of value-based decision making (and affect and cognition, more generally). It focuses attention at the process and representation level, captured primarily in terms of proposed networks of units and the activation processes that operate within them, with the possibility that these units and processes can ultimately be related to their underlying biological implementation, as illustrated in Figure 1a.

Although the framework is sensitive to optimality considerations, we see it as distinct from, and as an alternative to, the top-down approach to cognition and behavior advocated by those who focus primarily on optimality considerations, or on the computational level (similar to a focus on optimality) advocated by Marr and other proponents, as illustrated in Figure 1b. Instead of treating optimality considerations or what Marr (1982) called the computational level as primary, we place the greatest emphasis at the representation and process (algorithmic) level, while allowing that the representations and processes at this level are interdependent with context and history, optimality/computational-level constraints, and constraints imposed by our biology. With its openness to considerations other than optimization per se, we see our emergentist interactive activation framework as in line with the approach advocated by Rahnev and Denison (2018) in the realm of perceptual decision making, lowering the risk of deploying constructs that stipulate the target of explanation, and opening up a broader consideration of the best way to understand all of the factors that influence value-based decision making.

In reference to the two central “what” and “how” questions about decision making, the IAC framework proposes that (a) although decision making may sometimes be consistent with expected value maximization, at other times it may be nonoptimal due to activation of contextual variables and/or due to associative learning and (b) the process of value-based decision making fea-



**Figure 1.** Two ways of construing levels of description in value-based decision making (and other cognitive and affective domains). In the approach illustrated in Figure 1a, the process and representation level importantly informs and is informed by, other variables in the decision-making process. In the approach illustrated in Figure 1b, optimality considerations are the main driver of behavior. See the online article for the color version of this figure.

tures the reciprocal and interactive (i.e., not serial or modular) influence of many variables acting in parallel.

The rest of this article fleshes out this summary of our argument. We provide a high-level description of well-accepted accounts of value-based decision making and highlight two sets of empirical phenomena that challenge such accounts—that is, (a) examples of choices that do not appear to be guided by expected value maximization and (b) processing that does not appear to involve serial and modular processing. Next, we detail the IAC framework and provide a simulation that illustrates the operation of our framework which is consistent with expected value maximization. We use the interactive activation framework to present simulations of empirical phenomena in which expected value is not maximized and/or whose processing was not serial or modular (i.e., the phenomena described below). In Finally, we draw out several broader implications of the IAC model.

### A Consideration of Two Key Assumptions of Value-Based Decision Making

We define *value-based decisions* as decisions that involve subjective preferences (e.g., Do I want an apple or an orange?). These may be contrasted with perceptual decisions that involve decisions related to an objective state of the world (e.g., Is this an apple or an orange?; Polanía, Krajbich, Grueschow, & Ruff, 2014). Importantly, preferences may sometimes influence choice and sometimes influence evaluations. We see both these processes—evaluations and choice—as flowing from a common set of principles and consider them both to be relevant to the present work.

As described the *Summary of the Argument* section, value-based decision-making theories from neuroeconomics, psychology, and economics have been predicated upon two intuitively plausible assumptions: objective function maximization and serial/modular processing. *Objective function maximization* refers to the idea that decision makers prefer options that maximize expected value. *Serial/modular processing* refers to the idea that for each decision, the brain, in turn, represents the available options, values them, and then acts to choose the option with the highest valuation. Although it remains difficult to adjudicate the final ground truth, a gathering body of evidence is inconsistent with both these assumptions.

### Evidence Inconsistent With the Objective Function Maximization Assumption

Theories of decision making have often assumed that decision outcomes have a set of values and expectations (i.e., beliefs about the probability of an outcome being realized) associated with them, and that the overall utility of a decision outcome is simply its expected value—a probabilistically adjusted average value that could be expected if an option was selected (Glimcher, 2008; Samuelson, 1937; von Neumann & Morgenstern, 1944). It is often further assumed that decision makers select options that maximize their expected value. This view has been embraced and extended by neuro-economists seeking to find the brain correlates of expected value (Glimcher, 2008; Glimcher, Camerer, Fehr, & Pol-drack, 2009), decision scientists seeking to identify predictable irregularities in expected value computations (Tversky & Kahneman, 1981), psychologists seeking to identify the affective sources

of value (Higgins, 2015), and theorists studying decision making in the Bayesian tradition (Peterson, 2017; Pooseh et al., 2018).

In what follows, we review two sets of studies that jointly suggest that decision makers do not always maximize an objective function in order to make judgments and choices. In the first set, we shall consider examples in which behavioral choice *reflects attending to contextual factors that are unrelated* to the value of the options at hand. Here, we will describe effects related to anchoring to irrelevant variables, persistence with inferior default states, and the influence of increased visual gazes on choice. In the second set, we shall consider examples in which behavioral choice appears to emerge from (learned) associative constraints, rather than optimization related constraints. Here, we will describe examples related to negative auto-shaping, and associative learning related to prior rewards.

**Attending to contextual factors influences choice.** The first set of studies relates to the influence of attention to factors irrelevant to optimality. A well-known effect concerns the *anchoring heuristic*. In the classic demonstration by Tversky and Kahneman (1974; not in the domain of value-based decision making), for example, participants were asked whether the percentage of African countries in the United Nations was more or less than 10% (low anchor) or 65% (high anchor). Both anchors were understood by participants to be random, generated by a spinning wheel (rigged to stop at either 10 or 65). Results indicated that the arbitrarily selected number on the wheel affected estimates. Those whose spin had yielded a 10 provided a median estimate of 25% and those who spun a 65 provided a median estimate of 45%. Similar effects were seen in a value-based decision-making context when participants were asked to estimate whether the prices of several items were more or less than the last two digits of their social security numbers (Simulation 4, Ariely et al., 2003).

Some researchers have argued (Barkan, Ayal, & Ariely, 2016) that such examples may be thought of in terms of objective function maximization. It is possible that participants started off with an incidental anchor and then updated the anchor, trading off between the utility of being accurate and the effort of mentally iterating through considerations related to improving the quality of the estimate. However, such accounts do not address why a completely incidental number could serve as an anchor in an unrelated decision in the first place.

Parallel issues arise in the consideration of default effects (Simulation 3). Individuals are known to frequently persist with default options that appear to have inferior valuations to other possible actions in that particular context (Samuelson & Zeckhauser, 1988). Default preferences have been observed in diverse decision contexts such as organ donation (Johnson & Goldstein, 2003), savings behavior (Choi, Laibson, Madrian, & Metrick, 2004), voting patterns (Gow & Eubank, 1984), and choices in utility and insurance providers (Samuelson & Zeckhauser, 1988). Why do such default effects occur? To date, explanations consistent with the assumption of utility maximization have sought to argue that there are subtle benefits that increase the valuation of the default option and/or subtle costs that decrease the valuation of the nondefault options. Past research (Dinner, Johnson, Goldstein, & Liu, 2011; Kahneman, Knetsch, & Thaler, 1991; Samuelson & Zeckhauser, 1988) identified three such valuation-based factors: costs of evaluating the available options, an implied recommendation, and loss aversion associated with leaving a default state.



Although accounts based on these three factors successfully explain some behaviors that may appear counterintuitive or anomalous, there is growing evidence that default effects persist even when none of these three factors is present. In one such study (Suri, Sheppes, Schwartz, & Gross, 2013), when participants were asked to proactively press an easily accessible button to reduce the probability of being painfully shocked (if they did nothing they would be shocked), participants pressed the shock-probability-reducing button in only about half the trials. Postexperiment debriefing revealed that participants universally understood that button pressing was in their interest (thus decision costs were minimal), understood that experimenters were indifferent to whether they pressed the button or not (thus implied recommendations were minimal), and understood that there were no losses associated with leaving the default state of receiving an undesirable shock (thus loss aversion was unlikely to be an influencing factor). And yet, participants persisted with inferior default states for approximately half the trials. Similar effects were reported in a picture switching context (Suri & Gross, 2015) in which participants could avoid watching undesirable pictures (e.g., pictures featuring horrific mutilations), simply by pressing a button. However, they frequently did not avail themselves of this opportunity. Interestingly, drawing attention (via a flashing red border) to a message that stated image switching was possible via a button press, more than doubled switching rates—even though 100% of participants were fully aware that they had this option.

A final group of examples highlighting the importance of contextual factors concerns the influence of increased visual gazes and instructional cues on choice. Researchers have demonstrated that merely increasing the amount of time an option is available to view increases the rate at which that option is chosen. For example, a slightly less-liked item displayed for 900 ms was frequently chosen over a more liked item displayed for 300 ms (Simulation 10, Armel, Beaumel, & Rangel, 2008). Similarly, when asked to choose the more attractive of two faces, manipulation of gaze duration systematically biased observers' preference decisions (Shimojo, Simion, Shimojo, & Scheier, 2003). This gaze effect was also present when participants compared abstract, unfamiliar shapes for attractiveness. Relatedly, increasing the salience of particular features (e.g., health or taste of food items) has been shown to increase the influence of those features on ultimate choice (Hare, Malmaud, & Rangel, 2011). For example, in the presence of health-related instructions, participants were more likely to choose healthy items, even when the instructions emphasized that they should always make the decision that they prefer, regardless of the instruction (Simulation 9).

**Behavioral choice can emerge from (learned) associative constraints.** A second set of studies highlights how nonoptimality may arise due to learned associations. A well-known effect in which a learned associative constraint influences choice has been demonstrated in experiments involving negative auto-shaping. As described in the *Summary of the Argument* section, in a classic demonstration (Williams & Williams, 1969) placed pigeons in a box with a food hopper (which could deliver food) and an illuminated key (which the pigeons could peck on). One of the most important empirical conditions involved the illumination of the key and the delivery of food provided the pigeon did not pick on the illuminated key. Naïve pigeons, who did not have any prior associations between the illumination of the key and the delivery

of the food, at first performed this task effortlessly. They did not pick on the illuminated key and were rewarded with the delivery of food. However, after a few trial blocks, the pigeons spontaneously began to peck on the illuminated key—a behavior that was precisely nonoptimal since it caused the nondelivery of food.

Presumably, this behavior occurred because the pigeons associated the illuminated key with the delivery of the food in the initial trials in which they were not pecking on the illuminated key. After the association was established, the illuminated key caused approach behavior (due to its association with food), which led to pecking. This example suggests that behavior need not exclusively be shaped by valued reinforcements; rather associative constraints can and do shape behavior as well (Simulation 5).

A second example illustrating the impact of associative constraints concerns associations with prior rewards. Rats were placed in a T maze choosing between two courses of action that differed in their energetic demands and consequent reward sizes (Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). Energetic demands were varied based on barriers of different heights (e.g., no barrier, a 30-cm barrier, or a 40-cm barrier). Food rewards were varied based on the number of pellets available to the rat (they varied between two and six pellets). For example, one arm of the maze (the high reward [HR] arm) may offer four pellets and feature a 30-cm barrier, and the other arm (the low reward [LR] arm) would offer two pellets and feature no barrier. The HR side and the LR side were kept the same throughout the experiment. In early blocks, rats, consistent with optimality, preferred arms with more food pellets and lower (or no) barriers.

However, associative (not optimality) constraints came to the fore in the last block of trials. In this block, the HR side featured two pellets and a 30-cm barrier, and the LR side featured (as it did on most prior blocks) two pellets and no barrier. In this block, the rats preferred the HR block (60% for the HR side vs. 40% for the LR side) even though they had to climb a barrier to get two pellets, which were also on offer on the LR side without the barrier. Presumably, this occurred because the rats associated the HR side with a greater number of food pellets, even though this was not the case in the final block of trials. Prior associative learning related to the arm of the maze offering higher rewards prevented the rats from choosing the optimal option (Simulation 6).

## Evidence Inconsistent With Seriality/Modularity in Value-Based Decision Making

The process of value-based decision making is often assumed to consist of serially unfolding, modular subprocesses that include perception, valuation, and action (Gibson, 1979; Kable & Glimcher, 2009; Opris & Bruce, 2005; Rangel et al., 2008). According to these accounts, the value-based decision-making process begins with a subprocess in which the different available options in the decision-making process are perceived and represented (perception and representation). This is followed by a subprocess in which the different action outcomes are valued (valuation). Next, the outcome of the valuation process is implemented (action selection).

Proponents of serial and modular processing do allow for the possibility that the modularity of the subprocesses involved in decision making is not rigid, and that some subprocesses may proceed in parallel. However, at the core of their proposals is the

assumption that the subprocesses are qualitatively and computationally distinct. Each subprocess is thought to accept certain primitives and act on them via a set of ordered rules, to produce outputs that are accepted by the next subprocess. These computations continue until an action or an option is selected. Thus, in some respects, value-based decision making is assumed to be akin to symbolic, serial, and modular processes that are common in cognitive science (Fodor, 1983).

Although the serial and modular view of decision making has provided conceptual benefits, an increasing number of empirical studies suggest that decision making consists of richly interactive processing—so that processing corresponding to one function is constantly influencing and being influenced by processing related to other functions. Collectively, these empirical studies present evidence inconsistent with serial or modular processing between functionally distinct units. Importantly, such studies are not inconsistent with representational modularity (e.g., by brain region) but are inconsistent with process level modularity.

For example, several nonserial/modular effects have been noted in the motivational influences in visual processing literature. As described in the *Summary of the Argument* section, in one of the best known such studies (Balcetis & Dunning, 2006), participants perceived an ambiguous figure as either a horse or seal depending on whether a farm animal or a sea creature led to more rewards for them (Simulation 7). Related studies have demonstrated that valued objects (i.e., those that can fulfill immediate goals such as thirst), are perceived to be closer than they actually are (Balcetis & Dunning, 2010). Similarly, distances that require more effort to traverse are often judged to be longer (Hajnal, Bunch, & Kelty-Stephen, 2014). Relatedly, the probability of successfully performing an action has been shown to influence perceptual estimates of the target's size and speed (Lee, Lee, Carello, & Turvey, 2012; Witt & Sugovic, 2010, 2012). In these cases, participants' perception, valuation, and choices appeared to proceed in parallel and dynamically interact with one other. Finally, gathering neural evidence suggests simultaneous processing of perception-related and valuation-related elements (Hunt et al., 2018).

### The IAC Framework for Value-Based Decision Making

In this section, we develop the IAC framework for value-based decision making that, in addition to allowing for choices in which expected value is maximized, also allows for choices that are not consistent with expected value maximization, and allows for processing to occur in an interactive (i.e., not serial/modular) fashion.

The IAC framework embodies a specific set of principles developed within the PDP tradition (Rumelhart et al., 1986). Neural networks in the PDP tradition assume that all processing occurs within neuron-like elements called units. These units influence each other via weighted connections. All knowledge is resident within these weighted connections. Learning in the network occurs either by creating new connections between units or by updating existing connection weights. The IAC framework for value-based decision making makes these basic assumptions and is additionally constrained by further specifications described below.

### The IAC Network for Value-Based Decision Making

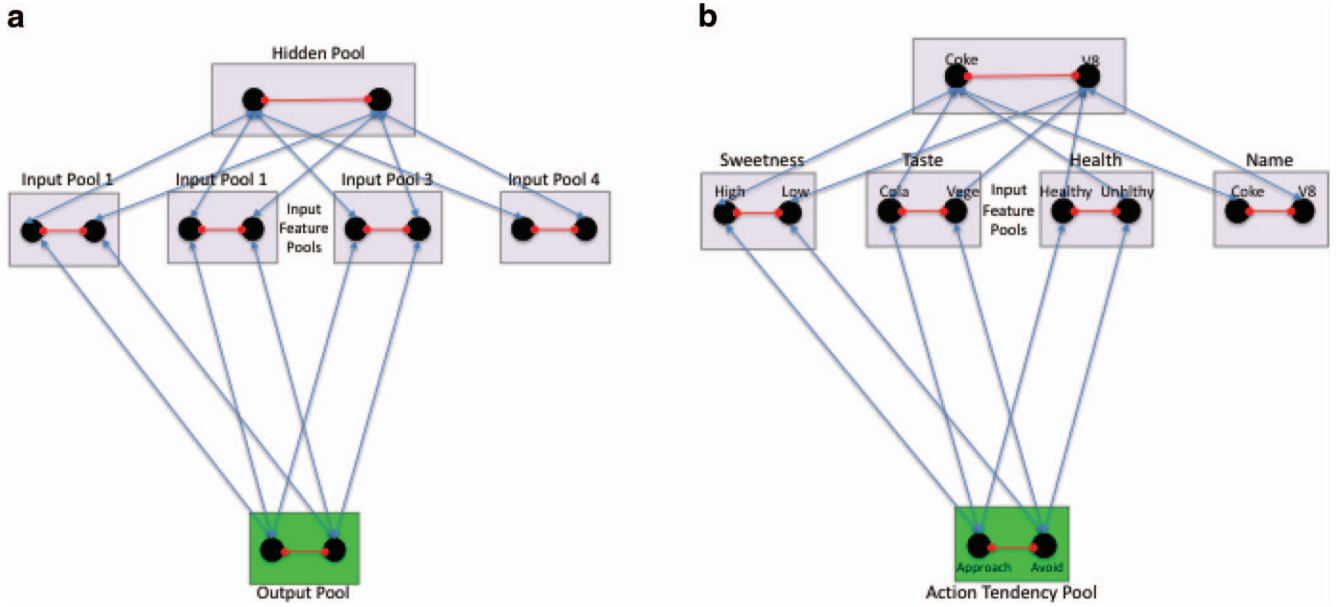
In the IAC network, units are organized into input, hidden, and output pools (see Figure 2). If units in a pool can receive input from sources outside the network, the pool is classified as an input pool; if the activation of the units in a pool is entirely determined by the activation of other units in the network, the pool is classified as a hidden pool. An output pool has units that provide activation into units outside the network. Although an input unit can receive activation from outside the network, its subsequent activation is influenced by the activation of other units that it may be connected with.

As shown in Figure 2a, each unit has inhibitory connections to other units in its pool and excitatory connections to units in different pools. The horizontally inhibiting and vertically exciting organization mimics the structure found throughout the mammalian nervous system (e.g., Kisvárdy, Tóth, Rausch, & Eysel, 1997). This implements a competition among the units such that the unit or units in the pool that receive the strongest activation tend to drive down the activation of the other units. Further and importantly, each connection is reciprocal and has weights in both directions. This makes the processing interactive in the sense that processing in each unit both influences and is influenced by processing in other units.

