A Theory of Dynamic Selective Vigilance and Preference Reversal, Based on the Example of New Coke

Daniel S. Levine
Dept. of Mathematics, University of Texas at Arlington, Arlington TX 76019-0408, b344dsl@utarlg.uta.edu

Samuel J. Leven
For a New Social Science, 4681 Leitner Drive West, Coral Springs, FL 33067

Abstract

A neural network theory of preference reversal is presented. This theory includes a model of why New Coke was preferred to Old Coke on taste tests but was unpopular in the market. The model uses competing drive loci representing "excitement" and "security." Context influences which drive wins the competition, hence, which stimulus attributes are attended to. Our network's design, outlined in stages, is based on Grossberg's gated dipole theory. Three sets of dipoles, representing attributes, categories, and drives, are connected by modifiable associative synapses. The network also includes competition among categories and enhancement of attention by mismatch of expectation.

Introduction: Modeling of Irrational Decisions

How rational are we? Tversky and Kahneman (1974, 1981) established that many human decisions do not maximize a measurable utility function. Moreover, deviations from rationality show patterns; for example, decisions among losses are more risk-taking than decisions among gains. Since decision irrationalities are repeatable, they lend themselves to quantitative modeling. Yet models of these effects lag behind models of other cognitive effects, such as pattern classification. Tversky and Kahneman modeled their own data using a non-connectionist theory whereby subjects maximize a nonlinear function of expected gains and losses. However, these authors did not explain how this function arose in the underlying system.

Differential reaction to gains and losses shows that the projected affective value of decisions depends on expectations generated by the current environment. Many neural networks compare current and ongoing values of stimulus or reinforcement variables (Grossberg, 1972; Sutton & Barto, 1981). We use Grossberg's gated dipole theory, to be described below, because it accounts best for stimulus duration effects in conditioning (see Grossberg & Levine, 1987). Grossberg and Gutowski (1987) constructed a gated dipole model for Tversky and Kahneman's data, including the data on gains versus losses. These authors captured the essence of decision under risk despite basing choices on maximizing a single function (affective value). Our model is in the spirit of Grossberg and Gutowski's but adds effects not present in their model: dynamically competing attractions to novel and to familiar stimuli, and competition between drive loci.

Much of Tversky and Kahneman's data involves imagined monetary gains and losses, so their results can be applied to economics. Leven (1987) argues that optimization theory, which dominates economic modeling, is not predictive and must be replaced by theories that include affective factors. Leven and Elsberry (1990) simulate "negotiations" between two neural networks that contain both rational and affective modules. We carry this work further by studying a famous economic example: the failure of New Coke in the market after it had defeated Old Coke in double-blind taste tests. The work of Tversky, Kahneman, and Grossberg readily suggests a qualitative model of the Coke data. Network instantiation of this model, however, led to a complex combination of three sets of gated dipoles representing attributes, categories, and drives; competition among categories and among drives; and associative learning of inter-dipole connections. We first describe the Coke data in detail, then develop our network in stages.

The Coke Data

When the Coca-Cola Company introduced New Coke, it was certain of the flavor's acceptance. Tens of thousands of subjects had undergone highly controlled taste tests. The new flavor had outscored all its competition, including victory over Old Coke by a margin of 2 to 1. Further tests hinted that less than ten percent of Old Coke drinkers would object to the new flavor combined with the old name. As most Americans know, the actual buying situation had very different results. New Coke was so unpopular that the company had to return Old Coke to the market (Oliver, 1986).

Coca-Cola had asked people, "If New Coke were introduced, would you like it?" But the influence of dynamic emotional states means that mental projections of the future are often inaccurate (Holbrook *et al.*, 1985). In the test situation, people based preferences on the direct appeal of taste. In the market, indirect emotional factors, such as memories associated with expected taste, were more important than taste itself. Moreover, buying was different from tests because the Coca-Cola Company was so confident in its research that *Old Coke was unavailable*. The public's reaction against buying New Coke was a *frustrative rebound*. The Coke label created expectation of a particular taste, and of the secure feeling it evoked, which led to frustration when this feeling was absent.

Results of Pierce (1987) support frustration theory. Pierce compared responses to advertisements of old and new versions of Coke by people who had been habitual Coke drinkers and by habitual drinkers of other drinks (such as Pepsi). By a small but significant margin, habitual Coke drinkers were more hostile to products they perceived as New Coke than were non-Coke drinkers.