**Units in input pools represent features of items.** To use the interactive activation network in a particular context of value-based decision making, we use the units of the input pool of the network to represent features of items that are relevant in the decision. For example, in the context of decisions involving choosing between Coke or V8 (shown in Figure 2b), the units of the input pool could represent properties of these two beverages such as sweetness, taste, health, and name (i.e., “Coke” or “V8”). Each feature (input) pool contains units representing different levels of the feature. For example, the feature pool representing taste may have units representing the levels “high” and “low.” The feature pool representing health may have units representing “healthy” and “unhealthy.” The units within each feature pool are inhibitory (e.g., the units for taste are inhibitory with units representing other tastes but have no connections with units representing health).

**Units in hidden pools represent a summary of an item.** The item hidden pool contains units that represent an approximate summary of instances of items relevant to a choice. For example, in the case of choosing between a Coke and a V8, the item hidden pool would contain units representing “Coke” and “V8” (note that these units represent the item, whereas the “name” input units represent what the item is called). Units in the feature pools that are consistent with a hidden representation have bidirectional excitatory links to it. For example, the healthy unit in the health feature pool may have a reciprocal connection with V8 but not with Coke. Importantly, it is not required for a feature to be a physical property of some underlying stimulus; rather an event may be coded as a feature associated with a stimulus if it happens to co-occur with that stimulus.

Having hidden units allows for the network to infer features, even when it does not receive perceptual evidence supporting those features. For example, the network may infer that a cola



**Figure 2.** An interactive activation network. a: Simple processing elements, called units, represented by black circles are organized into different pools. Units within a single pool have inhibitory connections between them (represented by red lines with circular terminals). Units across pools have reciprocal excitatory connections (represented by blue lines with arrow-shaped terminals). The number of units in each pool is illustrative. Not all feature units need be connected to output units. b: An illustrative instantiation of the interactive activation network, consisting of hidden pools representing Coke and V8, feature pools representing features of Coke and V8. Features consistent with the drink are connected with corresponding hidden units (e.g., the Coke hidden unit and the Sweetness feature are connected). In this illustration, the name feature pools not connected to the action tendency pool—potentially because the decision maker has no preferences evoked by the ‘name’ feature. This need not always be the case. See the online article for the color version of this figure.

drink it is not familiar with is likely to be sweet because most prior instances of such drinks have been sweet.

**Units in output pools represent an approach or avoid action tendency.** Finally, there is a pool of action tendency units in the network. Units in this output pool may represent the action tendency of approaching or avoiding a particular stimulus. Units in feature input pools may be connected with action tendency units with specific high or low weights. For example, activation in the “sweet taste” feature unit may result (depending on decision maker preferences) in more activation in the approach action tendency unit, and activation in the “high” price unit may result in activation in the “avoid” action tendency unit. However, there may be features that are not connected to either “approach” or “avoid” units. These units may exert their influence indirectly—via conjunctive connections to feature units that are directly connected with the output approach/avoid units. The approach and avoid units are mutually inhibitory.

**Activation.** The activation coming into each unit is summed algebraically, using a weighted average, constrained by maximum and minimum values, to yield a net input. The net input is combined with the existing activation to produce a new activation value. Activation of each unit decays at a rate specified by a model parameter (full quantitative details are provided in the *Algorithmic Details of the IAC Network* section). The output function of each unit is zero if the activation is below a specified threshold and is equal to the difference between the activation value and threshold

if the activation is above threshold. The activations of the units in an IAC network evolve gradually over time. The model updates activation in discrete steps called cycles, but by assuming each cycle to be a small unit of time, the activations in the model approximate a continuous updating process.

### Algorithmic Details of the IAC Network

As in most neural networks, the net input to a particular unit, say  $u_i$  in the IAC network, is the sum of the influences of all of the other units in the network plus any external input  $e_i$  from outside the network. The influence of some other unit (say,  $u_j$ ) is the product of that unit’s activation,  $a_j$ , times the strength or weight of the connection to unit  $i$  from unit  $j$ , designated by  $w_{ji}$ . Thus the net input,  $net_i$  to unit  $u_i$ , is given by

$$net_i = estr \times e_i + \sum_j w_{ji} \times a_j$$

where  $estr$  is a parameter that scales the contribution from external input relative to internally generated inputs to the unit.

Parameters  $\alpha$  (alpha) and  $\gamma$  (gamma) similarly scale the influence of excitatory and inhibitory inputs. If  $\alpha$  is equal to  $\gamma$  (as it is in many IAC implementations), then the above equation becomes:



$$\text{net}_i = \text{estr} \times e_i + \alpha \times \sum_j w_{ji} \times a_j$$

Our parameter values in this work follow the values chosen by McClelland and Rumelhart (1989) for prior IAC implementations. In particular, we've set 'estr' to 0.4, and set both alpha and gamma to 0.1. We also use four other parameters in our network implementations: max, min, rest, and decay; *max* and *min* are the maximum and minimum activation, respectively, that any unit can take. Their values are set to 1.0 and -0.2, respectively. At rest, a unit has an activation of -0.1 (thus "rest" = -0.1), and the decay rate parameter "decay" is set to 0.1. All parameter values are kept unchanged across all simulations in this work.

If the net input,  $\text{net}_i > 0$ , then the change in activation  $\Delta a_i$  is given by the following equation:

$$\Delta a_i = (\text{max} - a_i)\text{net}_i - \text{decay}(a_i - \text{rest})$$

Substituting the parameter values, we get:

$$\Delta a_i = (1 - a_i)\text{net}_i - 0.1(a_i - (-0.1))$$

Thus, if activation is near or equal to the maximum value, then  $\Delta a_i$  may be a negative value, even though  $\text{net}_i$  is positive. However, when  $a_i$  is not close to the max, then increases in  $\text{net}_i$  lead to approximately linear increases in  $a_i$ . In the absence of net input, the decay parameter drives down the activation of a unit toward rest.

A similar equation applies when  $\text{net}_i < 0$ :

$$\Delta a_i = (a_i - \text{min})\text{net}_i - \text{decay}(a_i - \text{rest})$$

A unit is said to converge when  $\Delta a_i = 0$ . When  $a_i > 0$ , this implies that

$$0 = (1 - a_i)\text{net}_i - 0.1(a_i - (-0.1))$$

Ignoring the 0.01 term, this simplifies to

$$a_i = \frac{\text{net}_i}{\text{net}_i + 0.1}$$

Analogous results are obtained when  $a_i < 0$ . Additional details related to calculations pertaining to competition between units are provided in the online supplementary materials.

Unit convergence does not imply network convergence. However, it is known (Perfetti, 1993) that any neural network with reciprocal weights does converge. When it converges, we may infer that the activation values of each unit are less than 1 (the maximum value allowed) and greater than min.

The activation update equations provide the backbone of the algorithm underlying the interactive activation network (Matlab code provided in online supplementary materials). The algorithm consists of two main routines: getnet, which computes the net input for each pool in the network based on the activations at the end of the previous cycle, and update, which updates the activation values for each unit based on the net inputs computed by getnet. 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. The update routine uses the net input value from the getnet routine and uses the above equation to update the activation levels for each unit in the network. A single getnet and update routine constitutes a cycle. The interactive activation net-

work may execute many cycles until convergence (if the time to make a decision is not of interest), or it may be terminated when an action tendency unit (output layer) crosses a threshold. The number of cycles is then considered to be proportional to the amount of time it takes to make the decision.

## Assumptions and Rationale

We next describe a set of assumptions related to structuring and using the network to simulate value-based decisions. Our goal is to use the assumptions presented here—and only these assumptions—to simulate a range of phenomena described in *A Consideration of Two Key Assumptions of Value-Based Decision Making* section. We begin by summarizing the assumptions and their sources in Table 1 below.

Many of the above assumptions were made in the original interactive activation model (McClelland & Rumelhart, 1981), others have frequently been made in papers using the IAC model (and other related papers), and a few assumptions related to value-based choice are new to this work. Certain assumptions in Table 1 (those entries italicized in Column 1) require a more detailed description of rationales and implications. We present these below.

**(A1) Hidden units.** As described in the *IAC Network for Value-based Decision Making* section, we have assumed that our networks have hidden pools. The units of these pools often represent an approximate summary of past experiences with instances of some type of item in the world. For example, a "Coke" hidden unit represents a summary of past experiences with instances with this drink; it acts as a hub through which various properties of Coke are connected to each other.

In line with exemplar models in use in many domains of psychology (e.g., Hintzman & Ludlam, 1980; Medin & Schaffer, 1978), the IAC model is based on the idea that there is a separate hidden unit for each experience someone has had with an item, even though we use a single hidden unit in the model to approximate the collective activity of the population of experience-specific hidden units that all share a presented property (e.g., the name Coke). On this way of thinking, a new hidden unit arises from associative conjunctions between co-occurring properties, each time an item is experienced. For example, we may frequently experience the name Coke along with sweetness, bubblyness, red packaging, and several other, perhaps less reliably co-occurring features. Each such encounter binds all of these elements together via its own experience-specific conjunction unit. The hidden unit for Coke that we use in our simulations represents an approximate summary of all of these conjunction units pertaining to Coke. Features that consistently co-occur with Coke (e.g., the color red) are strongly connected to the summary hidden unit (because they have many experience-specific conjunction units); features that may inconsistently co-occur with Coke (e.g., where it is purchased) will be weakly connected to the hidden unit (because they have fewer experience-specific conjunction units).

Connecting features to each other via conjunctive hidden units affords the opportunity to exploit malleability and context-sensitive associativity. Related to malleability, consider the case of a hypothetical decision-maker who has encountered many instances of drinks that are bubbly and unhealthy (e.g., Coke, Pepsi, and Sprite). On observing that an unknown drink is bubbly, the decision-maker is likely to conclude that it is unhealthy. In the IAC

Table 1

*Assumptions of the Interactive Activation and Competition (IAC) Model for Value-Based Decision Making*

No.	Assumption	Role in model	Source
A1	<i>Hidden units</i> represent an approximate summary of past experiences with instances of some type of item in the world.	Hidden units allow for context sensitivity and malleability in value-based decision making.	McClelland & Rumelhart (1981)
A2	<i>Feature units</i> represent summary representations of an item's cognizable property	Feature units are the only units that can receive input activation from outside the network. They influence and are influenced by hidden units and action-tendency units.	McClelland & Rumelhart (1981)
A3	<i>Action tendency</i> units represent common elements of approaching or avoiding items	Action tendency units are the output units of IAC networks. Their activation represents approach/avoid tendencies and are used to identify choice. Importantly, they are not valuation units as they can receive activation that is value-related or otherwise.	Read et al., 2010
A4	Units within an input pool, hidden pool or action tendency pool have <i>bidirectional inhibitory</i> connections with each other.	Inhibitory connections ensure that processing in a network is mutually competitive. For example, if item feature units represent the property of 'red' and 'blue', competition ensures that both units are not concurrently highly activated.	McClelland & Rumelhart (1981)
A5	<i>Weights between hidden units and feature units</i> are positive if a property represented by a feature unit is consistent with an item represented by a hidden unit.	Feature unit—hidden unit weights represent the extent of association between a property represented by a feature unit and an item represented by a hidden unit. They are always between 0 and 1, and always bidirectional.	McClelland & Rumelhart (1981)
A6	<i>Feature units may influence action tendency units</i> as follows: (a) innate connections, (b) indirect influencers via hidden units, (c) learned direct connections over a broad range of experiences.	(a) Innate feature/output connections correspond to evolutionary-conserved tendencies related to certain features (e.g. food, physical pain). (b) Features that are not directly connected to action tendencies (e.g. a brand name) can exert indirect influence on action tendency units via their connection with hidden units, which in turn are connected to other feature units with direct connections. (c) Some features (e.g. money) may often co-occur with approach or avoid tendencies. Such associations can lead to direct feature-action tendency connections that are not innate.	(a) New to current model; (b) McClelland & Rumelhart (1981); (c) New to current model
A7	<i>Input</i> into the network may only enter the network via the feature pools.	Increased network input may correspond to three situations: (a) any empirical manipulation that increases exogenous attention (e.g. a beep), (b) direct experimental instruction, and (c) increased presentation time.	Adapted from IAC input usage
A8	Activation at convergence in approach and avoid units determines <i>choice</i> . <i>Convergence</i> is defined as activation values stabilizing within an interval of $\pm 10^{-4}$ for ten consecutive network cycles.	In go/no-go decisions an action is assumed to occur if, and only if, activation in the approach unit exceeds a threshold parameter. In the multi-alternative choice context, two or more instances of the model (one with the inputs appropriate for each option) are run in parallel, and the approach output unit with the highest activation at convergence is selected. The approach units of each network have an inhibitory weight ( $-1$ ) with each other.	New to current model
A9	The cumulative density function (cdf) is used to capture <i>probabilistic processing</i> in the network.	Each activation value represents a point on a normal distribution, and the corresponding cdf value is used to determine the associated probability. The cdf is mathematically equivalent to the Softmax function, which in turn is equivalent to the Luce Choice rule (McClelland, 2013)	New to current model
A10	Inter-trial effects <i>learning</i> effects within a single experiment are captured by increased connection weights.	Repeated experiences with the same associative relationship result in a strengthening of connection weights between a hidden unit encoding the associative relationship and the elements that are involved in the association. Here, we assume the increase in connection weights without specifying a particular learning algorithm.	Kumaran & McClelland, 2012

framework, this occurs because the “bubblyness” feature activates the Coke, Pepsi, and Sprite hidden units. These units are, in our hypothetical network, all associated with the unhealthy feature unit, which is likely to receive strong activation. However, this association is malleable: imagine that the decision-maker next encounters several other drinks that have the bubblyness feature, but are also healthy (e.g., mineral water). In these drinks, bubblyness is associated with the healthy feature (via new conjunction units). Now the bubblyness features would activate conjunction units that vote for healthy (as well as earlier conjunctions voting for unhealthy). This would weaken the association between bubblyness and unhealthiness. With enough new conjunctions, bubblyness may even become associated with healthiness. Hidden units thus play the role of casting top-down votes that determine the effective associative strength between different features.

Related to context-sensitive associativity, associations involving hidden units are also subject to selective activation based on combinations of cues (Medin & Schaffer, 1978; McClelland, 1981). Increasing specificity of input can alter the pattern of activated conjunctions in the hidden pool. Continuing with our bubbly drinks example, providing input into the “sweet” and “bubbly” feature units will activate hidden units corresponding to Coke, Pepsi, and Sprite; this activation may in-turn activate the unhealthy feature unit. On the other hand, input into the nonsweet and bubbly feature units will activate the hidden unit corresponding to mineral water; this activation may in-turn, activate the healthy feature unit.

Exactly what summary representations should we posit when we formulate simulations of specific situations? Our approach to this is based on the idea that behavior in accordance with summary representations at different granularities is an emergent property of our brain-based associative learning systems, and that it is therefore justifiable to choose a granularity appropriate for the specific situation. For example, if someone sometimes drinks Coke from a soda machine and other times from a can, and the ones from the soda machine are often flat, then activating the name Coke and “from a soda machine” will have the effect of weighting the resulting pattern of activation in accordance with the features of experiences with Cokes from the soda machine, so that the feature “flat” would be active instead of bubbly. In that case, we might simulate a choice between a coke from a soda machine and a coke from a can using separate hidden units for the two types of cokes, each approximating the collective experiences with the appropriate subset of Cokes. Other times, when simply simulating the choice between Coke and V8, we can approximate the influence of the entire ensemble of Coke experiences with a single summary hidden unit.

**(A2) Feature units.** As described above, feature units represent properties of the world. Examples of feature properties in the context of drinks may include the name of a drink, its sweetness, bubblyness, packaging, and the extent of its healthiness. Feature units can receive input from outside the network, as well as excitatory activation from hidden units and inhibitory activation from other units within the same pool.

Similar to hidden units, feature units are best conceptualized as summary representations of a property. A property may be generally specified involving a range of subcomponents (e.g., healthiness) or it may be tightly specified (e.g., the particular shade of red used by Coke). For our purposes, features are not necessarily

simple values on a dimension but can be any cognizable property (Collins & Loftus, 1975).

Whether two intersecting features (e.g., different shades of red) are separated or joined in the feature units will depend on whether there are items to be distinguished within the setting of the experiment that differ in that particular feature (e.g., an experiment involving two slightly different shades of red). In both cases, the underlying network is assumed to be the same—the approximation to it that comes to the forefront in an experiment is the only thing that changes.

**(A6) Weights between feature units and action tendency units.** These weights represent the action tendencies associated with certain features. Activation in some features may result in increased activation in the ‘approach’ action tendency unit, whereas activation in other features may result in increased activation in the ‘avoid’ action tendency unit. Not all features need be connected to either output unit.

We assume that features may influence output in three ways: First, some features may be innately connected to approach and avoid action tendencies. We assume that there are only a small set of such feature units with direct, evolutionarily conserved, connections to action tendency units. These may include units representing features related to food (connected to the approach action tendency unit), physical pain (connected to the avoid action tendency unit), and particular bodily states—such as feeling energetic or sluggish.

A second set of features may influence approach units indirectly—mediated by conjunctions. For example, let us say that a tone reliably precedes the delivery of food for an animal. Here we assume that food-related features are innately connected to that approach unit, but that the tone unit is not innately connected to the approach unit. However, in the IAC perspective, the tone and the food are connected to a common hidden unit, and activation in the tone unit would elicit increased approach behavior via the food-approach connection. Thus, the tone feature would elicit approach behavior even though it is not directly connected to it. A brand name (e.g., Coke) may also exert its influence in this manner.

A third possible way for feature units to influence output units occurs when features (e.g., money, or pushing a button to obtain a desired item) co-occur with approach behavior over a broad range of experiences. We assume that this repeated co-occurrence can directly bind certain features with output units.

All feature-output weights are excitatory and bidirectional. Greater weights between one feature and an action tendency unit (compared to weights between another feature unit and that action tendency unit) indicate a relatively greater influence of that feature with respect to that action tendency.

**(A7) Network input.** Input into the network may only enter the network via the feature pools. When modeling empirical data, increased input activation may correspond to one or more of the following three scenarios:

- Any empirical manipulation that increases exogenous attention (Theeuwes, 1991): for example, a loud tone, or a flashing red border around a stimulus would result in more input for that stimulus compared to a situation in which the stimulus is not accompanied by a tone or a flashing red border (Glöckner & Herbold, 2011). Similarly, requiring

participants to interact with an experimental feature will increase activation related to that feature. For example, asking participants to make a judgment about a feature (e.g., its color or size) will increase input activation related to that feature proportional to the level of interaction observed by the experimenter.

- Direct experimental instruction: for example, instructions from the experimenter to focus on a particular feature (e.g., the healthiness of a food item) would result in more network input compared to a situation in which the stimulus is not accompanied by experimenter instructions.
- Increased presentation time: for example, presentation of good features for 900ms would result in a (3×) longer duration of network input than presentation for 300 ms.

**(A9) Transforming activation values into probabilities for comparison with empirical data.** Empirical variables over multiple trials and participants are inherently probabilistic. Although we believe mental processes are inherently stochastic (Usher & McClelland, 2001), we have simulated the activation process as deterministic, for simplicity. To capture probabilistic processing, we used the cumulative density function (cdf) of a normal distribution to translate activation values into response probabilities. For example, in the go–no-go context we used the activation value of the approach unit as the mean of a normal distribution (with a fixed standard deviation) and calculated the area under the curve (i.e., the cdf) greater than a threshold parameter to estimate the proportion of the time an action would occur. In the context of experiments involving multiple variables we used a single normal distribution (with a mean and standard deviation kept fixed throughout the simulation), and used the cdf for output activation value corresponding to each variable to estimate its probability. This method has the benefit of simplicity and is mathematically equivalent to other commonly used methods such as the SoftMax function, which in turn is equivalent to the Luce choice rule (McClelland, 2013).

### Simulation #1: Coke Versus V8. Illustrating Usage of the IAC Network to Simulate Choice

We will next demonstrate several properties of the interactive activation network using an illustrative model involving preferences between beverages. The assumptions used in setting up and using this network are indicated in parentheses (e.g., A1 indicates a reference to Assumption 1 from the *Assumptions and Rationale* section).

**Network structure.** The network corresponds to the structure displayed in Figure 3a: A hidden pool contains two units corresponding to Coke and V8 (A1). There are four feature pools for sweetness, taste, health, and name. The sweetness pool consists of two units corresponding to high and low sweetness; the taste pool consists of two units corresponding to cola taste and vegetable taste; there are two health units corresponding to high health and low health; there are two name units corresponding to Coke and V8. The name units are activated only by the perception of the name “Coke” or the name “V8.” This is akin to a unit that is activated by hearing the name “John” but whose activation does not represent the sum of cognitions about the person John (A2). The action tendency pool contains two units corresponding to

approach and avoid action tendencies toward the option under consideration (i.e., Coke or V8; A3).

As in all networks in the present work, the feature pools can receive input from outside the network (i.e., they are input pools; A7), and the action tendency pool can provide output outside the network (i.e., it is the output pool; A3).

The units within each pool have bidirectional inhibitory weights of strength  $-1$  (A4). The units in each feature pool bidirectionally connect (with excitatory weights of strength 1) to consistent units in their corresponding hidden pool (A5). For example, the hidden unit for Coke has bidirectional excitatory weights (+1) to the “high” sweetness unit, the “cola” taste unit, the “low” health unit, and the “Coke” name unit. The hidden unit for V8 has bidirectional excitatory weights (+1) to the “low” sweetness unit, the “vege” taste unit, the “high” health unit, and the “V8” name unit.

Each unit in the feature pool bidirectionally connects with units in the action readiness pool with varying weights shown in Table 1 (A6). These weights are assumed to reflect the preferences of a hypothetical decision maker. Such a decision maker prefers more sweet things over less sweet things, the cola taste over the vege taste, and healthy items over unhealthy items. Here, we are assuming that the connections between the sweetness pool, taste pool, and the action tendency pool are direct and innate (A6i). The direct connection between the health pool and the action tendency pool is associative formed by a preference to eat healthily over a broad range of experiences (A6iii)

Many features either have an approach connection, or an avoid connection, but this need not always be the case. For example, the vege taste unit has both approach and avoid connections. The name “Coke” and the name “V8” are assumed not to be linked to the output units (i.e., have zero weights). However, the name units can exert their influence indirectly via the hidden units (A6ii)

We describe three aspects of the IAC network in the context of this illustrative model. First, we will use the model to illustrate decision-making—both in the multialternative choice context and the go/no-go context. Our simulation will show that choice between two options can be modeled via levels of activation in the action-tendency units, and choice involving go–no-go decisions can be modeled by using a threshold level in the action tendency unit. Second, we will illustrate the ability of the model to infer properties from incomplete information. We will input certain features into the model (e.g., high health), and examine how the model generalizes these features to other likely features (e.g., whether the beverage has a cola taste or a vegetable taste). Third, we shall examine the consequences of competition in competing networks, and in competing units in feature pools.

**Simulating decisions using interactive activation.** The decision maker in the illustrative model depicted in Figure 3a prefers features associated with Coke to those associated with V8. This preference is captured by more features associated with Coke being connected with the approach unit with greater weight (see Table 2). The features associated with V8 (other than calories “low”) are associated with the approach unit with less weight and are associated with the avoid unit with greater weight.

In this simulation, we first consider two decision making contexts: Two-alternative choice and go/no-go choice (A8). To model the choice between Coke and V8 we ran two parallel instances of the model, activating the Coke hidden unit via its name unit, and the V8 hidden unit via its name unit. Activation



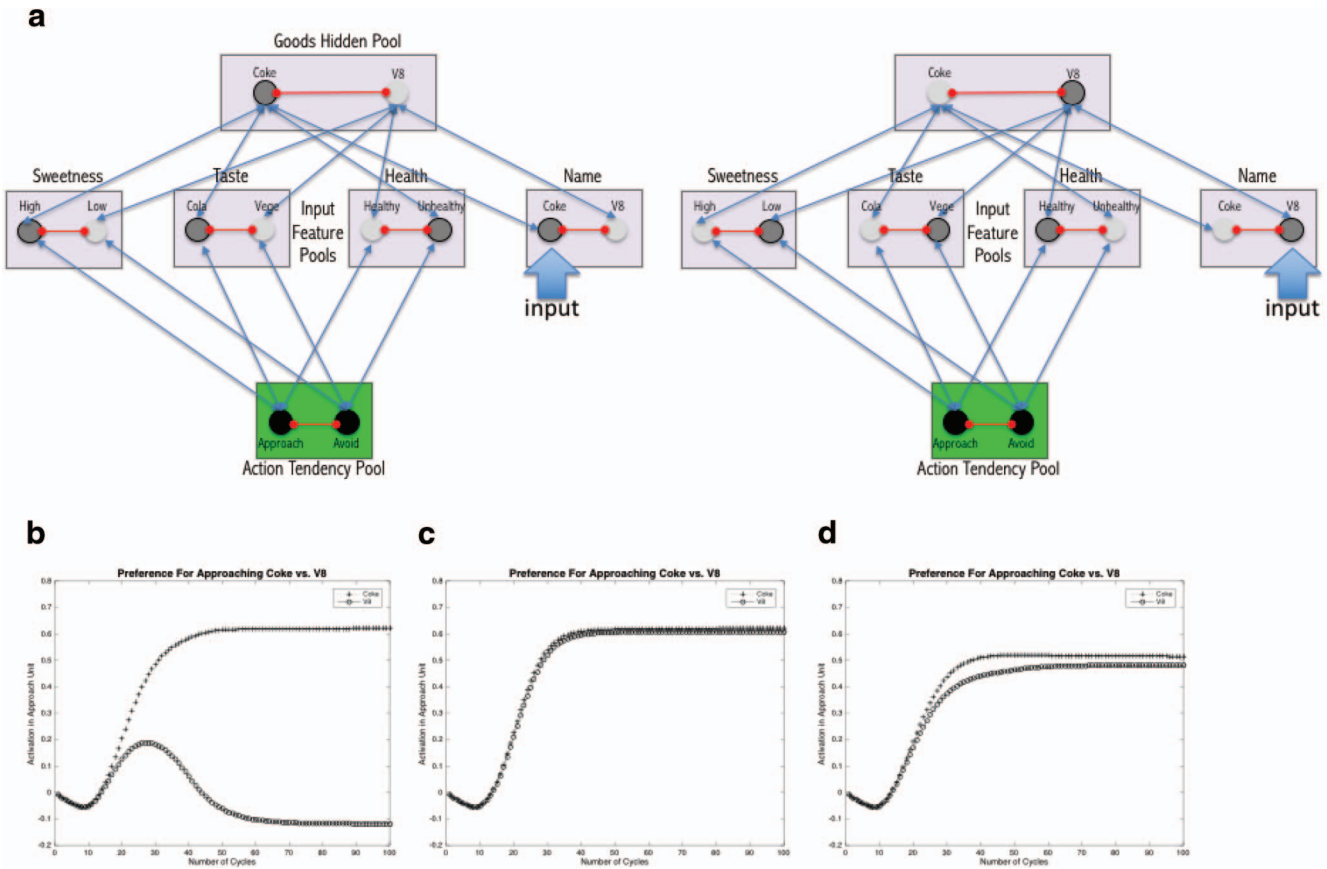


Figure 3. Network structures and activations in Simulation 1: Coke versus V8. Panel a shows two identical networks. The left network gets activation for the “Coke” name unit, and the right network gets activation for the “V8” name unit. In each network, feature and hidden units with greater activation are shaded darker. Panels b–d show activation levels related to the approach unit for Coke and V8 in Simulation 1. Panel b shows the resulting difference in approach activation when there are clear differences in weights between Coke (+) related features and V8 (o) related features (in favor of Coke; the weights are shown in Table 2). Panel c shows the effect of minimizing the weight differences and removing competition between the approach units in the two networks. This makes the activation in the approach Coke and V8 units become nearly identical. Panel d shows the effect of introducing inhibition between the two approach units while keeping all other weights unchanged from Panel c. Panel a does not represent mutual inhibition between the approach units of the two units since it is operational in Panels b and d, but not in c. See the online article for the color version of this figure.

flowed from the name units to their respective hidden units, and then in-turn into the feature units representing each drink. Activation in the feature units then propagated to the action tendency units. The action tendency units for each network inhibited each other (i.e., the approach unit in the Coke network competed with the approach unit in the V8 network and the avoid unit in the Coke network competed with the avoid unit in the V8 network).

At convergence for the Coke network, the activation in the approach action tendency unit is 0.66 and the activation in

the avoid action tendency unit is  $-0.14$ . For the V8 network, the activation in the approach unit is  $-0.12$  and the activation in the avoid unit is 0.55, showing that all else equal (e.g., actions required to procure the drink), the network would approach Coke and avoid V8 (Figure 3b shows activations in the approach units).

In our simulation, we elected to input the name unit, but this need not always be the case. For example, the network input corresponding to a situation in which a decision-maker attends to something sweet would be provided into the high unit in the

Table 2  
Feature-Action Tendency Weight Matrix for a Hypothetical Decision Maker

Action	SwHigh	SwLow	TasteCola	TasteVege	HealthHigh	HealthLow	NameCoke	NameV8
Approach	1	0	1	.75	1	0	0	0
Avoid	0	1	0	.5	0	1	0	0

sweetness pool. This would activate the respective hidden units (for sweet drinks encountered in the past) which in turn would activate the feature units for the drinks in the consideration set.

Go-no-go decisions (i.e., where the decision is about either obtaining the item, or doing nothing) may be simulated by assuming a threshold in the action tendency units (A8). For example, assuming a threshold of 0.65 in the approach unit would cause a no-go decision for V8 and a go decision for Coke. Here there is only one network processing activation, not two, as was the case for two-item choice.

**Inferring features from incomplete information.** A second performance aspect of the interactive activation network is that it can infer feature properties of newly encountered items based on prior knowledge. Consider the scenario in which the network encounters a mystery drink that has vege taste, but an unknown health level. To examine network behavior in this context, we provide input to the vege taste feature unit. This unit causes activation in the V8 hidden unit (this activation represents the prior knowledge of the network). The activation in the V8 hidden unit in-turn activates the high health unit over the low health unit (0.64 vs. -0.14 at network convergence). Thus, the network assumes that the mystery drink is healthy. This may not actually turn out to be the case, but it represents the network's estimate of likely properties. If the network later encounters healthy beverages with a cola taste, it would update its estimates.

**Competition.** A third performance aspect of the interactive activation network concerns competition between units. Imagine that V8 has the weights shown in Table 2, except the decision maker has a weight of 0.90 from the vege taste unit to the approach unit and has zero weight to the avoid unit (instead of the 0.75 and 0.5 weights in Table 2). Under these conditions, the weights from the Coke features into the output pools are nearly identical to the weights from the V8 features. (The sweetness pool in favor of Coke is canceled out by the symmetrical health pool in favor of V8, and the taste weights are now nearly identical). Further, imagine that (unlike in the original simulation above) there is no mutual inhibition between the approach units for Coke and V8. This results in a nearly identical approach unit activation for Coke and V8 (Figure 3c).

Reintroducing inhibition between the approach units results in clear separation (Figure 3d) even though the weights in the Coke network are only very slightly more favorable than the weights in the V8 network. This separation occurs because the two approach units are mutually inhibiting each other and even a small advantage for one of the units (Coke in this case) is amplified over time.

Competition similarly enables mutually inhibiting units in a feature pool to have clearly differentiated activation levels. Imagine for example that the input activation for a drink is nearly identical for the high health and low health units, except that the high health unit receives slightly greater activation. This initial difference will be amplified over time (i.e., network cycles) until the network represents features related to high health.

**Significance.** In Simulation 1, preference emerges via the interaction of features (e.g., taste, health) that have a direct bearing on utility. In this case activation in the output units precisely represents the value of the drink (Coke or V8). More generally, if a network only includes features related to the utility of the item represented in the conjunction unit, then the activation level in the output units represents the overall value of that item. However, if

the network includes features and connections reflecting the influence of items unrelated to utility, then the output units do not only represent value-related activation. Thus, the IAC framework can account for all patterns of choice observed in models that are based on the computation of value and it can also account for patterns of choice created by activations unrelated to value. Later simulations will show that IAC networks can also simulate contexts in which choice is not consistent with expected value maximization.

In addition, and notably, Simulation 1 did not include serial or modular processing. All units participated in influencing each other from the very start, and there were no separable submodules related to perception, valuation, and action selection.

## Situating the Present Work in Prior Literature

Beyond its initial application to context effects in perception, the IAC framework has been used to capture a wide range of psychological phenomena, including perception (McClelland, Mirman, Bolger, & Khaitan, 2014) emergent category formation and category-based inference (McClelland, 1981), social cognition (Freeman & Ambady, 2011), memory (Kumaran & McClelland, 2012), social behavior (Ehret, Monroe, & Read, 2015; Read & Miller, 1998), legal judgments (Simon, Stenstrom, & Read, 2015), emotional consciousness (Thagard, 2008), and probabilistic inference (Glöckner, Betsch, & Schindler, 2010). Unsurprisingly, we are sympathetic to these efforts and consider their contributions to be substantial and useful.

However, there are also some important differences between our work and prior approaches. First, we have extended the IAC framework to examine two core questions in value-based decision making: whether such decisions are always in the service of maximizing expected value, and whether they follow a serial/modular process. Although these points have not been generally emphasized, we note that the absence of seriality/modularity has also been observed by others (e.g., Simon, Snow, & Read, 2004). Second, although we share some assumptions with the papers above, we do have a unique set of assumptions and those assumptions do have a unique set of implications with respect to value-based decision making. For example, we do not have units that are solely tasked with evaluating stimuli. This assumption has generally not been made in the value-based decision-making domain (but a similar approach has been taken in the motivation and personality domains, e.g., Read et al., 2010). As we discuss in the *General Discussion*, the parsimony of not assuming units exclusively dedicated to computing value invites consideration of the possibility that value-based choices may involve variables other than value and they may not unfold in a serial and modular process.

## Simulations of Empirical Findings Using IAC

In reference to the two central what and how questions about decision making, the IAC framework makes two proposals: First, although decision making may sometimes be consistent with expected value maximization (e.g., in Simulation 1), at other times it may be nonoptimal. This may occur because of attention to irrelevant contextual variables or because of learned associative constraints. Second, the process of value-based decision-making features the reciprocal (nonserial) and interactive (nonmodular)

influence of many variables acting in parallel. Next, we will feature simulations of empirical phenomena (previously described in *A Consideration of Two Key Assumptions of Value-Based Decision Making* section) that lend support to these proposals.

## Overview of the Simulations

We selected experimental targets for simulation based on the following three criteria: (a) The empirical phenomenon should pertain to (1) a violation of optimality due to attending to contextual factors and/or (2) a violation of optimality due to constraints caused by associative learning and/or (3) interactive effects of the type that are unlikely to feature in processes that are serial/modular. Together the chosen simulations should span these categories. (b) The mechanistic rationale offered by the IAC simulation for the empirical phenomenon should advance the current set of explanations in the field, and should richly leverage the assumption set underlying the IAC model. Phenomena that can be simulated by a simple assumption related to free parameters in the model and/or ones that do not amplify existing analysis in the field would not be ideal simulation targets. (c) The underlying empirical phenomenon should be well-known in the value-based decision-making literature.

First, related to nonoptimality, we showcase empirical phenomena that are not consistent with expected value maximization. We begin with examples in which nonoptimality effects may be attributed to the activation of units representing contextual variables that are irrelevant to utility. These include the Halo Effect (Simulation 2), default effects (Simulation 3), and anchoring (Simulation 4). Simulations related to instructional cue-effects (Simulation 9 and Simulation 10) also rely on activation of incidental stimuli

and are detailed in the online supplementary materials. Second, we simulate examples in which nonoptimality effects may be attributed to learned associative constraints. These include negative auto-shaping (Simulation 5) and reward and path associations (Simulation 6). Third, related to the absence of serial/modular processing, we showcase a phenomenon in which processing appears to be interactive and reciprocal (Simulation 7). Fourth, and finally, we propose that in addition to providing new viewpoints on the what and how questions described above, the IAC framework can provide a plausible mechanism for dynamic process in goal-directed behavior (Simulation 8)—a key component of value-based decision making. The list of all simulations is summarized in Table 3.

## Simulation #2: Halo Effect

This simulation was designed to showcase how the interactivity of the IAC model is not consistent with choices that are optimal or with decision making processes that are serial/modular. The empirical structure in this classic demonstration of the Halo Effect invited evaluation and not choice. However, since evaluations in this experiment were shaped by preferences, and preferences are a part of value-based decision making, we consider this demonstration of the Halo Effect to be an important part of the present work even though participants were not invited to make a choice.