Frustrative rebound is an example of comparing current with expected or ongoing reinforcement. Just as cessation of a negative reinforcer (*e.g.*, electric shock) is positively reinforcing (provides relief), cessation of a positive reinforcer, or its absence when it is expected, is negatively reinforcing (provides frustration). We will now review how gated dipole networks model both effects.

Background: Gated Dipole Networks

How can a response associated with *offset* of *negative* reinforcement become itself positively reinforcing? To answer this question, Grossberg (1972) introduced the network shown in Fig. 1. The synapses w_1 and w_2 have a chemical transmitter that is depleted with activity. The input J could be shock, for example. The input I is nonspecific arousal to both channels y_1-x_1-x_3 and y_2-x_2-x_4. While shock is on, left channel activity x_1 exceeds right channel activity x_2, leading to net positive activity of the left channel output node x_3. For a short time after shock ceases, both channels receive equal inputs I but the right channel is less depleted of transmitter than the left channel. Hence, x_2 now exceeds x_1, leading to net positive activity of the right channel output node x_4. The active output node excites or inhibits x_5, thus enhancing or suppressing some motor response. The network is called a gated dipole because it has two opposite ("negative" and "positive") channels that "gate" signals based on amounts of transmitter. If the two channels in Fig. 1 are reversed in sign so the channel receiving input is positive, the network explains frustration when positive reinforcement either ceases or is absent when expected.

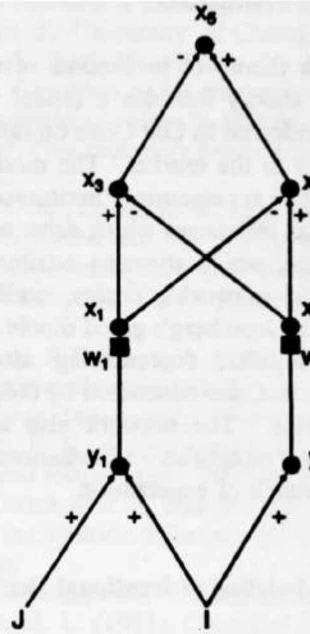

Fig. 1. Schematic gated dipole. "+" denotes excitation, "-" inhibition. Other symbols are explained in text. (From Levine, 1991, with permission of Lawrence Erlbaum Associates.)

Transmitter depletion is not yet verified in many actual synapses; the qualitative effect we model may be based instead on conformation changes at membrane receptors (Changeux, 1981). However, a network principle that models a range of cognitive data can be useful before its biological basis is known, and its later verification is likely even if in a different form than first proposed.

Network Modeling of Coke Data: Combining Sensory and Motivational Dipoles

We build our network for modeling the Coke data in several stages. The network has submodules denoting sensory features, object categories, and drives; simulations are in progress. The rationale for our network architecture can best be shown by starting with simpler networks that model the data partially, then modifying those networks to fit further details of the data.

Gated dipoles instantiate the idea of opponent processing, which applies to vision as well as motivation. For example, there are pairs of opponent colors (*e.g.*, green and red), and each color is transiently perceived after removal of the other. Grossberg (1980) introduced dipoles whose channels consist of "on" and "off" nodes encoding presence or absence of specific stimuli, then joined channel pairs for various stimuli into a *dipole field*. Leven and Levine (1987) discussed how a dipole field could embody competing attractions to previously reinforced stimuli (here, Old Coke) and to novel stimuli (here, New Coke). If drive is high, or reward signals strong, previously reinforced stimuli are favored. If drive is low, novel stimuli are favored. Leven and Levine's first approximation to a model of the Coke data treated testing as a low-motivation state and buying as a high-motivation state, hence explaining preference reversal between the two contexts. They noted an analogy to some monkey data on novelty preference.

Pribram (1961) compared normal rhesus monkeys and those with frontal lobe lesions in a scene with several objects. Successive objects are added to the scene, unobserved by the monkey. Each time a novel object is introduced, a reward (peanut) is placed under the novel object. When the monkey has lifted this object a fixed number of times, the next object is added. Pribram measured the number of errors (liftings of a familiar object) before the monkey first selects the novel object. Frontally lesioned animals are more attracted to novelty than normals so make fewer errors. Fig. 2 shows the dipole field used to model Pribram's data, which was simulated in Levine and Prueitt (1989). The dipole channel pairs in Fig. 2 correspond to an old cue and a novel cue. The nodes $x_{1,5}$ and $x_{2,5}$ represent tendencies to approach given cues. Inhibition between these nodes, and a node $x_{3,5}$ coding some other environmental cue, denote competition between attractions to different cues. The cue with largest $x_{i,5}$ is approached.