**Target experiment.** The *Halo Effect* (Nisbett & Wilson, 1977) is generally defined as the tendency for an impression created in one area to influence opinion in another area. To specify this effect, researchers videotaped two different interviews that were staged with the same individual—a college instructor who spoke English with a European accent. In one of the interviews, the

Table 3  
*List of Simulations*

Number	Section	Topic	Phenomenon	Purpose
1	3.4	Illustrative	Choice of Coke vs. V8	Functionality consistent with optimality (also: illustrate basic IAC mechanisms)
2	4.2	Halo effect	The attractiveness of one feature influences perceptions of other unrelated features	Nonoptimality due to context variable activation & absence of serial/modular processing
3	4.3	Default effects	Valued actions remain unchosen until attention is directed towards them	Nonoptimality due to context variable activation
4	4.4	Anchoring (a & b)	Estimates of number of (a) physicians in a city biased by incidental ID# and (b) prices via social security number	Nonoptimality due to context variable activation and attentional effects
5	4.5	Negative auto-shaping	Associative constraints between an illuminated key and food delivery can cause non-optimal behavior	Nonoptimality due to associative constraints
6	4.6	Reward and path associations	Prior associations cause animal to undertake extra costs for rewards	Nonoptimality due to associative constraints (also, effect of action costs)
7	4.7	Motivated influences on visual stimulus categorization	The classification of a visual stimulus is dependent on the outcomes associated with each classification	Interactive and reciprocal (not serial/modular) processing
8	4.8	Dynamic processes in goal-directed choice	Prior goals can bias activation towards the option they are associated with	Extension: Influence of goals
9/10	SM	Instructional cue effects	Cuing taste/health features of food items and input duration amplifies their influence on choice	Nonoptimality due to context variable activation and input duration

Note. IAC = interactive activation and competition.

instructor was warm and friendly, in the other, cold and distant. Participants were assigned to separate conditions—and viewed one of the two videos in each condition. They were then asked to classify the instructor’s appearance, mannerism, and accent on a scale that ranged from *very appealing* to *very irritating*.

**Empirical results.** The participants who saw the warm instructor rated his appearance, mannerisms, and accent as appealing, whereas those who saw the cold instructor rated these attributes as irritating. Seventy percent of the participants who had seen the “warm” interview rated the instructor’s appearance as appealing (the rest rated it irritating); 62% rated his mannerisms as appealing (the rest rated it irritating) and 48% rated his accent as appealing (the rest rated it irritating). Of the participants viewing the “cold” interview, the ratings were generally flipped: 68% rated the instructor’s appearance as irritating, 60% rated his mannerisms as irritating, and 81% rated his accent as irritating.

**Network structure.** We used a network consisting of a hidden pool and a set of input feature pools (see Figure 4). The network structure corresponds to the idea that over time we develop favorable impressions of people who are warm, attractive, and with appealing mannerisms and accents and unfavorable impressions of people who are not. The former set is represented by the favorable hidden unit, and the latter by the unfavorable hidden unit. These units are connected to features of people we generally find favorable or unfavorable. There are four such feature pools each with two units, as follows: a warmth feature pool (two units varying between warm and cold), an attractiveness feature pool (two units varying between attractive and unattractive), a mannerisms feature pool (two units varying between appealing and irritating), and an accent feature pool (two units varying between appealing and irritating).

Weights between units in the same pool are competitive (equal to  $-1$ ). Weights from the features to the hidden pool units are excitatory—the positive units (e.g., attractive, appealing) are connected to the favorable unit and the negative units (e.g., unattractive, irritating) are connected to the unfavorable unit. Each favorable weight was assumed to be equal to its unfavorable counterpart. Since

humans are most used to making judgments about attractiveness, we assumed those weights to be the highest ( $+1$ ); we assumed mannerism weights were lower ( $+0.8$ ) and accent weights—since they are less frequently perceived—to be the lowest ( $+0.7$ ). All weights were reciprocal.


**Network dynamics.** In the “warm and friendly” condition, the external activation for the warm feature unit was set to 1, and the external activation for the cold feature unit was set to 0. In the “cold and unfriendly” condition the external activation for the cold feature unit was set to 1, and the external activation for the ‘warm’ feature unit was set to 0.

Further, we set the inputs units related to the participant’s attractiveness and mannerisms to 0. This was equivalent to assuming that the instructor was seen to be neither attractive, nor unattractive, and his mannerisms neither appealing nor unappealing. However, because most unfamiliar accents are perceived slightly negatively (Gluszek & Dovidio, 2010), we gave the “unappealing” unit in the accent pool a small external activation (equal to 0.09) in both the favorable and unfavorable conditions.

In the warm and friendly condition, activation from the warm feature unit flowed to the favorable hidden unit. The activation in the favorable hidden unit in turn activated the “attractive,” “appealing mannerism,” and “appealing accent” feature units. In the cold and unfriendly condition, the unfavorable unit received activation from the cold feature unit. This activation in the avoid unit in turn activated the “unattractive,” “unappealing mannerism,” and “unappealing accent” feature units.

Values in the feature pool units were measured after 400 network cycles by which time all network units had converged.

**Simulation results.** After convergence, we measured the activations in the feature pools related to attractiveness, mannerisms, and accents. The activation level for the winning unit (i.e., the unit with the greater activation in the pair of units representing a trait) was assumed to be a point on a normal distribution with mean of 0.2 and standard deviation of 0.5. The cumulative density function corresponding to the activation level was used to calculate the probability estimate for that activation level. The losing unit in the



**Figure 4.** Structure and network dynamics for Simulation 2: The Halo Effect. The left panel depicts the network structure and activation flows. The black arrows in the figure (partially) depict the early activation flows in the network. Initially, the observed “warm” feature unit activates the “favorable” unit in the hidden pool. These units in-turn reciprocally activate the “attractive,” “appealing mannerism,” and the “appealing accent” units (depicted by the dashed black arrows). Similar activation flows emanate from the “cold” feature units. The right panel depicts the activation in the attractiveness unit in the warm and cold conditions. See the online article for the color version of this figure.



pair was assigned the complement of the percentage assigned to the paired winning unit. As shown in Table 4, simulation results matched the pattern of the empirical data.

**Significance.** The Halo Effect simulation highlights the importance of interactivity in our network. Features that were neither perceived as positive or negative received a favorable or unfavorable cue since activation from a single feature pool activated a hidden unit (either favorable or unfavorable), which in-turn activated other (corresponding) feature units. This showcases a mechanism in which perceptions about a particular attribute (e.g., how appealing an accent is) are made in parallel with, and are influenced by, judgments about other attributes (e.g., warmth). This is not consistent with decision-specific optimality or with a serial/modular process.

Simulation #3: Default Effects

As described in the *Evidence Inconsistent With the Objective Function Maximization Assumption* section, people retain default states for a number of reasons that are consistent with utility maximization. However, there is a growing list of empirical contexts in which leaving a default state appears to have positive utility, and yet the participant does not initiate actions to leave the default. In such cases, changing the level of attention toward a previously understood cue changes behavior. We next simulate an experiment demonstrating this effect.

**Target experiment.** Participants were shown either a negative image or a neutral image (Suri & Gross, 2015; Suri, Sheppes, & Gross, 2015; Suri, Shine, & Gross, 2018) from a database of affective images, classified into positive, neutral, and negative categories. In each of 40 trials, they were shown a negative or neutral image as a default. If they did nothing, they viewed the default image for the 15-s duration of the trial. However, they were given the option to press a button to switch from the default image to a higher-valenced image. For example, participants could switch from viewing a negative image to viewing a neutral image, or from viewing a neutral image to viewing a positive image. A caption under each default image reminded participants that they could switch away from the current image to a higher-valenced image. Postexperiment debriefing confirmed that 100% of participants understood (from the caption as well as from instructions and practice trials at the start of the experiment) that they could switch away from the default image. Prior studies had also demonstrated that when given a (two-alternative) choice, participants overwhelmingly chose to view neutral images over negative ones, and positive images over neutral ones. However, in the context of a go-no-go choice, where pressing a button could have enabled

participants to leave the inferior default, they often did not do so (switching rate was 29%). In postexperiment debriefings, participants did not identify any utility maximizing consideration that could have driven their inaction.

To test whether attention was a factor in participant nonpresses, participants were randomly assigned to a high-attention group and a low attention group. Participants in both groups were shown a caption below the picture stating that a button press enabled the lower-valenced image to be replaced by the higher-valenced image. In the high-attention group (if the participant had not elected to switch images), a red-border around the caption was shown 5 s after the commencement of each trial; if the participant had still not pressed the button 10 s into the trial, the red border was briefly flashed. In the low attention group, no red border was shown.

**Empirical results.** In the low-attention group participants switched images in only 29% of the trials. In the high attention group, participants switched images in 50% of trials. Notably, the rate of switching in the first 5 s (prered border) of the high attention group was indistinguishable from the low attention group. High attention group participants switched images for 17% (6.76 out of 40) trials within the first 5 s compared with 19% (7.4 out of 40) switches in the low-attention group within the first 5 s.

**Network structure.** We used the network structure shown in Figure 5 to simulate the above results. There were two units in the hidden pool—one representing an instance of a negative-to-neutral trial (Neg2Neut), and the other representing an instance of a neutral-to-positive trial (Neut2Pos). There were three feature pools where the first pool represented the current image. It had two units; one unit represented a default negative image, and the other represented a default neutral image. The second pool represented the caption stating that the participant can press a button to switch images. The third feature pool was activated by the expected features of the alternative future image. It had one unit for a potential neutral image, and one for a potential positive image. The currently viewed negative image and the positive potential future image had weak (+0.1) associations with the (button press) action unit. The caption unit had a strong (+1) connection with the action unit. These weights referred to weak associations between viewing a low-valenced image (or having the opportunity to view a higher-valenced image) and the action of pressing a button to switch images. The caption was a direct cue, learned over time, and therefore had a stronger weight with the action unit (A6 iii, Table 1).

**Network dynamics.** We simulated the probability of a button press in any given trial by computing whether the activation in the output unit exceeded the threshold activation (a free parameter, set

Table 4  
Simulation #2: Halo Effect, Results

Features	Warm condition				Cold condition			
	Empirical data		Simulation		Empirical data		Simulation	
	Appealing	Irritating	Appealing	Irritating	Appealing	Irritating	Appealing	Irritating
Appearance	70	30	72	28	32	68	27	73
Mannerism	62	38	69	31	40	60	30	70
Accent	48	52	48	52	19	81	24	76




Figure 5. Network structure for Simulation 3: Default effects. The hidden pools correspond to the trial type: There are two competing units for each trial type. The feature pools correspond to representations involving the current image, the image-switching caption and the future image (the nature of which is known to the participant). Image features are linked with the “press button” output action with weak weights (+0.1), whereas the connection between the caption feature and the ‘press button’ action are strong (+1). Increasing the activation (with red border/flashes) in the caption feature increases button-pressing rates. See the online article for the color version of this figure.

equal to 0.38) at any time in the 15-s trial (represented by 60 network cycles). We simulated a total of 1,000 trials (each with varying levels of input activation—drawn from a random distribution—into the switch cue feature unit) to capture variance related to the fluctuations of attention. In each simulation run, we recorded whether activation in the output unit exceeded the threshold parameter. This was assumed to correspond to a button press. Simulation runs in which the activation in the output unit did not exceed the threshold parameter for 60 network cycles were assumed to correspond to nonpresses.

The external input into the units representing the (viewed) negative and the (potential) positive image feature units was assumed to be 1. This activation reflected the salience of the affective images. The external input in the switch-cue unit was assumed to vary with varying levels of attention toward it. In the low-attention condition, input was assumed to be constant throughout the trial. It was drawn from a normal distribution centered at  $-0.1$  with a standard deviation of 1. In the high-attention condition, the

input prior to the display of the red square (5 s or 20 network cycles) was assumed to be identical to the input in the low attention condition. After the introduction of the red-border however, the level of input was increased. It was drawn from a normal distribution centered at 0 with a standard deviation of 1.

**Simulation results.** In the low-attention group, the network activation in the output unit corresponded to a switch in 31% of the trials (compared to 29% in the empirical data). In the high attention group, network activation in the output unit corresponded to a switch in 48.8% of the trials (compared to 50% in the empirical data). Unsurprisingly, the rate of switching in the pre-red border period of the high attention group was identical to switching in the same time frame in the low attention group (both groups had a switch rate of 19.5% in the first 5 s of the trial, compared to an average of 18% in the empirical data).

**Significance.** Importantly, this simulation sought to explain persistence with inferior default states not by finding hidden variables that impact the valuation calculus, but by assuming differing

levels of input activation into the switch cue unit. This in-turn produced differing activation levels in the action-tendency unit enabling the action (i.e., button press) to occur or not occur. In cases of insufficient input into the caption unit the action was less likely to occur; when input was higher (in the high-attention condition), the action was more likely to occur.

#### Simulation #4: Anchoring

While using the same network, we simulate two experiments related to anchoring. The first experiment (in the value-based domain) shows that comparing price estimates to an irrelevant quantity (the last two digits of one's social security number), can influence willingness to pay. The second experiment (not in the value-based domain), and described in the online supplementary materials, shows that increasing levels of attention toward an incidental and irrelevant variable (a participant identification number) can result in increasing levels of influence of that variable.

**Target experiment.** In a widely cited study (Ariely et al., 2003), participants were shown six products (computer accessories, wine bottles, luxury chocolates, and books), which were briefly described without mentioning their market price. After introducing the products, participants were asked whether they would buy each good for a dollar figure equal to the last two digits of their social security number. After this Accept/Reject response, they stated their dollar maximum willingness-to-pay for the product. Experimenters ensured that participant responses were incentive compatible (i.e., they had to actually complete the transaction from a randomly chosen trial).

**Empirical results.** The willingness-to-pay numbers were examined by quintiles of the social security number distribution. The values of the top quintile participants were typically greater by a factor of three. For example, subjects with social security numbers in the top quintile were willing to pay \$37.55 on average for a bottle of wine, compared with \$11.73 on average for subjects with bottom quintile numbers. In general, there was an unmistakable correlation (approximately 0.4) between the quintile of a participant's social security number and their willingness-to-pay (see Table 5).

**Network structure.** We used the network structure in Figure 6. The hidden layer had two hidden pools with five units each, representing five quintiles of encountered quantities from low to high. These pools represent two sets of conjunctions between feature units. The input layer had three feature pools, each with five units. The first pool contained units representing evaluations

derived from prior experiences related to the estimate at hand. In the present experiments, these estimates were assumed to be noisy. The second pool contained units representing the estimates of the quantity (i.e., prices of various goods). The third pool contained units representing the value of an anchor—relative to other values that the anchor might take. Units in the estimate pool were influenced by activations in the experience pool and in the anchor pool.

The experience-based estimate “est” (the price estimate corresponding to the 3rd quintile in Experiment 2) 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 units of the anchor pool represented the relative size of the contextual number relative to other potential numbers in the same category. In the context of the last two digits of social security number, we assumed that each unit represented a quintile of the range between 0 and 100. For example, a social security number ending 12 would activate the first anchor unit, and a social security number ending in 94 would activate the fifth unit.

The bidirectional weights between units in each pool were inhibitory (−1). The *i*th unit in the experience pool was connected to the *i*th unit of the first hidden pool (with excitatory +1 weights) and its immediate neighbors (with excitatory +0.5 weights). For example, the 3rd unit in the experience pool was connected to the 3rd unit in Hidden Pool 1 with weight +1, and to the second and fourth units with weight +0.5. Weights from Hidden Pool 1 to the estimate pool followed an identical scheme. The +0.5 weights to neighboring units reflected the noisiness of the representation of experience-based evaluations. Further, the *i*th estimate unit and the *i*th anchor unit were connected to the *i*th unit in Hidden Pool 2 for  $1 \leq i \leq 5$ . These bidirectional weights were all excitatory (+1).

**Network dynamics.** The external input in the central experience unit was assumed to be proportional to the extent of relevant experience a decision maker could bring to bear regarding an estimate (e.g., equal to +1 for estimates that the decision maker had a great deal of information about). The decision maker estimates were noisy and we assumed the activation into the central (third) estimation unit to be equal to 0.1. Neighboring units (i.e., Units 2 and 4 also received activation equal to 0.05, again reflect-

Table 5  
Comparison of Simulated Values and Empirical Price Estimates in Simulation #4

SS quintile	Cordless trackball		Cordless keyboard		Average wine		Rare wine		Design book		Belgian chocolates	
	Simulated value	Observed value	Simulated value	Observed value	Simulated value	Observed value	Simulated value	Observed value	Simulated value	Observed value	Simulated value	Observed value
1	5.63	8.64	12.26	16.09	5.25	8.64	7.57	11.73	6.62	12.82	5.21	9.55
2	12.10	11.82	26.34	26.82	11.29	14.45	16.28	22.45	14.24	16.18	11.20	10.64
3	14.79	13.45	32.20	29.27	13.81	12.55	19.90	18.09	17.40	15.82	13.70	12.45
4	20.17	21.18	43.90	34.55	18.82	15.45	27.13	24.55	23.73	19.27	18.67	13.27
5	37.07	26.18	80.67	55.64	34.59	27.91	49.85	37.55	43.60	30.00	34.31	20.64

Note. SS = social security.




Figure 6. The network structure in Simulation 4: Anchoring. All weights within the same pool are inhibitory ( $-1$ ) and every unit is connected to every other unit in the same pool (in order to promote clarity, not all inhibitory connections are shown). Connections between the pools are excitatory. Solid lines denote a weight of  $+1$ , and dashed lines denote a weight of  $0.5$ . Each unit in the experience pool activates its corresponding unit in Hidden Pool 1 with a connection of weight  $+1$ , and its immediate neighbors with a weight of  $0.5$ . Similar connections are present between the “Estimates” feature pool and Hidden Pool 1. See the online article for the color version of this figure.

ing the noisiness of the estimates.). Units in the anchor pool received an input activation of  $0.6$ .

We ran the network to convergence and used the settled activation values in the units to calculate the model’s willingness to pay.

**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 cdf of the activation value of the unit on a normal distribution ( $M = 0.2$ ,  $SD = 0.05$ ) kept fixed for all estimates. 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.,  $0.7 \times \text{est}$ ). Table 5 describes the simulated and empirically observed value for each item.

**Significance.** This simulation showed that the size of the anchor is proportional to its impact on the estimate of interest. This resulted from interactivity between the anchor pool and the estimate pool. The simulation had an additional notable element: the value of the anchor not only altered the estimate but also had the potential to alter the recollection of experiences that the decision maker may have used to generate an unanchored estimate. For example, let us examine the case in which the unanchored estimate corresponded to  $0.1$  input of the third unit and  $0.07$  activation of the fourth unit (in the experience pool), and the anchor corresponded to  $1$  unit of input in the 5th unit (of the anchor pool). Upon convergence, we observe not only a changed estimate but also a changed recollection of the experiences that were relevant to the unanchored estimate: upon convergence the winning experience unit is the fourth unit (activation  $0.19$  vs.  $0.14$  for the third unit)—even though the third unit had more external input to start with. This feature of the model suggests a mechanism for how events in the world may cause some memories to be more accessible than others in that they are more likely to come to mind (Radvansky, 2017).

### Simulation #5: Negative Auto-Shaping

This simulation featured a context in which intertrial learning effects between previously associated features could result in behavior which was not consistent with expected value maximization.

**Target experiment.** Based on prior work (Brown & Jenkins, 1968), it was known that if a response key was regularly illuminated for several seconds before food was presented, pigeons would reliably peck at the key after a moderate number of pairings, especially if the delivery of food was contingent on the birds pecking at the illuminated key. This behavior was seen as a robust example of instrumental learning in which the pigeons acted in a way to maximize their rewards. Williams and Williams (1969) sought to determine whether the pecking behavior would be maintained even when pecks on the illuminated response key prevented the delivery of food.

They created two groups of pigeons: the first group (the reinforced group) was initially trained via auto-shaping trials in which an illuminated key (displayed for  $6$  s) was followed by the presentation of the food hopper. By pecking on the illuminated key prior to  $6$  s, the pigeons could make the food tray appear instantly. The pigeons in the reinforced group were then additionally trained to actively peck at an illuminated key for food delivery to occur. In (an average of) two sessions consisting of  $50$  trials each, the pigeons were shown an illuminated response key. If they pecked at the key, they were rewarded with the delivery of food via a grain containing hopper. However, if they did not peck at the illuminated key, then the food delivery did not occur. After these sessions, these pigeons were placed in a negative response condition. Now, food delivery was contingent upon the pigeon not pecking at the illuminated key. If they did nothing, they were rewarded at the end of every trial. However, if they chose to peck at the illuminated key, the food delivery did not occur. A second group of pigeons (the naïve group) was directly placed in the negative response condition.

**Empirical results.** Pigeons in the reinforced group (that associated pecking at a lighted key with delivery of food), pecked in



~80% of the trials in the first two sessions of the negative response condition (all data in this simulation has been inferred from figures in the original paper; exact numbers of pecks per session were not provided). After that, most pigeons showed a pattern of recovery (in which the extent of their pecking decreased), and regressions (in which pecking went backup to the early rate of ~80%).

Pigeons in the naïve group did not initially tend to peck on the illuminated key. Across all pigeons in this group, pecking in the first two sessions of the negative response condition was less than ~2%. The surprising result of the experiment concerned the behavior of the naïve group after the first two sessions. Their pressing rate increased (from almost completely absent) to ~50%. This result was particularly noteworthy since pecking led to the absence of food reward, and the optimal behavior for these pigeons was to not peck on the illuminated key (as they had not been doing), and simply wait for the food to arrive. Yet, the pigeons pecked at high rates.

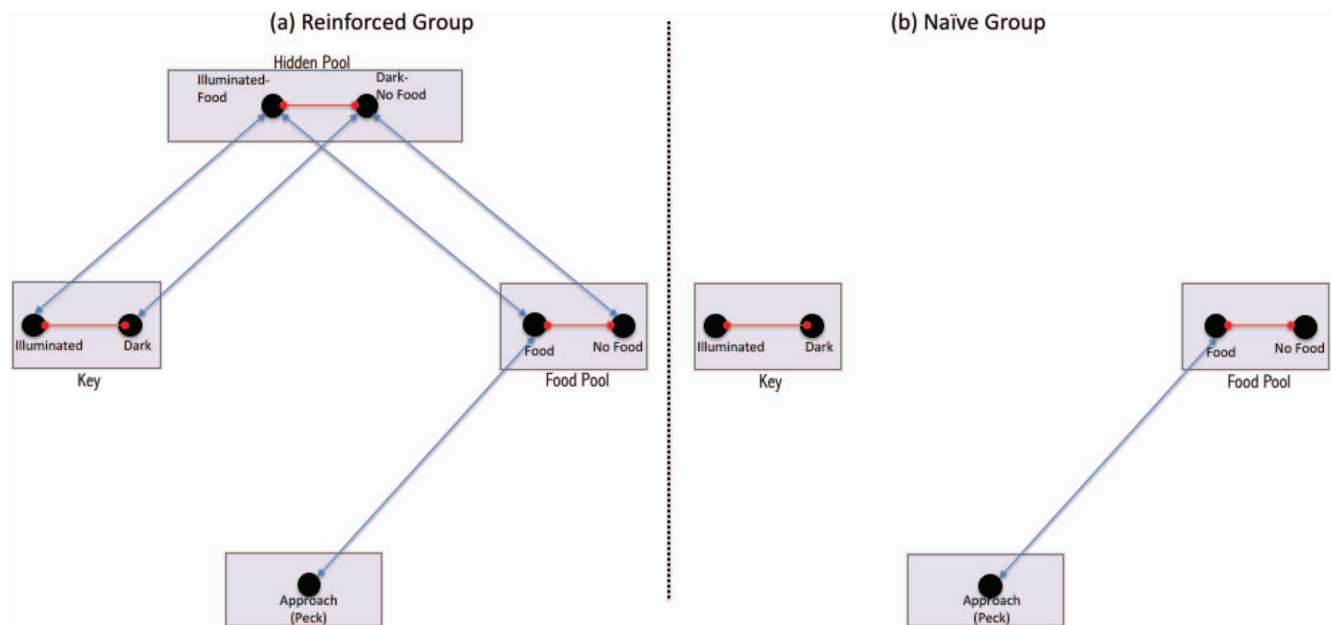
**Network structure.** We separately describe the network structures (prior to the start of the negative response condition) for the reinforced and naïve groups (see Figure 7). The structure for the reinforced group (Figure 7a) involved two hidden units, the first representing a conjunction between an illuminated key and food availability, and the second representing a conjunction between a dark (nonilluminated) key and the absence of food availability. There were two feature pools: a food pool (consisting of food available unit, and a food not available unit), and a key pool

(consisting of a unit representing an illuminated key, and a unit representing a nonilluminated key). A single action tendency unit represented the approach tendency toward the environment. The prototypical pigeon's approach action in the context of food is a particular type of pecking, and this pecking action—directed at the food or at the illuminated key—was represented by the approach unit. Weights between units in the same pool were competitive (equal to  $-1$ ). The feature units were connected to corresponding hidden units with weight  $+1$ . We assumed this connectivity was established in the training phases (auto-shaping and reinforcement). Only the food availability unit was connected to the approach action tendency unit with an excitatory connection ( $+1$ ).

The initial network structure for the naïve group (Figure 7b) was identical to the reinforced group except the features were not associated with each other, since there were no conjunction (hidden) units. Only the food availability unit was connected to the approach action tendency unit with an excitatory connection ( $+1$ ).

**Network dynamics.** In the Reinforced group input in the illuminated key unit (corresponding to the perception of the illuminated key), flowed into the unit representing food availability. This in-turn caused activation in the approach action tendency unit (representing pecking). No such interactive activation occurred for the naïve group network since there were no conjunctions between feature units at the start of the negative response trials.

In the first two sessions, the pigeons in the naïve group gained experience with trials in which an illuminated key was associated with (i.e., preceded) the delivery of food (provided they did not



*Figure 7.* The network structure in Simulation 5: Negative Auto-shaping. In Panel a the pigeons in the reinforced group had an association of food availability with the illuminated key, and an absence of food availability when the key was dark (these associations were via the hidden units). In Panel b the key and the food units were not associated with each other. In Panel a activation in the “illuminated” key unit caused activation to flow to the “food” unit. This caused approach behavior in the form of pecking. In Panel b activation in the “illuminated” key unit did not cause any activation in the approach unit. After a few blocks, the network for the naïve group resembled the reinforced group because of associative learning (A10). See the online article for the color version of this figure.

peck at the key—which was the case in most initial negative response trials). After the first two sessions, this joint activation resulted (Assumption A10) in experience-based connections between the feature pairs (food and illuminated; no food and not illuminated)—albeit via hidden units. This updated the initial network for the Naïve group to become structurally similar to the Reinforced group network depicted in Figure 7a. There was one important difference, however: since the experience resulting in associations between the feature was less than the experience of pigeons in the Reinforced group, we assumed that feature unit/hidden unit connections had a weight of 0.25 (less than the corresponding weights of +1 in the reinforced network).

The gradually increasing connection strength between feature units representing food and feature units representing the illuminated key led to a gradual increase in pecking rates in the Naïve group until it eventually resembled the reinforced group. This transition occurred even though it resulted in behavior that did not maximize expected value.

**Simulation results.** To calculate the percentage of trials in which a pigeon pecked at the illuminated key, we calculated the cumulative density function of the activation value of the approach unit using a normal distribution with a mean (0.515) and standard deviation (0.025) that was kept fixed across all trials, sessions and conditions. These assumptions resulted in (a) simulated pecking rates of 84% in Sessions 1 and 2 for the reinforced group for negative response trials (compared to the empirically observed value of ~80%), (b) simulated pecking rates of 0% in Sessions 1 and 2 for the naïve group for negative response trials (compared to the empirically observed value of ~2%), and (c) simulated pecking rates of 50% in Sessions 3 and following for the naïve group for negative response trials (compared to the empirically observed value of ~50%).

**Significance.** This simulation proposed a plausible mechanism that could cause naïve and untrained pigeons to begin pecking at the illuminated key, even though that action was not optimal (rather, not-pecking was the optimal action). The mechanism relied on experience-based conjunctions developing between the unit representing an illuminated key and the unit representing food availability. Once such conjunctions developed, interactive activation caused approach behavior to occur when the illuminated key unit was provided external input. Associative constraints led to nonoptimal choice.

### Simulation #6: Reward and Path Associations

In the present simulation, our goal was to simulate an empirical context in which the presence of associative constraints led to decision makers not maximizing expected value. Specifically, rats, who had previously learned to associate a high reward with a

particular arm in a maze continued to prefer that arm, even when doing so was no longer optimal.

**Target experiment.** Using a simple T maze, researchers examined the decisions of rats choosing between two courses of action that differed in their respective energetic demands and consequent reward sizes (Salamone, Cousins, & Bucher, 1994; Walton et al., 2006). One arm of the T maze (counterbalanced between rats) was designated the high cost/high reward (HR) arm and the other the low cost/low reward (LR) arm. The location of these arms was kept fixed throughout the experiment. Rewards varied based on the number of food pellets. The cost of obtaining the reward was manipulated by requiring rats to scale a wire-mesh barrier to reach the food.

**Empirical results.** Various configurations of costs and rewards were used to measure rat choices. In a first series of tests, the researchers measured the effects of placing barriers in arms of a maze. In Test 1, there was a 30-cm barrier in the HR arm while the LR arm was unoccupied; in Test 2, an identical 30-cm barrier was present in both arms; in Test 3, the LR arm was again vacant (i.e., did not have a barrier) but the HR now contained a 40-cm barrier. In all three tests, there were four food pellets in the HR arm and two in the LR arm.

In Test 1, the rats selected the HR arm in an average of 67% of trials, and the LR arm in the remaining 33% of trials. In Test 2 (in which both arms had equal barriers), the HR arm was selected in approximately 86% of trials. In Test 3, the increased barrier resulted in the preference for the HR arm dropping to approximately 31% (See Top row of Table 6). The same set of rats participated in all tests

In a subsequent second series of tests, researchers tested the effects of different rewards. Rats chose between climbing a 30-cm barrier in the HR arm or selecting the unoccupied LR arm. However, the reward in the HR arm varied across tests, with the ratio between the HR and LR decreasing from 6:2 (Test 4), to 3:2 (Test 5), and finally to 2:2 (Test 6). In this series of tests, the obstacle was only used in the HR arm. The preference for the HR arm in the three Series 2 tests was approximately 80%, 72%, and 60%, respectively.


There are two notable features of these results: First, it is clear (see Table 6) that both rewards (number of pellets) and action costs (height of barrier) have a profound influence on choice. In all tests, except the last, rats preferred greater rewards to lesser rewards and lower barriers to higher barriers. Second, the rats did not always maximize reward and minimize cost: in the last test, rats frequently chose to climb up a barrier to get 2 pellets, when they could easily have obtained 2 pellets in the LR arm without barriers.

**Network structure.** This simulation featured the structure shown in Figure 8. In the hidden pool, four units represented

Table 6  
*Empirical Data and Simulation Results for Simulation 6*

Condition	4/30 vs 2/0	4/30 vs 2/30	4/40 vs 2/0	6/30 vs 2/0	3/30 vs 2/0	2/30 vs 2/0
Empirical: Preference for HR option	.67	.86	.31	.8	.72	.60
Simulation: Preference for HR option	.64	.80	.20	.73	.57	.63

Note. HR = high reward.



**Figure 8.** The network structure for Simulation 6: Reward and path associations. Panel a depicts the network for Test 1 (a similar network was used for Test 2–5). Increased rewards (i.e., more food pellets) and lower costs (i.e., lower barriers) increase the probability of choice. In Panel b, memory of prior results captured via the “time” feature pool leads to the effect of rats preferring the high reward (HR) option even though it has a barrier to collect two pellets that are available in the left reward (LR) option without the barrier. See text for details. See the online article for the color version of this figure.

conjunctions between the available reward (HR and LR) and temporal aspects related to the trials (Current and Previous). The current units were mutually competitive with each other ( $-1$ ), and the past units were mutually competitive with each other ( $-1$ ).

There were four feature pools. The first pool represented the position of the HR and LR arms. Within each block, these positions were kept fixed for each rat (though varying between rats). In nonchoice trials, rats had acquired knowledge of the location of each arm (e.g., HR arm on the left). The second pool (food pellets) had six units corresponding to six, five, four, three, two, or one food pellets. The third feature pool consisted of three units that represented barriers that were either 30 cm or 40 cm tall, or the absence of a barrier. The fourth feature pool represented time, in which one unit represented the current test, and the other unit represented the influence of past tests. All units in the same feature pool were competitive with each other, units within a  $\pm 2$  neighborhood had a competitive connection of  $-0.5$ , and units further away had a competitive connection of  $-1$ .

The input units were connected to an approach and an avoid unit in the output action-tendency pool with the following weights: the six pellet unit had a  $+0.8$  wt with the approach unit and the 5, 4, 3, and 2 pellet units had connections with the approach unit with positive weights of  $+0.65$ ,  $+0.5$ ,  $+0.35$ , and  $+0.2$ , respectively (a reduction of 0.15 connection weight per food pellet). On the cost side, the 40-cm barrier unit had a connection with the avoid unit of  $+0.6$ , the 30-cm barrier unit had a connection with the avoid unit of  $+0.2$ , and the “no barrier” unit was not connected to the approach or avoid unit.

Figure 8a displays the connection units relevant to Test 1 (four pellets/30-cm barrier vs. two pellets/0-cm barrier). Because this was the first test, there were no connections with the feature unit representing the influence of past trials. The high reward-current conjunction (hidden) unit was connected to the four food pellet

unit and the 30-cm barrier (corresponding to the four of 30-condition). The low reward-current conjunction (hidden) unit was connected to the two food pellet unit and the no barrier unit (corresponding to the two pellets/0-cm barrier condition). A similar network was used for Tests 2–5 with the appropriate linking of foods and costs to the options being tested.

Figure 8b displays the connection units relevant to Test 6 (two pellets/30-cm barrier vs. two pellets/0-cm barrier). This was the final test, and associations from prior tests were assumed to exert an influence on the current choice. We acknowledge that it is possible that some prior test effects influence Tests 4–5, but the design of the experiment does not afford the testing of such possibilities. In Test 6, relevant to the current test (i.e., two pellets/30-cm barrier vs. two pellets/0-cm barrier), the high reward-current conjunction (hidden) unit was connected to the two food pellet unit and the 30-cm barrier (corresponding to the 2/30 condition). The low reward-current conjunction (hidden) unit was connected to the two food pellet unit and the no barrier unit (corresponding to the two pellets/0-cm barrier condition). Relevant to the influence of past tests, the high reward-past conjunction (hidden) unit was connected to the four food pellet unit and the 30-cm barrier (corresponding to the four pellets/30-cm barrier condition which was the most common HR condition in prior tests). The low reward-current conjunction (hidden) unit was connected to the two food pellet unit and the no barrier unit (corresponding to the two pellets/0-cm barrier condition that was the most common LR condition in prior tests).

**Network dynamics.** We discuss the dynamics corresponding to Test 1 (Figure 8a) and Test 6 (Figure 8b) separately. Tests 2–5 had dynamic very similar to those described for Test 1.

In Test 1, consistent with the discussion in Simulation 1, we used two identical, competing networks corresponding to the two options being tested. Thus, for the HR arm in Test 1 (four pellets/

30-cm barrier vs. two pellets/0-cm barrier) we provided input into the current and 30-cm feature units in one network and into the current, no obstacles unit in the competing network. The output action tendency of each approach unit was mutually competitive (as described in Simulation 1).

For the HR network, activation poured into the four pellet unit (no food unit received input since the number of pellets was not observable by the rat) via the HR-current hidden unit. For the LR network, activation poured into the two pellet unit via the LR-current hidden unit. The greater connection weight between the four pellet unit and the approach unit in the HR network, helped the HR option overcome the disadvantage induced by the connection weight between the 30-cm barrier, and the avoid unit. Thus the four pellets/30-cm barrier option was preferred over the two pellets/0-cm barrier option. In Test 2 (four pellets/30-cm barrier vs. two pellets/30-cm barrier), the increased avoid activation in the LR arm increased the proportion of the trials in which the HR arm was selected; this pattern was reversed in Test 3 due to high connection weight between the 40-cm barrier and the avoid unit. Test 4 and Test 5 simulation results (see Table 6) are attributable to the strength of the connection weight between the food pellet units and the approach unit.

The simulation dynamics of Test 6, in which a two pellet reward in the HR condition (featuring a 30-cm obstacle) was preferred over a two pellet reward in the LR condition (with no obstacle), bear a closer look. Here the influence of past tests (represented by the past feature unit) proved influential. In this simulation, the HR arm feature unit, and the 30-cm barrier feature unit activated the HR past hidden unit. This unit, in turn, activated the four pellet unit (and the pas' feature unit). The current hidden units were also activated (similar to Test 1), but the activation of the four pellet unit via the HR-past unit was enough for the rats to prefer the HR condition, even though they had to scale a 30-cm barrier to get 2 pellets of food, which were available in the LR condition without the presence of a barrier.

**Simulation results.** The cumulative density function of positive activation of output units (of a normal distribution, with a mean of 0.25 and a standard deviation of 0.4, kept fixed for all six tests) was used to calculate the choice rates shown in Table 6.

**Significance.** This simulation highlights how associative learning related to prior rewards can lead to nonoptimal decision making in new contexts. It also shows how actions and their associated features—often involving costs—can be integrated into the interactive activation framework.

### Simulation #7: Motivational Influences on Information Processing

This simulation was designed to showcase how the interactivity of the IAC model is not consistent with decision making processes that are serial/modular.

Participants were asked to complete a task ostensibly about differences in predictions of and actual taste experiences (Balcetis & Dunning, 2006). They were shown pictures of two categories of animals worth positive and negative points. In one typical study (Study 2), for half of the participants, farm animals were worth positive points, whereas sea creatures were worth negative points. For the other half of the participants, this was reversed.