The network of Fig. 2 incorporates two competing rules. The on channel corresponding to the novel cue is less depleted than the on channel for the old cue, because that channel has not been active as long. Hence, competition among $x_{i,5}$ nodes favors those corresponding to novel cues, all else equal. But also the reward node is active when the monkey finds the peanut. Each $x_{i,5}$ connects with the reward node via synapses which are strengthened when the corresponding cue is rewarded. Hence, competition also favors $x_{i,5}$'s with strong links to the reward node, all else equal. The frontal lobes are identified with gain from reward nodes to sensory dipoles. If this gain is high, as in a normal monkey, the dipole output $x_{1,5}$ for the previously rewarded cue is the larger. If gain is low, as in a lesioned monkey, the output $x_{2,5}$ for the novel cue is larger.

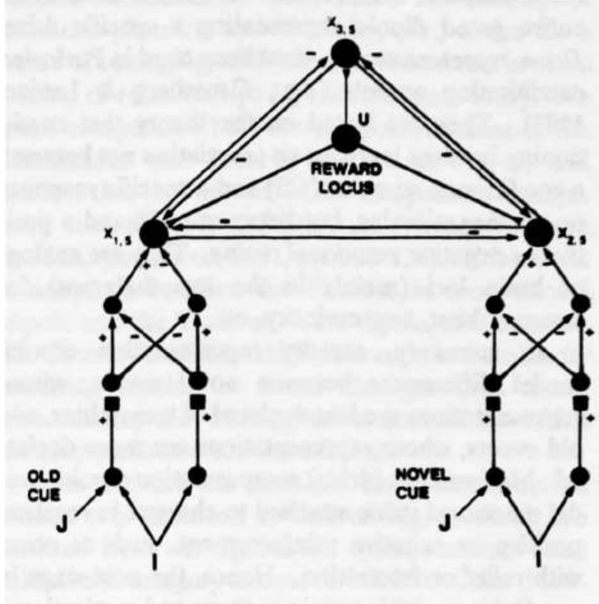

Fig. 2. Dipole field used to simulate novelty data. Semicircles denote modifiable weights. (Adapted from Levine & Prueitt, 1989, with permission of Pergamon Press.)

Leven and Levine (1987) noted that the Coke data could be approximately modeled by the network of Fig. 2, with "New Coke" identified with "novel cue," "Old Coke" with "old cue," "testing" with "frontally damaged," and "buying" with "normal." Of course, most people taking the taste test are not brain-damaged. The analogy is plausible, though, because humans with frontal damage tend to be less goal-directed than normal humans (Fuster, 1989); hence, their day-to-day life is closer to a "play" than to a "serious" situation.

Yet the network of Fig. 2 is inadequate to model the Coke data for at least two reasons. First, in the market, there was not a choice between New Coke and Old Coke as in tests. Hence, relative value attached by buyers to the two drinks must be inferred indirectly from relative preference for New Coke and for non-Coke drinks (*e.g.*, Pepsi). Second, consumers' angry reaction to the change in Coke was not based on taste alone. As one Coca-Cola executive said later, "We were spitting on the American flag and didn't know it." Hence, a realistic model of the Coke data incorporates two competing drives: one for taste, the other for a range of feelings which we label "Security"; these will be added below.

The network of Fig. 2 contains a node that represents a reward signal. We can model frustrative rebound if we replace the reward node by an entire gated dipole representing a specific drive. *Drive representations* have been used in Pavlovian conditioning models (*e.g.*, Grossberg & Levine, 1987). They are based on the theory that conditioning involves learning an association not between a conditioned stimulus (CS) and a specific response or another stimulus, but between a CS and a positive or negative *emotional value*. They are analogs of brain loci (mainly in the hypothalamus) for hunger, thirst, sex, curiosity, *etc*.

In summary, sensory representation dipoles model differences between novel events, whose representations are less depleted of transmitter, and old events, whose representations are more depleted. Motivational (drive) representation dipoles model emotional value attached to changes in received positive or negative reinforcement, such as occur with relief or frustration. Hence, the next stage in our Coke model is to join sensory and motivational dipoles (Fig. 3).