If participants ended the experiment with a positive score, they would consume candy, but if their score at the end was negative, participants would consume less desirable canned beans. Participants completed several trials in which one type of animal provided positive points, and the other type of animal provided negative points. As the game progressed, the last three rounds brought increasingly negative point totals and it became ever more suggestive that participants would consume the canned beans. By the final trial, there was precisely one animal who could bring the cumulative points to positive territory. For half of the participants, this animal was a horse (in the farm animal condition); for the other half, it was a seal (in the sea creature condition).

In the final trial, participants were shown an ambiguous image that could be interpreted as either the head of a horse or the full body of a seal (Figure 9b). The experimenters sought to determine whether participants who would benefit from seeing a horse were more likely to see a horse and whether participants who would benefit from seeing a seal were more likely to see a seal.

**Empirical results.** Participants' interpretations depended on what category of animal was worth positive points. When hoping

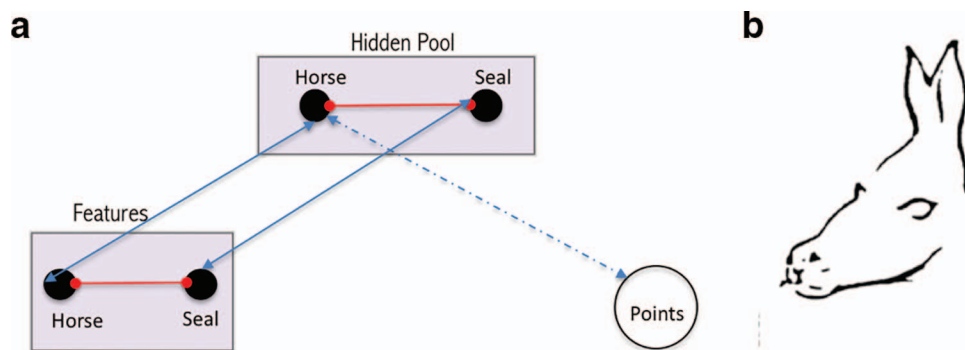


Figure 9. a: Network structure for Simulation 7: Motivational Influences on Information Processing. The network shown here is for a participant in the farm animal condition. The connection between points and the hidden horse unit is developed in the last three trials. A similar structure would apply in the sea creature condition, except the points feature unit would be connected to the seal hidden unit. b: The ambiguous horse/seal figure used by Balcetis and Dunning (2006). From "Ambiguity of Form: Old and New," by G. H. Fisher, 1968, *Perception & Psychophysics*, 4, p. 191. Copyright, 1968 by the Psychonomics Society. Figure used with permission. See the online article for the color version of this figure.



to see a horse, 66.7% of participants saw the ambiguous figure as a horse, and 33.3% saw a seal. However, this bias reversed when participants were hoping to see a seal. Only 27.3% of this group saw a horse, but 72.7% reported a seal.

**Network structure.** Figure 9a shows the network structure used in the farm animal condition. Two competing input feature units in the feature pool were activated with evidence consistent with either a horse or a seal. There was also a single unit representing the possibility of points. Two units in the hidden pool corresponded to a horse or seal. In the farm animal condition, the points unit developed a connection to the horse hidden unit (and similarly to the seal in the sea creature condition). These associations were developed in the three penultimate trials in which it became clear that only one animal would be worth sufficient points to bring the points total to positive territory.

**Network dynamics.** At the time of input, equal external activation (0.2) was given to features corresponding to both the horse and the seal. The points unit was also provided activation (1). In the farm animal condition, this activation flowed into the horse feature unit and increased visual processing related to a horse. Corresponding flows occurred in the sea animal condition.

**Simulation results.** As shown in Table 7, activation in the feature unit for the horse was high (0.63) when the horse could provide the desired points; conversely, when recognizing a seal provided points, activation in the feature unit for the seal was symmetrically high. The probabilities in Table 7 were obtained by calculating the cumulative distribution function of each activation level on N (0.4, 0.4).

**Significance.** This experiment and the mechanism described in the simulation suggested that perception and valuation-related processing may interact with each other and value-based decision making need not be serial or modular. Participants preferred to consume jelly beans over canned beans (a value-based decision). This preference motivated them to end the empirical game with positive points, which in turn influenced how they processed the ambiguous image. If perception and valuation were serial and modular processes, one would not expect (putatively downstream) value-related preferences to influence (putatively upstream) perception-related processes.

### Simulation #8: Dynamic Processes in Goal-Directed Choice

The prior simulations developed computational evidence to suggest that value-based decision making need not always maximize expected value, and its associated processes need not be serial/modular. In this simulation, we sought to demonstrate that the IAC, as constructed, can also uncover dynamic mechanisms un-

derlying goal-directed choice, and thus have the potential of serving as a general framework for value-based decision making.

We sought to examine whether effects consistent with the IAC framework could be observed in participant mouse-tracks as they decided between tasty and healthy (goal congruent) food items. Experimenters have frequently used mouse tracking paths to detect dynamic patterns in the choice process (Freeman & Ambady, 2011). Mouse tracking experiments provide insight into the unfolding of the choice process (in addition to recording the choice itself). For example, the degree of curvature in the mouse-track is thought to represent the spatial attraction toward the nonchosen option (Freeman & Ambady, 2011; Gillebaart, Schneider, & De Ridder, 2016; Spivey & Dale, 2006).

**Target experiments.** Mouse tracking trials in the domain of food choice required participants to choose between a tasty item and a healthy item by dragging their mouse toward their selection (Sullivan, Hutcherson, Harris, & Rangel, 2015). Such choices were recorded over multiple food pairs in multiple trials. Experimenters wished to measure speed with which the decision-making circuitry processes basic attributes like taste, versus more abstract, goal-related attributes such as health.

**Empirical results.** Many mouse tracks in which the healthy option is accepted revealed a pattern in which the decision maker, starting from a neutral position, initially swung toward the tasty item (or toward a response button indicating liking of the tasty item), and then veered toward the healthy item (or toward a response button indicated disliking of the tasty item). In cases where the tasty item was selected or preferred, this veering did not take place and the decision maker tracked a direct path from the neutral position to the tasty item/preference.

Across all trials, the influence of the healthy option was detected later than the influence of the tasty option (using a regression featuring the trajectory of the mouse track), even when the healthy option was preferred.

**Simulation approach.** We sought to determine whether we could model the empirically observed phenomenon of the earlier influence of taste-related variables and the later influence of goal-congruent health-related variables. To do this we constructed a model in which a decision-maker had to choose between a tastier but less healthy item (hamburger), and a less tasty but healthier item (salad).

**Network structure.** Some feature units (e.g., flavor, taste, and nutrition) represent properties of the food item (i.e., hamburger or salad) and are connected to their corresponding hidden representations and to appropriate output approach/avoid action tendencies. Importantly, as shown in Figure 10, some of these units (e.g., those representing high nutrition, but not those representing low nutrition) are connected with the abstract representation of healthy. Importantly, the abstract representation of healthy is a hidden unit in a pool of a separate subnetwork (i.e., not a direct part of the network representing the choice between the hamburger and the salad). This hidden unit healthy may be connected to features such as “optimal weight” or “feeling energetic,” both of which may have connections with the approach action tendency unit (for consideration of an alternative structure, please see online supplementary materials).

**Network dynamics.** As in choice-related simulations involving two separate options, we used two copies of the network. In each copy of the network, we input one unit of

Table 7  
*Simulation 7, Motivational Influences on Visual Processing: Results*

Condition	Empirical data		Simulation	
	Horse	Seal	Horse	Seal
Farm animal → points	66.7%	33.3%	71.7%	28.3%
Sea creature → points	27.3%	72.7%	28.3%	71.7%

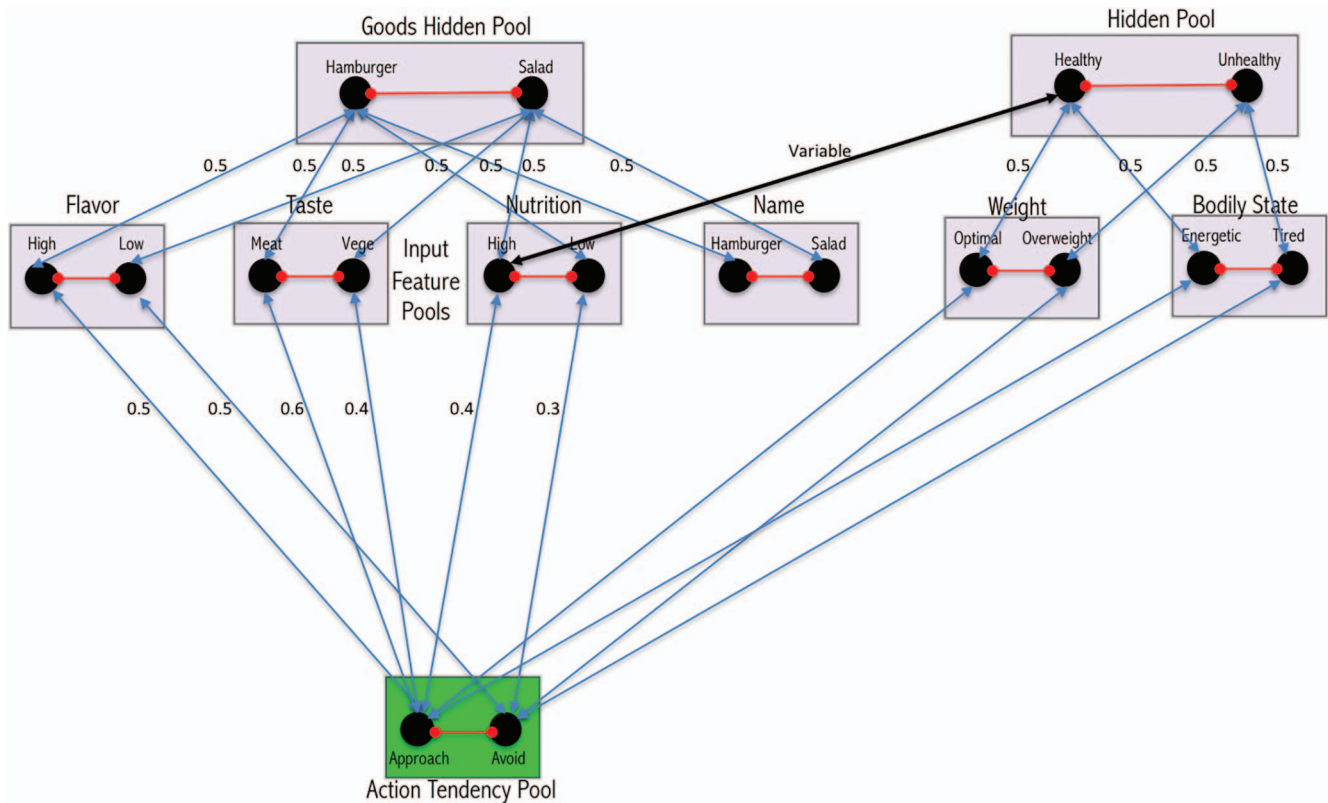


Figure 10. The network structure and dynamics for Simulation 8: Dynamic processes in goal-directed choice. We simulate the choice between a tasty/less healthy hamburger and a less tasty/healthy salad. Two copies of the network run in parallel; one network copy receives input in the hamburger unit of the name pool and the other copy receives input in the salad unit of the name pool. In both cases, activation flows into the respective hidden units (hamburger or salad) and then into the corresponding features. Only in the salad network, activation flows from the “high” nutrition unit to another subnetwork related to the abstract goal of health. In principle, there may be other connections between the two subnetworks (e.g., between the low nutrition unit and the unhealthy unit), but we focus on one for simplicity. See the online article for the color version of this figure.

activation into one of the “name” units (either hamburger or salad). This activation flowed—via the hidden units—to the respective feature units, including into the taste and nutrition pools. Initially, because of the higher weight between the high taste/flavor units and the approach unit (compared to the weight between the high nutrition unit), the preference for hamburger (i.e., more favorable weights to output units compared to weights between salad features and output units) made hamburger the more preferred choice over salad.


However, the extended network related to the healthy unit made its presence felt later in the process. Activation from the salad high nutrition feature unit flowed into the healthy hidden unit (see Figure 10). This unit activated the optimal weight unit and the energetic bodily state unit; these units in-turn strengthened the approach unit. The hamburger feature units were not connected to the healthy unit and thus did not receive a corresponding activation boost in its approach unit. This activation flow enabled activation into the approach unit of the salad network to overcome its early disadvantage.

Importantly, this activation pattern only occurred for networks in which the connection between the high nutrition unit and the

healthy unit had sufficient weight. When this connection was absent or had a low weight, the early advantage of the hamburger over the salad was never overcome, and the decision maker chose the hamburger. Related proposals have highlighted the asymmetry of temptation-related associations and goal-related associations (Fujita, 2011; Fujita & Sasota, 2011).

**Simulation results.** As shown in Figure 11, a connection strength of +1 between the nutrition feature unit and the healthy hidden unit was enough for salad to be preferred over hamburger. For connection strengths less than or equal to 0.5, the hamburger was chosen. In cases in which the healthy item (salad) was chosen, its approach activation initially lagged the approach activation for the hamburger but overtook it later (the crossover occurred at cycle number 79).

**Significance.** The simulation captured two important elements of health-related goal-directed behavior: first, it offered an explanation of why such behavior occurs in some cases, but not in others (i.e., due to differences in weights between health-goal-related units and feature units); second, it offered a mechanism to explain why health-goal-related features may come online later than more salient taste related features (i.e., due to



*Figure 11.* Activation related to approaching the hamburger and salad for low and high nutrition-feature/healthy unit connection weights. At a connection weight of 0.5, hamburger is preferred over salad (higher activation for the hamburger approach unit). At a weight of 1.0 salad is preferred hamburger (higher activation for salad approach unit). In the latter case, even though salad has a higher activation at convergence, the activation for hamburger is higher in earlier network cycles. There is a cross over at Cycle #79. Thus, consistent with empirical data, features related to taste and health come online sooner than features related to health. See the online article for the color version of this figure.

indirect activation into the ‘healthy’ unit). Importantly, the representation related to the goal of being healthy in this simulation did not require the computation of a desired state (which presupposes the capability to calculate value and/or stipulates the target of its explanation). On our account, any representation, if appropriately connected to feature units, can result in goal-directed behavior.

### Contributions of the Simulations

Collectively, our simulations offer a concrete mechanism related to how decisions may be nonoptimal and how their process is not serial/modular. This is their central purpose.

In some cases, the simulations also make counterintuitive predictions about behavior, and these predictions are supported by gathering evidence. We next highlight three such examples: First, unlike prior literature (Dinner et al., 2011), our model (Simulation 2) proposes that many default effects are not caused by factors such as loss aversion or implied recommendations of the experimenter; rather, on our account, default effects often occur due to insufficient attention toward the available nondefault option. Our model predicts that increasing activation toward the nondefault option should decrease default effects. Gathering evidence supports this view (Suri & Gross, 2015). Second, our framework, for the first time to our knowledge, provides a mechanistic account of how prior estimates of a certain quantity and levels of attention to unrelated anchors interact with each other to produce the anchoring effects (Simulation 4). Finally, Simulation 8, for the first time to our knowledge, provides a mechanistically explicit, neural network-based account of how some goal-related behaviors might be activated by particular features of one of the items being evaluated. These features may in-turn activate a goal-congruent

subnetwork (see Figure 10). There is emerging evidence for this view of goal-directed behavior (e.g., Fujita & Sasota, 2011).

### General Discussion

In this article, we have described an IAC framework for value-based decision making. In this section, we first summarize the key features of this framework and highlight its important benefits. We will then consider how the IAC framework implies that choice may be emergent. Next, we will outline two opportunities for extending the scope of the IAC framework. Finally, we will discuss important limitations and future directions related to our approach.

### A General Framework for Value-Based Decision Making

We began this work to consider alternatives to the view that value-based choices are governed solely by their utility and are made via a serial and modular process. We specified the IAC framework in which choice emerges via the reciprocal interaction between units of the network. In this framework, utility-related variables can play a part in shaping choice outcomes—but so can context variables that are unrelated to utility. The process by which choice emerges in this framework is not serial; rather all units can interact with each other and influence each other throughout the decision-making process.

Because the IAC framework spans the key elements of value-based decisions—that is, option benefits, action costs, contextual variables, and goals—we propose that the IAC framework may be considered as a viable framework for the domain of value-based decision making as a whole.

## Benefits of the IAC Approach

Marr (1982) famously posited three distinct levels of analysis of information-processing systems. This first, the computational level, is a description of what information processing problem is being addressed. The second, the algorithmic level, is a description of what mechanistically explicit steps are being carried out to solve the problem. The third, the implementational level, is a description of the physical characteristics of the information processing system (e.g., the brain).

Some theories of value-based decision making, such as those in the Bayesian tradition (e.g., Solway & Botvinick, 2012), have advanced theories predicated upon objective function maximization. These theories have used biologically plausible neural networks (at the algorithmic level) that are geared toward satisfying the core constraint of objective function maximization at the computational level. They therefore effectively assume that their proposals at the computational level can be implemented unchanged at the algorithmic level. As such, these approaches assume a “triumphant cascade” (Dennett, 1987, p. 227) through Marr’s three levels so that constraints identified at the computational level (via the objective function) are the main constraints at the algorithmic level and implementational level. We propose that there is no reason that specifications at the computational level should be privileged over specifications from other levels, and assuming so may lead to frameworks that are incomplete or require post hoc just-so stories to provide an explanation purportedly grounded in an optimization framework. Although this approach can lead to insights into previously neglected factors that could be subject to optimization, it can also, we argue, obscure recognition of the importance of other factors that could contribute to behavior.

Unlike many other approaches to value-based decision making, the interactive activation approach is based on a fully specified process that resides in a network that is biologically informed. It is thus at-least partially constrained at all three levels of analysis. Although the IAC model is fully specified at the algorithmic level, its constraints at the biological level are weaker. It is not intended to fully capture the biological processes underlying value-based decision making. Rather, it is an exploration of ideas about how choice develops. In this exploration, simplification is essential, because simplification enables a closer look at the consequences of the central ideas being put forward.

Constraints at multiple levels of analysis offer several advantages. First, the assumptions underlying our framework are less reliant on a set of phenomenological intuitions (e.g., maximization of value) because, in addition to being computationally specified, these assumptions must also satisfy the requirements of the specified algorithm and the physical system in which the processing takes place. With fewer degrees of freedom, such sets of assumptions are more likely to be more parsimonious and have explanatory power at multiple levels of analysis.