Combined sensory and motivational dipoles may also account for conditioning phenomena such as *unblocking* (Kamin, 1969). *Blocking* has been simulated in neural networks (Grossberg & Levine, 1987; Sutton & Barto, 1981). If a bell, say, is paired with shock and an animal learns a fear response to the bell, then a bell-light combination is paired with the same shock, no fear is learned to the light. The light is *unblocked* when the shock level paired with the bell-light compound is unequal to the level paired with bell alone. Since unblocking involves associating a novel stimulus (light) and a changed affective value, it may be modeled by the network of Fig. 3.

Network Modeling of Coke Data: Categorizations and Multiple Attributes

New Coke elicited strong reactions because of how it was *different* from Old Coke, but also because of how it was *like* Old Coke! The public reacted to a new taste *combined with an old label*. Pierce's (1987) data show that the reaction was stronger when larger positive affect was attached to the old taste. Hence, our modeling problem becomes how to build a network to respond to stimuli that contain both a novel element and a familiar, significant element. For example, Robert Dawes (personal communication) has discussed a network model of seeing the Mona Lisa with a mustache added. We find that picture grotesque because the mustache mismatches expectations produced by the rest of the picture.

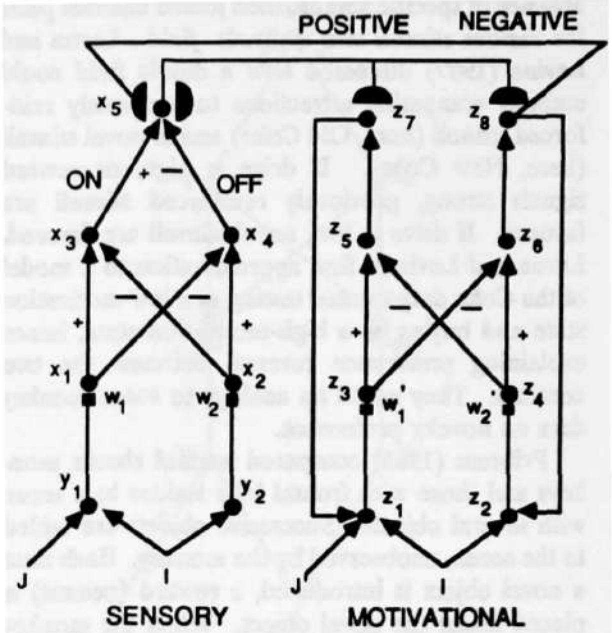

Fig. 3. Combination of sensory and motivational dipoles. Dipole output for a sensory stimulus can be conditioned to positive or negative reinforcement, as shown by modifiable connections at top.

To deal with expectation, we refine our model so that New and Old Coke are no longer single stimuli but vectors of attributes, *each attribute represented by its own gated dipole* (Fig. 4). We use the minimal set of attributes needed: Coke Label; Familiarity; Taste; Pepsi Label. The latter is introduced to model the switch from New Coke to competing cola drinks (lumped together as "Pepsi" for simplicity), or else to avoidance of all soft drinks, when Old Coke was unavailable.

How does this explain the data of Pierce (1987) on habitual Coke drinkers versus habitual Pepsi drinkers? If a network is to represent general cognitive principles, it should also, if possible, account for individual differences. Major behavioral differences can arise from differences in one or a few network parameter values. In Fig. 4, let weights (in both directions) between the on side of the Coke Label attribute dipole and the positive side of the motivational dipole be higher in one copy of the network than in another. Then the first network models a (generic) habitual Coke drinkers, whereas the second models a habitual Pepsi drinker. Analogously, let corresponding weights to and from the Pepsi Label attribute dipole be higher in the second network. Because of feedback from drive to sensory loci in the network of Fig. 4, the expected positive affective value from seeing the Coke label is greater in the habitual Coke drinker. Hence, frustrative rebound from mismatching expectations generated by that label is also greater in habitual Coke drinkers.