Second, our approach is less likely to stipulate what we’re trying to explain as part of the explanation. Some theories implicitly smuggle in new constructs that explain value-based decision making using constructs of parallel complexity. For example, value-based decisions are often said to be made in the pursuit of valued goals (e.g., Austin & Vancouver, 1996). This

begs the question of why (and how) some goals are deemed valuable in the first place. The IAC approach proposes a mechanism for goal pursuit that does not rely on the prior overall valuation of a desirable end state. In the context of Simulation 8, for example, one need not represent the overall idea of eating a salad to become healthy to pursue a health-related goal. Rather interactive activations between activated features can lead to behavior congruent with goal pursuit.

Finally, the incorporation of algorithmic constraints reduces opportunities to implicitly rely on a homunculus that makes trade-offs, is sensitive to transient states, and pursues goals. Postulating a homunculus is universally recognized as unscientific, and no scientist explicitly refers to a homunculus to explain behavior. However, as Hazy, Frank, & O’Reilly (2006) have noted, psychological models that exclusively stay at a computational level often tacitly evoke a homunculus since the capabilities they assign to the mind (at the computational level) are unconstrained. The interactive activation model, with its transparent assumptions, makes it less likely for a homunculus to be inadvertently used to understand value-based decision making.

## Implications of Interactive Activation: Choice Is Emergent and Value May Not Be Represented

An important implication of the interactive activation approach is that the action of choosing one alternative over another, or of deciding whether or not to take a particular action, is a consequence of the input into the network and the nature of connections between the units of the network. Choice emerges via a distributed interactive process among diverse populations of units that are subject to multiple influences and is, therefore, an emergent property of the system (Chialvo, 2010). Furthermore, there is no separate computation of value, in that no units in the network are making calculations exclusively related to the calculation of value independent from other influences.

The units for approaching and avoiding a choice alternative that we use in our models may initially appear to represent value, and they are indeed influenced by value-related variables, but they do not represent value alone. Indeed, in our simulations, they are influenced by other factors, such as the recent activation of the particular choice response associated with an option. Thus, in our models, there is no pure representation of value.

The observation that there are many parts of the brain that are active in choice situations, and whose activation may be associated with some measure of value is consistent with this view since these activations may not represent value per se or may be affected by variables other than value. These regions may, for example, code representations related to features of items, features of actions required to obtain those items, active needs, goal representations, and action tendencies—all of which are related to value, but separable from it.

Even brain regions previously thought to code for a ‘pure’ value signal such as the ventromedial prefrontal cortex and the orbitofrontal cortex may in fact be subject to influence by an array of alternative nonvalue related factors and processes, such as outcome-identity coding, informational coding, and encoding of autonomic and skeletomotor consequences, and may also be subject to salience or attentional effects (Knutson, Taylor, Kaufman,



Peterson, & Glover, 2005; O'Doherty, 2014). Further, consistent with the IAC framework in which many units participate in the emergence of preference-based choice, there are numerous correlates of aspects of value in many other brain regions including the insula, the dorsal striatum, the anterior and posterior cingulate cortices, the ventrolateral and dorsolateral parts of the prefrontal cortex, the sensory cortex, the motor cortex, and the Intra Parietal Sulcus (IPS) (Hunt & Hayden, 2017). For example, an fMRI study using multi-voxel-patterns-analysis (MVPA) showed that 30% of all voxels showed a statistically significant ability to decode value, and that value related signals were observable in the gray matter of nearly every region of the brain (Vickery, Chun, & Lee, 2011) consistent with the proposition that the brain makes value-related decisions via an activation-fueled distributed consensus, rather than by computing and maximizing value per se.

Importantly, we are not claiming that distributed processing implies a complete lack of specialization. On the contrary, certain brain regions likely do specialize in certain preference-related tasks, and this specialization may be a function of the anatomical inputs received by a region (Neubert, Mars, Sallet, & Rushworth, 2015). For example, OFC neurons may specialize in food-related stimulus features (Padoa-Schioppa, 2011), whereas neurons in the anterior cingulate may specialize in processing related to action features such as physical effort (Hosokawa, Kennerley, Sloan, & Wallis, 2013). Although the inputs and therefore the factors affecting activations in different regions may differ, our theory holds that these activations nevertheless reflect a range of factors, and indeed the evidence is consistent with this, as indicated above.

Some brain areas such as the vmPFC do often show a linear relationship between the choice of an item on offer and BOLD activation related to it. However, this need not imply that the vmPFC codes for explicit value representation. For example, if the vmPFC coded for action tendencies in value-based decisions (corresponding to the approach/avoid units in the IAC network), then in situations in which value-related constraints determined action tendencies, the vmPFC would appear to code for a value signal. However, in situations in which other variables influenced action tendencies, the vmPFC would best be viewed as coding for an action integration signal which would be distinct from a value signal (since it would include the influence of contextual variables). Consistent with this view, evidence implicates the vmPFC in an integrative role that includes value related factors such as reward, as well as contextual factors such as recent behavior and fatigue (San-Galli, Varazzani, Abitbol, Pessiglione, & Bouret, 2018).

### Opportunities for Extending the IAC Framework

The interactive activation approach may provide at least two opportunities for extending its breadth of application. First, future computational evidence may make it possible to include goal-directed and habitual behavior under a single integrated framework. Second, it may be possible to use the IAC principle of conjunctive association to include the effects of physiological needs on decision making. We consider each of these two points in turn.

**Goal-directed and habitual behaviors unfold within the same system.** Several prior accounts have proposed that there are qualitatively different types of decisions and valuation systems

(Balleine & O'Doherty, 2010; Dickinson, 1985; Kahneman, 2011). Furthermore, different valuation systems are thought to separately process goal-directed and habitual behaviors and actions (Dolan & Dayan, 2013; Rangel et al., 2008). An underlying assumption of such accounts is that many behaviors fit into one category—they are either goal-directed or habitual.

In the interactive activation approach, goal-directed behaviors occur because of weights between feature units and action tendency units, or the influence of goal representations in memory. Habitual behaviors occur because deepening experience of two associated representations is manifested via increasing weights between units. For example, in Simulation 6, rats traversing a maze tended to prefer larger pellet rewards over smaller ones. This was an instance of goal-directed behavior and was driven by connections between units representing food pellets and units representing 'approach' tendencies. At the same time, rats showed a tendency to return to the previously highly rewarding arm of the maze, even though it was no longer optimal to do so. This is an example of habitual behavior and was driven by associativity between the representation of the high-reward arm and a larger number of food pellets. In Simulation 6, both goal-directed behavior and habitual behavior unfolded within a single, interacting system.

Development of the IAC approach may more generally show that goal-directed and habitual behaviors unfold in the context of the same framework and are guided by the same set of assumptions. They may not require processing via different valuation systems.

**Physiological needs.** Ubiquitous everyday experiences demonstrate that our decisions vary with our physiological needs. In a very thirsty state, we may choose water over any other item, but in a nonthirsty state this may not be the case at all. For present purposes, we constrain the set to include only those needs that are intrinsic (i.e., present from birth in some form), evolutionarily conserved across species, and unmistakably influential in an animal's value-based decisions.

We believe that it is possible to use the principles of interactive activation to include the effect of such physiological needs. This could be enabled via associating features with units representing active physiological needs via a conjunctive pool (similar to the hidden pool). This would cause weights between feature units and action tendency units to vary depending on whether or not a physiological need is active (see Read, Smith, Droutman, & Miller, 2017 for a similar approach).

### Limitations and Future Directions

Despite its explanatory power, the IAC framework, as described above, has some important limitations. We believe that future work pertaining to these limitations (and their potential mitigations) will deepen and broaden the applicability of the interactive activation approach in the domain of value-based decision making. In this section, we will consider three limitations of the interactive activation network: (a) its localist (as opposed to distributed) nature, (b) the requirement that multiple instances of the network—one corresponding to each available option—must run in parallel to simulate multiitem choice, and (c) the fact that in the model we stipulate the values of connection strengths, rather than specifying processes whereby they might be acquired through

experience. In each case, we will propose a future research direction that may help mitigate this limitation.

**Localist nature of the interactive activation network.** A localist network has the property that its units can represent a cognizable concept. In contrast, a distributed network represents concepts as patterns of activity over a collection of units—and therefore more than one unit is required to represent a concept. Further, each unit participates in the representation of more than one concept (Hinton, McClelland, & Rumelhart, 1986; Plate, 2002).

Individual connections between units in a localist network mediate meaningful associative relationships. In distributed models, however, the situation is more complex. In such systems, if one wishes to associate, for example, the taste of pizza with the sight of pizza, and if the taste and sight are each represented as a pattern of activation over a set of units, then the connection weight changes required to store the association may involve many, or even all, of the weights involved in other associations (McClelland & Cleeremans, 2009; McClelland & Rumelhart, 1985).

Some have argued that the brain relies on localist representations (Bowers, 2009) especially in the medial temporal lobe (Roy, 2012), while others argue that representations in the brain are generally distributed (Plaut & McClelland, 2010). Our perspective, shared with others (e.g., Kumaran, Hassabis, & McClelland, 2016), is that representations fall on a continuum of degree of overlap, with ‘localist’ representation being a simple and intuitively graspable approximation that can be useful as a guide for explanation.

The interactive activation network is a localist network and therefore does not directly rely on distributed representations. However, the mechanisms crucial to the operation of our network—including interactive activation, competition, and reciprocity—are applicable in distributed networks in a manner similar to how they operate in the network developed in this work. For example, very early work by Anderson and colleagues (1977) illustrates how competition between entire patterns of activation can occur in networks of densely interconnected neurons.

In a sense, localist networks (including the interactive activation network) may be thought of as approximate characterizations of more complex distributed networks. They provide the convenience of being able to render all of the dynamics in terms of conceptual entities, rather than in terms of the individual neuronal-level dynamics (Smolensky, 1986). They are therefore easier to comprehend than essentially equivalent distributed networks. For some purposes, localist networks may not be the best way to capture a phenomenon of interest, however: as modelers, we choose a particular level of description to suit a particular purpose. In our context, the special strengths of distributed networks (e.g., capacity to simulate gradual degradation of cognitive capacity) were unnecessary; further, the relative conceptual and computational simplicity of a localist network supported our objective of expositional clarity. We therefore chose the localist approach, despite its limitations.

Nevertheless, we acknowledge that the brain is unlikely to follow purely localist representations. Future research should endeavor to build fully distributed models. Such models are likely to be constrained by implementation level considerations to a greater extent than the present work.

**Multiple instances of the same network.** When simulating choice, say between a Coke and a V8, our approach requires the instantiation of two copies of the same network. In this example, one copy of the network receives and processes the input features of Coke and the other receives and processes the input features of V8. The network that converges to higher activation in its action tendency represents the “winning option.”

We do not think it likely the brain literally contains multiple copies of the same network. One prior approach to solve a similar issue (in the reading domain), known as the *programmable blackboard*, McClelland (1986), focused on creating a central repository of knowledge that could be made available for processing different stimuli presented at different locations in the visual field. However, it is unclear whether proposals similar to the programmable board can in fact be implemented in the brain. Another possibility is that alternatives are evaluated in alternation, and there is some evidence consistent with this possibility. For example, ensembles of choice-encoding neurons (e.g., in the OFC) appear to alternate in the evaluation of each of the two items, shifting activity patterns as the network evaluates each option (Rich & Wallis, 2016). An exploration of how such alternation could mimic parallel instantiation, as well as a consideration of other implementations that might capture the activation and competition processes we have modeled by instantiating multiple copies of the same network, remain to be more fully explored.

**A nonlearning network.** Our account has not addressed the learning processes that lead to establishing and adapting the strengths of the connections between network units. In most cases, we have stipulated connections that might plausibly have been established either by developmental programs or experience-dependent learning processes, whereas in others we have assumed that connections between neurons that are active at the same time increase the strengths of their mutual connections (in accordance with Hebb’s (1949) postulate), without actually implementing an activity-dependent learning rule. The details of the rules of synaptic plasticity remain a rich and ongoing area of research both in the computational (LeCun, Bengio, & Hinton, 2015) and the biological sciences (Bliss, Collingridge, & Morris, 2003). Our underlying assumption has been that the brain can and does flexibly develop excitatory and inhibitory weights of varying strengths between neurons, and our focus in the present work has been to offer an existence proof that it is possible to construct a framework with simple, transparent, and biologically informed assumptions that can simulate the outcomes and dynamics of a broad range of behavior related to value-based decision making. Our simple approach here is a limitation, however, even for the goal of accounting for overt behavior. The details of synaptic plasticity rules will be important for understanding the exact conditions under which prior activation results in strengthening versus weakening of connections, and these conditions will affect the details of the patterns of behavior that are observed when people make choices in value-based decision-making settings. Although we do not yet know exactly the rules the brain uses for connection weight changes, we note that negative auto-shaping and many of the other phenomena we have simulated support the view that co-occurrence, and not simply outcome maximization, contributes to the strengths of the connections that influence the value-based choices we and other organisms make.

Another limitation of the present work related to learning concerns experiments in which decision makers must rapidly react to changing contingencies. Complementary learning systems (CLS) theory (Kumaran et al., 2016) suggests one mechanism by which rapid contingency adjustments may occur in frameworks featuring interactive activation. According to CLS, there are two complementary learning systems, one instantiated primarily in the neocortex and the other in the medial temporal lobes. The first gradually acquires structured knowledge representations while the second quickly learns the specifics of individual experiences. Importantly, neocortical learning can be rapid for information that is consistent with known structure (McClelland, 2013; Tse et al., 2007). This suggests that it is possible that interactive activation-based mechanisms, when combined with the CLS, are able to computationally model situations in which a change in action-outcome contingencies causes a decision maker to rapidly adjust behavior. Future work is required to specify and develop such models.

### Concluding Comment

We believe that the purpose of models is to explore the implications of ideas. Here, we have developed a model based on a fairly small set of simple ideas and have used these ideas to simulate a broad range of empirical phenomena in value-based decision making. The scope of these simulations invites a reconsideration of the view that motivated behaviors are undertaken in the service of maximizing expected value, and that the process of decision making proceeds via the serial processing of functionally distinct submodules. It also suggests that interactive activation is a viable framework to integrate and inform the field of value-based decision making.

### References

- Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review*, 84, 413–451. <http://dx.doi.org/10.1037/0033-295X.84.5.413>
- Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, 118, 73–106. <http://dx.doi.org/10.1162/00335530360535153>
- Armell, K. C., Beaumel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, 3, 396–403.
- Atkinson, J. W. (1957). Motivational determinants of risk-taking behavior. *Psychological Review*, 64, 359–372. <http://dx.doi.org/10.1037/h0043445>
- Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological Bulletin*, 120, 338–375. <http://dx.doi.org/10.1037/0033-2909.120.3.338>
- Balceris, E., & Dunning, D. (2006). See what you want to see: Motivational influences on visual perception. *Journal of Personality and Social Psychology*, 91, 612–625. <http://dx.doi.org/10.1037/0022-3514.91.4.612>
- Balceris, E., & Dunning, D. (2010). Wishful seeing: More desired objects are seen as closer. *Psychological Science*, 21, 147–152. <http://dx.doi.org/10.1177/0956797609356283>
- Balleine, B. W., & O'Doherty, J. P. (2010). Human and rodent homologies in action control: Corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*, 35, 48–69. <http://dx.doi.org/10.1038/npp.2009.131>
- Barkan, R., Ayal, S., & Ariely, D. (2016). Revisiting constructed preferences: Extrapolating preferences from relevant reminders. *Decision*, 3, 281–294. <http://dx.doi.org/10.1037/dec0000051>
- Bell, D. E. (1988). *Decision making: Descriptive, normative, and prescriptive interactions*. New York, NY: Cambridge University Press. <http://dx.doi.org/10.1017/CBO9780511598951>
- Bliss, T. V. P., Collingridge, G. L., & Morris, R. G. M. (2003). Synaptic plasticity in health and disease: Introduction and overview. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 358, 607–611. <http://dx.doi.org/10.1098/rstb.2003.1282>
- Boureau, Y. L., Sokol-Hessner, P., & Daw, N. D. (2015). Deciding how to decide: Self-control and meta-decision making. *Trends in Cognitive Sciences*, 19, 700–710. <http://dx.doi.org/10.1016/j.tics.2015.08.013>
- Bowers, J. S. (2009). On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological Review*, 116, 220–251. <http://dx.doi.org/10.1037/a0014462>
- Brown, P. L., & Jenkins, H. M. (1968). Auto-shaping of the pigeon's key-peck. *Journal of the Experimental Analysis of Behavior*, 11, 1–8. <http://dx.doi.org/10.1901/jeab.1968.11-1>
- Camerer, C. F., Loewenstein, G., & Rabin, M. (Eds.). (2011). *Advances in behavioral economics*. Princeton, NJ: University Press. <http://dx.doi.org/10.2307/j.ctvc4m4j8j>
- Chialvo, D. R. (2010). Emergent complex neural dynamics. *Nature Physics*, 6, 744–750. <http://dx.doi.org/10.1038/nphys1803>
- Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2004). For better or for worse: Default effects and 401 (k) savings behavior. In David A. Wise (Ed.), *Perspectives on the economics of aging* (pp. 81–126). Chicago, IL: University of Chicago Press. <http://dx.doi.org/10.7208/chicago/9780226903286.003.0003>
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407–428. <http://dx.doi.org/10.1037/0033-295X.82.6.407>
- Dehaene, S., & Mehler, J. (1992). Cross-linguistic regularities in the frequency of number words. *Cognition*, 43, 1–29. [http://dx.doi.org/10.1016/0010-0277\(92\)90030-L](http://dx.doi.org/10.1016/0010-0277(92)90030-L)
- Dennett, D. (1987). *The intentional stance*. Cambridge, MA: The MIT Press.
- Dickinson, A. (1985). Actions and habits: The development of behavioural autonomy. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 308, 67–78. <http://dx.doi.org/10.1098/rstb.1985.0010>
- Dinner, I., Johnson, E. J., Goldstein, D. G., & Liu, K. (2011). Partitioning default effects: Why people choose not to choose. *Journal of Experimental Psychology: Applied*, 17, 332–341. <http://dx.doi.org/10.1037/a0024354>
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80, 312–325. <http://dx.doi.org/10.1016/j.neuron.2013.09.007>
- Eccles, J. S., & Wigfield, A. (1995). In the mind of the actor: The structure of adolescents' achievement task values and expectancy-related beliefs. *Personality and Social Psychology Bulletin*, 21, 215–225. <http://dx.doi.org/10.1177/0146167295213003>
- Ehret, P. J., Monroe, B. M., & Read, S. J. (2015). Modeling the dynamics of evaluation: A multilevel neural network implementation of the iterative reprocessing model. *Personality and Social Psychology Review*, 19, 148–176. <http://dx.doi.org/10.1177/1088868314544221>
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press. <http://dx.doi.org/10.7551/mitpress/4737.001.0001>
- Freeman, J. B., & Ambady, N. (2011). A dynamic interactive theory of person construal. *Psychological Review*, 118, 247–279. <http://dx.doi.org/10.1037/a0022327>
- Friedman, M., & Friedman, M. (1953). *Essays in positive economics*. Chicago, IL: University of Chicago Press.



- Fujita, K. (2011). On conceptualizing self-control as more than the effortful inhibition of impulses. *Personality and Social Psychology Review*, 15, 352–366. <http://dx.doi.org/10.1177/1088868311411165>
- Fujita, K., & Sasota, J. A. (2011). The effects of construal levels on asymmetric temptation-goal cognitive associations. *Social Cognition*, 29, 125–146. <http://dx.doi.org/10.1521/soco.2011.29.2.125>
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin Harcourt.
- Gillebaart, M., Schneider, I. K., & De Ridder, D. T. (2016). Effects of trait self-control on response conflict about healthy and unhealthy food. *Journal of Personality*, 84, 789–798. <http://dx.doi.org/10.1111/jopy.12219>
- Glimcher, P. W. (2008). Neuroeconomics. *Scholarpedia*, 3, 1759. <http://dx.doi.org/10.4249/scholarpedia.1759>
- Glimcher, P. W., Camerer, C. F., Fehr, E., & Poldrack, R. A. (2009). Introduction: A brief history of neuroeconomics. In P. W. Glimcher, C. F. Camerer, E. Fehr, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision making and the brain* (pp. 1–12). Oxford, UK: Academic Press. <http://dx.doi.org/10.1016/B978-0-12-374176-9.00001-4>
- Glöckner, A., Betsch, T., & Schindler, N. (2010). Coherence shifts in probabilistic inference tasks. *Journal of Behavioral Decision Making*, 23, 439–462. <http://dx.doi.org/10.1002/bdm.668>
- Glöckner, A., & Herbold, A. K. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *Journal of Behavioral Decision Making*, 24, 71–98. <http://dx.doi.org/10.1002/bdm.684>
- Gluszek, A., & Dovidio, J. F. (2010). The way they speak: A social psychological perspective on the stigma of nonnative accents in communication. *Personality and Social Psychology Review*, 14, 214–237. <http://dx.doi.org/10.1177/1088868309359288>
- Gow, D. J., & Eubank, R. B. (1984). The pro-incumbent bias in the 1982 national election study. *American Journal of Political Science*, 28, 224–230. <http://dx.doi.org/10.2307/2110795>
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23, 121–134.
- Hajnal, A., Bunch, D. A., & Kelty-Stephen, D. G. (2014). Going for distance and going for speed: Effort and optical variables shape information for distance perception from observation to response. *Attention, Perception, & Psychophysics*, 76, 1015–1035. <http://dx.doi.org/10.3758/s13414-014-0629-z>
- Hare, T. A., Malmaud, J., & Rangel, A. (2011). Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *The Journal of Neuroscience*, 31, 11077–11087. <http://dx.doi.org/10.1523/JNEUROSCI.6383-10.2011>
- Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2006). Banishing the homunculus: Making working memory work. *Neuroscience*, 139, 105–118. <http://dx.doi.org/10.1016/j.neuroscience.2005.04.067>
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York, NY: Wiley.
- Higgins, E. T. (2015). What is value? Where does it come from? A psychological perspective. In T. Brosch & D. Sander (Eds.), *Handbook of value* (pp. 43–63). Oxford, UK: Oxford University Press.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group, (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 77–109). Cambridge, MA: MIT Press.
- Hintzman, D. L., & Ludlam, G. (1980). Differential forgetting of prototypes and old instances: Simulation by an exemplar-based classification model. *Memory & Cognition*, 8, 378–382. <http://dx.doi.org/10.3758/BF03198278>
- Hosokawa, T., Kennerley, S. W., Sloan, J., & Wallis, J. D. (2013). Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *The Journal of Neuroscience*, 33, 17385–17397. <http://dx.doi.org/10.1523/JNEUROSCI.2221-13.2013>
- Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. *Nature Reviews Neuroscience*, 18, 172–182. <http://dx.doi.org/10.1038/nrn.2017.7>
- Hunt, L. T., Malalasekera, W. M. N., de Berker, A. O., Miranda, B., Farmer, S. F., Behrens, T. E. J., & Kennerley, S. W. (2018). Triple dissociation of attention and decision computations across prefrontal cortex. *Nature Neuroscience*, 21, 1471–1481. <http://dx.doi.org/10.1038/s41593-018-0239-5>
- Huygens, C. (1657). De ratiociniis in aleæ ludo [The value of all chances in games of fortune]. In F. van Schooten & L. Batava's (Eds.), *Exercitationum mathematicarum*. Leiden, the Netherlands: Elsevier.
- Johnson, E. J., & Goldstein, D. (2003). Medicine. Do defaults save lives? *Science*, 302, 1338–1339. <http://dx.doi.org/10.1126/science.1091721>
- Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: Consensus and controversy. *Neuron*, 63, 733–745. <http://dx.doi.org/10.1016/j.neuron.2009.09.003>
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus, and Giroux.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). The endowment effect, loss aversion, and status quo bias. *The Journal of Economic Perspectives*, 5, 193–206. <http://dx.doi.org/10.1257/jep.5.1.193>
- Kisvárdy, Z. F., Tóth, E., Rausch, M., & Eysel, U. T. (1997). Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cerebral Cortex*, 7, 605–618. <http://dx.doi.org/10.1093/cercor/7.7.605>
- Knutson, B., Taylor, J., Kaufman, M., Peterson, R., & Glover, G. (2005). Distributed neural representation of expected value. *The Journal of Neuroscience*, 25, 4806–4812. <http://dx.doi.org/10.1523/JNEUROSCI.0642-05.2005>
- Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20, 512–534. <http://dx.doi.org/10.1016/j.tics.2016.05.004>
- Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, 119, 573–616. <http://dx.doi.org/10.1037/a0028681>
- Laplace, P.-S. (1814). *Essai philosophique sur les probabilités*. Paris, France: Courcier. Retrieved from <http://books.google.fr/books>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <http://dx.doi.org/10.1038/nature14539>
- Lee, Y., Lee, S., Carello, C., & Turvey, M. T. (2012). An archer's perceived form scales the “hitableness” of archery targets. *Journal of Experimental Psychology: Human Perception and Performance*, 38, 1125–1131. <http://dx.doi.org/10.1037/a0029036>
- Locke, E. A., & Latham, G. P. (2006). New directions in goal-setting theory. *Current Directions in Psychological Science*, 15, 265–268. <http://dx.doi.org/10.1111/j.1467-8721.2006.00449.x>
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA: Freeman.
- Marteau, T. M., Hollands, G. J., & Fletcher, P. C. (2012). Changing human behavior to prevent disease: The importance of targeting automatic processes. *Science*, 337, 1492–1495. <http://dx.doi.org/10.1126/science.1226918>
- McClelland, J. L. (1981). Retrieving general and specific information from stored knowledge of specifics. *Proceedings of the Third Annual Conference of the Cognitive Science Society* (pp. 170–172). Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.126.504>
- McClelland, J. L. (1986). The programmable blackboard model of reading.



- In J. L. McClelland & D. E. Rumelhart, & the PDP Research Group. (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 122–169). Cambridge, MA: MIT Press.
- McClelland, J. L. (2013). Integrating probabilistic models of perception and interactive neural networks: A historical and tutorial review. *Frontiers in Psychology*, 4, 503. <http://dx.doi.org/10.3389/fpsyg.2013.00503>
- McClelland, J. L., & Cleeremans, A. (2009). Connectionist models. In T. Byrne, A. Cleeremans, & P. Wilken (Eds.), *Oxford companion to consciousness* (pp. 177–181). New York, NY: Oxford University Press.
- McClelland, J. L., Mirman, D., Bolger, D. J., & Khaitan, P. (2014). Interactive activation and mutual constraint satisfaction in perception and cognition. *Cognitive Science*, 38, 1139–1189. <http://dx.doi.org/10.1111/cogs.12146>
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part I. An account of basic findings. *Psychological Review*, 88, 375–407. <http://dx.doi.org/10.1037/0033-295X.88.5.375>
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159–188. <http://dx.doi.org/10.1037/0096-3445.114.2.159>
- McClelland, J. L., & Rumelhart, D. E. (1989). *Explorations in parallel distributed processing: A handbook of models, programs, and exercises*. Cambridge, MA: MIT press.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207–238. <http://dx.doi.org/10.1037/0033-295X.85.3.207>
- Neubert, F. X., Mars, R. B., Sallet, J., & Rushworth, M. F. (2015). Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 112, E2695–E2704. <http://dx.doi.org/10.1073/pnas.1410767112>
- Nisbett, R. E., & Wilson, T. D. (1977). The Halo Effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology*, 35, 250–256. <http://dx.doi.org/10.1037/0022-3514.35.4.250>
- O'Doherty, J. P. (2014). The problem with value. *Neuroscience and Biobehavioral Reviews*, 43, 259–268. <http://dx.doi.org/10.1016/j.neubiorev.2014.03.027>
- Opris, I., & Bruce, C. J. (2005). Neural circuitry of judgment and decision mechanisms. *Brain Research Reviews*, 48, 509–526. <http://dx.doi.org/10.1016/j.brainresrev.2004.11.001>
- Orr, D., & Guthrie, C. (2005). Anchoring, information, expertise, and negotiation: New insights from meta-analysis. *Ohio State Journal on Dispute Resolution*, 21, 597–628.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience*, 34, 333–359. <http://dx.doi.org/10.1146/annurev-neuro-061010-113648>
- Perfetti, R. (1993). On the convergence of reciprocal discrete-time cellular neural networks. *IEEE Transactions on Circuits and Systems I, Fundamental Theory and Applications*, 40, 286–287. <http://dx.doi.org/10.1109/81.224306>
- Peterson, M. (2017). *An introduction to decision theory*. New York, NY: Cambridge University Press. <http://dx.doi.org/10.1017/9781316585061>
- Plate, T. (2002). Distributed representations. *Encyclopedia of Cognitive Science*. London: Macmillan.
- Plaut, D. C., & McClelland, J. L. (2010). Locating object knowledge in the brain: Comment on Bowers's (2009) Attempt to revive the grandmother cell hypothesis. *Psychological Review*, 117, 284–288. <http://dx.doi.org/10.1037/a0017101>
- Polanía, R., Krajchich, I., Grueschow, M., & Ruff, C. C. (2014). Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron*, 82, 709–720. <http://dx.doi.org/10.1016/j.neuron.2014.03.014>
- Poosch, S., Bernhardt, N., Guevara, A., Huys, Q. J. M., & Smolka, M. N. (2018). Value-based decision-making battery: A Bayesian adaptive approach to assess impulsive and risky behavior. *Behavior Research Methods*, 50, 236–249. <http://dx.doi.org/10.3758/s13428-017-0866-x>
- Radvansky, G. A. (2017). *Human memory*. New York, NY: Routledge. <http://dx.doi.org/10.4324/9781315542768>
- Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41, e223. <http://dx.doi.org/10.1017/s0140525x18000936>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9, 545–556. <http://dx.doi.org/10.1038/nrn2357>
- Read, S. J., & Miller, L. C. (1998). On the dynamic construction of meaning: An interactive activation and competition model of social perception. In S. J. Read & L. C. Miller (Eds.), *Connectionist models of social reasoning and social behavior* (pp. 27–68). Mahwah, NJ: Lawrence Erlbaum Associates Publishers.
- Read, S. J., Monroe, B. M., Brownstein, A. L., Yang, Y., Chopra, G., & Miller, L. C. (2010). A neural network model of the structure and dynamics of human personality. *Psychological Review*, 117, 61–92. <http://dx.doi.org/10.1037/a0018131>
- Read, S. J., Smith, B., Droutman, V., & Miller, L. C. (2017). Virtual personalities: Using computational modeling to understand within-person variability. *Journal of Research in Personality*, 69, 237–249. <http://dx.doi.org/10.1016/j.jrp.2016.10.005>
- Rich, E. L., & Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. *Nature Neuroscience*, 19, 973–980. <http://dx.doi.org/10.1038/nn.4320>
- Roy, A. (2012). A theory of the brain: Localist representation is used widely in the brain. *Frontiers in Psychology*, 3, 551. <http://dx.doi.org/10.3389/fpsyg.2012.00551>
- Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group. (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 45–76). Cambridge, MA: MIT Press.
- Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1986). *Parallel distributed processing* (Vol. I and II). Cambridge, MA: MIT press.
- Salamone, J. D., Cousins, M. S., & Bucher, S. (1994). Anhedonia or anergia? Effects of haloperidol and nucleus accumbens dopamine depletion on instrumental response selection in a T-maze cost/benefit procedure. *Behavioural Brain Research*, 65, 221–229. [http://dx.doi.org/10.1016/0166-4328\(94\)90108-2](http://dx.doi.org/10.1016/0166-4328(94)90108-2)
- Samuelson, P. A. (1937). A note on measurement of utility. *The Review of Economic Studies*, 4, 155–161. <http://dx.doi.org/10.2307/2967612>
- Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1, 7–59. <http://dx.doi.org/10.1007/BF00055564>
- San-Galli, A., Varazzani, C., Abitbol, R., Pessiglione, M., & Bouret, S. (2018). Primate ventromedial prefrontal cortex neurons continuously encode the willingness to engage in reward-directed behavior. *Cerebral Cortex*, 28, 73–89. <http://dx.doi.org/10.1093/cercor/bhw351>
- Shimojo, S., Simion, C., Shimojo, E., & Scheier, C. (2003). Gaze bias both reflects and influences preference. *Nature Neuroscience*, 6, 1317–1322. <http://dx.doi.org/10.1038/nn1150>
- Simon, D., Snow, C. J., & Read, S. J. (2004). The redux of cognitive consistency theories: Evidence judgments by constraint satisfaction. *Journal of Personality and Social Psychology*, 86, 814–837. <http://dx.doi.org/10.1037/0022-3514.86.6.814>

- Simon, D., Stenstrom, D. M., & Read, S. J. (2015). The coherence effect: Blending cold and hot cognitions. *Journal of Personality and Social Psychology*, 109, 369–394. <http://dx.doi.org/10.1037/pspa0000029>
- Simon, H. A. (1986). Rationality in psychology and economics. *The Journal of Business*, 59, S209–S224. <http://dx.doi.org/10.1086/296363>
- Slovic, P. (1995). The construction of preference. *American Psychologist*, 50, 364–371. <http://dx.doi.org/10.1037/0003-066X.50.5.364>
- Smolensky, P. (1986). Neural and conceptual interpretations of parallel distributed processing models. In J. L. McClelland & D. E. Rumelhart, & the PDP Research Group. (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 390–431). Cambridge, MA: MIT Press.
- Solway, A., & Botvinick, M. M. (2012). Goal-directed decision making as probabilistic inference: A computational framework and potential neural correlates. *Psychological Review*, 119, 120–154. <http://dx.doi.org/10.1037/a0026435>
- Spivey, M. J., & Dale, R. (2006). Continuous dynamics in real-time cognition. *Current Directions in Psychological Science*, 15, 207–211. <http://dx.doi.org/10.1111/j.1467-8721.2006.00437.x>
- Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychological Science*, 26, 122–134. <http://dx.doi.org/10.1177/0956797614559543>
- Suri, G., & Gross, J. J. (2015). The role of attention in motivated behavior. *Journal of Experimental Psychology: General*, 144, 864–872. <http://dx.doi.org/10.1037/xge0000088>
- Suri, G., Sheppes, G., & Gross, J. J. (2015). The role of action readiness in motivated behavior. *Journal of Experimental Psychology: General*, 144, 1105–1113. <http://dx.doi.org/10.1037/xge0000114>
- Suri, G., Sheppes, G., Schwartz, C., & Gross, J. J. (2013). Patient inertia and the status quo bias: When an inferior option is preferred. *Psychological Science*, 24, 1763–1769. <http://dx.doi.org/10.1177/0956797613479976>
- Suri, G., Shine, J. M., & Gross, J. J. (2018). Why do we do what we do? The Attention–Readiness–Motivation framework. *Social and Personality Psychology Compass*, 12, e12382. <http://dx.doi.org/10.1111/spc3.12382>
- Thagard, P. (2008). *Hot thought: Mechanisms and applications of emotional cognition*. Cambridge, MA: Bradford Book, MIT Press.
- Theeuwes, J. (1991). Exogenous and endogenous control of attention: The effect of visual onsets and offsets. *Perception & Psychophysics*, 49, 83–90. <http://dx.doi.org/10.3758/BF03211619>
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. M. (2007). Schemas and memory consolidation. *Science*, 316, 76–82. <http://dx.doi.org/10.1126/science.1135935>
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131. <http://dx.doi.org/10.1126/science.185.4157.1124>
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458. <http://dx.doi.org/10.1126/science.7455683>
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108, 550–592. <http://dx.doi.org/10.1037/0033-295X.108.3.550>
- Vickery, T. J., Chun, M. M., & Lee, D. (2011). Ubiquity and specificity of reinforcement signals throughout the human brain. *Neuron*, 72, 166–177. <http://dx.doi.org/10.1016/j.neuron.2011.08.011>
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Walton, M. E., Kennerley, S. W., Bannerman, D. M., Phillips, P. E. M., & Rushworth, M. F. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, 19, 1302–1314. <http://dx.doi.org/10.1016/j.neunet.2006.03.005>
- Williams, D. R., & Williams, H. (1969). Auto-maintenance in the pigeon: Sustained pecking despite contingent non-reinforcement. *Journal of the Experimental Analysis of Behavior*, 12, 511–520. <http://dx.doi.org/10.1901/jeab.1969.12-511>
- Wilson, T. D., Houston, C. E., Etling, K. M., & Brekke, N. (1996). A new look at anchoring effects: Basic anchoring and its antecedents. *Journal of Experimental Psychology: General*, 125, 387–402. <http://dx.doi.org/10.1037/0096-3445.125.4.387>
- Witt, J. K., & Sugovic, M. (2010). Performance and ease influence perceived speed. *Perception*, 39, 1341–1353. <http://dx.doi.org/10.1068/p6699>
- Witt, J. K., & Sugovic, M. (2012). Does ease to block a ball affect perceived ball speed? Examination of alternative hypotheses. *Journal of Experimental Psychology: Human Perception and Performance*, 38, 1202–1214. <http://dx.doi.org/10.1037/a0026512>

Received July 11, 2018

Revision received July 15, 2019

Accepted July 16, 2019 ■