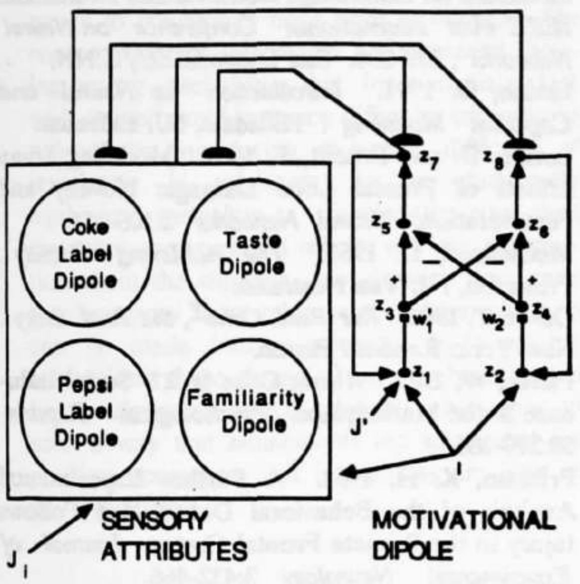

Fig. 4. Extension of network of Fig. 3 to include dipoles for each stimulus attribute, with inputs J_i to each. Circles represent sensory dipoles, not shown in full here. Each sensory dipole has modifiable reciprocal connections with positive and negative motivational channel outputs, z_7 and z_8.

At this point, we must revise our account of the difference between testing and buying. To a first approximation, we have treated testing as a "low motivation" context which thereby disinhibits the attraction to novelty. Yet when we look at attributes, what is actually different between the two contexts is not the *amount* of motivation but rather the *focus* of motivation. During testing, the *intrinsic* (taste-related) attractiveness of the product is important, and the *socially learned* attractiveness of the product much less so. Hence, the Taste attribute (Fig. 4) plays a larger role in categorizations and decisions during testing than does the Familiarity attribute. The Familiarity attribute, by contrast, plays a larger role during buying.

Now we need a theory of context-based attentional switches between attributes. Such a theory includes multiple sensory dipoles *and* multiple motivational dipoles (Fig. 5). Here, two dipoles are labeled "Excitement," *i.e.*, desire for sensory or aesthetic pleasure, and "Security," *i.e.*, desire for a sense of belonging, affiliation, or rootedness in one's society or relationships (*cf*. McClelland, 1961). We posit competition between the positive sides of the Excitement and Security dipoles in Fig. 5, and assume that the "winner" of the competition changes with context (For a history of relevant network models, see Levine, 1991, pp. 133-134). If feedback connections between the Excitement motivational dipole and the Taste attribute dipole, and between the Security motivational dipole and the Familiarity attribute dipole, are much stronger than cross-connections, the winning drive determines which sensory attributes are attended to.

The network we simulate for a future article makes two additions to that of Fig. 5, which are omitted from our figures for space reasons. One addition is category nodes. If habitual Coke drinkers attach positive affect to the *Coke category* as well as the Coke Label attribute, this enhances expectation of positive value from drinking any Coke product, thus increasing frustration when New Coke mismatches that expectation. Our current network includes modifiable feedback between category nodes and attribute nodes, in the manner of ART networks (Carpenter & Grossberg, 1987). Category nodes also connect directly with motivational dipoles. Hence, the affective value of a category can differ from values of the category's exemplars (*e.g.*, one can love humanity and hate people, or vice versa). In ART, the input vector is compared with stored category prototypes, and classified with any prototype that it mismatches to less than a prescribed amount (*vigilance*). Our model posits that vigilance is dynamically feature-selective. If, for example, the current attentional bias favors the Familiarity attribute over the Taste attribute, the network is selectively sensitive to mismatch with the Coke prototype in the Familiarity dimension. The amygdala might be a brain locus for such a bias mechanism (Pribram, 1991).

Our second addition is modulation by mismatch signals of perceived time durations of inputs, which was introduced by Ricart (1992). Ricart identified the modulatory node with the midbrain locus ceruleus, which produces norepinephrine and focuses attention on significant or novel stimuli. This node nonspecifically sharpens perception of both stimuli and reinforcements after mismatch generated by the attribute-category subsystem.

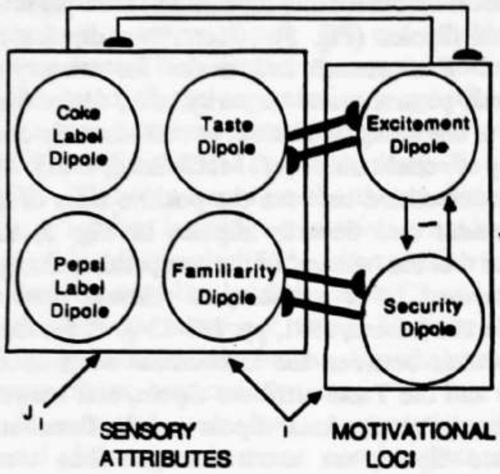

Fig. 5. Extension of network of Fig. 4 to include competing motivational dipoles (shown as circles). Darker lines indicate stronger connections.

References

Carpenter, G., and Grossberg, S. 1987. A Massively Parallel Architecture For a Self-organizing Neural Pattern Recognition Machine. *Computer Vision, Graphics, and Image Processing* 37: 54-115.

Changeux, J.-P. 1981. The Acetylcholine Receptor: an "Allosteric" Membrane Protein. *Harvey Lectures* 75:85-254.

Fuster, J. 1989. *The Prefrontal Cortex*. New York: Raven.

Grossberg, S. 1972. A Neural Theory of Punishment and Avoidance. II. Quantitative Theory. *Mathematical Biosciences* 15, 253-285.

Grossberg, S. 1980. How Does a Brain Build a Cognitive Code? *Psychological Review* 87:1-51.

Grossberg, S. and Gutowski, W. 1987. Neural Dynamics of Decision Making Under Risk: Affective Balance and Cognitive-Emotional Interactions. *Psychological Review* 94:300-318.

Grossberg, S. and Levine, D. 1987. Neural Dynamics of Attentionally Modulated Pavlovian Conditioning: Blocking, Interstimulus Interval, and Secondary Reinforcement. *Applied Optics* 26:5015-5030.

Holbrook, M.; Moore, W.; Dodgen, G.; & Havlena, W. 1985. Nonisomorphism, Shadow Features, and Imputed Preferences. *Marketing Science* 4:215--233.

Kamin, L. 1969. Predictability, Surprise, Attention, and Conditioning. In B. Campbell & R. Church (Eds.), *Punishment and Aversive Behavior*, 279-296. New York: Appleton-Century-Crofts.

Leven, S. 1987. Choice and Neural Process. Ph. D. diss., School of Urban Studies, University of Texas at Arlington.

Leven, S. and Elsberry, W. 1990. Interactions Among Embedded Networks Under Uncertainty. In International Joint Conference on Neural Networks, San Diego, Vol. III, 739-746. Ann Arbor, MI: IEEE Neural Network Council.

Leven, S. and Levine, D. 1987. Effects of Reinforcement on Knowledge Retrieval and Evaluation. *IEEE First International Conference on Neural Networks*, 279-296. San Diego:IEEE/ICNN.

Levine, D. 1991. *Introduction to Neural and Cognitive Modeling*. Hillsdale, NJ: Erlbaum.

Levine, D. and Prueitt, P. 1989. Modeling Some Effects of Frontal Lobe Damage: Novelty and Perseveration. *Neural Networks* 2:103-116.

McClelland, D. 1961. *The Achieving Society*. Princeton, NJ: Van Nostrand.

Oliver, T. 1986. *The Real Coke®, the Real Story*. New York: Random House.

Pierce, W. 1987. Whose Coke Is It? Social Influence in the Marketplace. *Psychological Reports* 60:279-286.

Pribram, K. H. 1961. A Further Experimental Analysis of the Behavioral Deficit that Follows Injury to the Primate Frontal Cortex. *Journal of Experimental Neurology* 3:432-466.

Pribram, K. H. 1991. *Brain and Perception*. Hillsdale, NJ: Erlbaum.

Ricart, R. 1992. Neuromodulatory Mechanisms in Neural Networks and Their Influence on Interstimulus Effects in Pavlovian Conditioning. In D. Levine & S. Leven, eds., *Motivation, Emotion, and Goal Direction in Neural Networks*, 117-166. Hillsdale, NJ: Erlbaum.

Sutton, R. S. and Barto, A. G. 1981. Toward a Modern Theory of Adaptive Networks: Expectation and Prediction. *Psychological Review* 88:135-170.

Tversky, A. and Kahneman, D. 1974. Judgment Under Uncertainty: Heuristics and Biases. *Science* 185:1124-1131.

Tversky, A. and Kahneman, D. 1981. The Framing of Decisions and the Rationality of Choice. *Science* 211:453-458.