Chapter to appear in 2012: Cognitive Consistency: A Fundamental Principle in Social Cognition. Edited by Bertram Gawronski and Fritz Strack.

5

Fluency of Consistency

When Thoughts Fit Nicely and Flow Smoothly

Piotr Winkielman
David E. Huber
Liam Kavanagh
Norbert Schwarz

Cognition influences affect—sometimes dramatically, as when a crafty political speech rouses passions and moves people to lofty or dangerous endeavors. However, many cognition–affect interactions are subtle. A fuzzy text or a noisy reception triggers annoyance. A harmonious chord or a symmetric design elicits pleasure. Recognizing a familiar face produces a sense of warm glow. There is something "off" in our colleague's explanation of an event. Our chapter deals with these subtle interactions. Specifically, we address how *everyday evaluative responses* depend on fluency—ease or difficulty of information processing. Furthermore, we show how the mechanisms linking fluency and evaluation can be understood using a combination of psychological experimentation and computer modeling.

The unique aspect of our contribution, besides offering several theoretical refinements and an empirical update on our respective programs of research, is that we address how the notion of fluency relates to the notion of consistency and related concepts, such as coherence and dissonance; that is, we explore how fluency and consistency phenomena are similar and different from each other, both conceptually and empirically. We also show how they can be precisely modeled using a computational framework.

The structure of our contribution is roughly as follows. First, we distinguish various sources of evaluative responses—nonspecific processing dynamics and specific feature-based information. This then allows us to address the similarities and differences between the concepts of fluency and consistency. Next, we describe empirical work suggesting that evaluative reactions to fluency can explain several common preference phenomena. Finally, we describe some computational models of fluency and consistency.

FLUENCY, CONSISTENCY, AND RELATED CONCEPTS

Fluency and Features

A visitor arrives in Warsaw and sees several faces in the airport's arrival hall. Some she identifies as Polish, others she does not. Some she recognizes, others she does not. Some she likes, others she does not. Why? One obvious source of her reactions is the "what" of processing—the stimulus's specific *features*. For example, detecting a round face with high cheekbones suggests that the person is Polish. Detecting a particular mustache on a waving greeter triggers recognition of the visitor's host. Detecting a smile—"the curve that sets everything straight"—triggers positive affect. In addition, however, there is the "how" of processing—the dynamics of perceptual and conceptual operations on the stimulus. These dynamic parameters include processing speed and processing ease—how much effort the process requires (subjectively and objectively). For example, how fast and effortful was the process of identifying, recognizing, and categorizing the face? These attributes are usually referred to with the general term *fluency*. Fluency is often divided into *perceptual* fluency, which relates to basic stimulus identification processes, and *conceptual* fluency, which relates to higher-order interpretational and reasoning processes (Jacoby, Kelley, & Dywan, 1989).

As we describe shortly, one idea guiding this chapter is that fluency is typically associated with positive evaluations. Thus, returning to our initial example, some faces might be easier to recognize, and this ease may trigger positive affect. We also discuss how fluency links to other experiences and judgments, from familiarity to truth.

It is useful to highlight a few things about the relation between features and fluency. First, both can be available simultaneously, with each contributing to the final reaction. For example, positivity from detecting a "smile" feature can combine with positivity from ease of face recognition. Second, features and fluency can play off each other. For example, the same feature, such as symmetry, might not only create a positive reaction because of its cognitive implications (e.g., health) but also make the face easier to recognize. In other words, a feature might create an evaluative reaction not only directly but also indirectly, via its influence on fluency.

Fluency and Consistency

The relation between the concepts of *fluency* and *consistency* is in some way straightforward. Fluency is about "how"—the speed and ease with which a particular cognitive element, or a set of elements, is processed. Consistency is about "what"—a content *match* between cognitive elements. As such, consistency is one of the relations that may result in fluency. However, what exactly is a "consistent match?" This issue also touches on some of the trickiest problems in psychology. So before we discuss this, let's briefly recall how *consistency* is generally used in psychology.

What Is Consistency?

The idea of "consistency" goes back to the days of gestalt psychology, which emphasized that perception and cognition tend to organize in a way that achieves *prägnanz* (a concise, pithy, "good" structure). Heider (1946) proposed that mental organization of social relations uses similar principles. His balance theory mostly deals with evaluative consistency, in which relationships between elements (e.g., people) are balanced when their attitudes (likes, dislikes) toward each other *match* in a relevant respect. For example, Norbert likes Liam, Liam likes Dave, Dave likes Norbert, and Norbert and Dave dislike Piotr. A *mismatch* in the structure (e.g., Liam likes Piotr) is posited to cause some level of psychological discomfort, or at least a motivation to restore balance. Balance theory was followed by Festinger's (1957) famous dissonance theory and various congruity theories, but the basic logic stayed the same (Abelson, 1983). Interestingly, consistency theories were very popular in the 1960s, but they declined in popularity in the 1970s and 1980s, being replaced by more content-oriented cognitive approaches (Abelson, 1983). Reasons for the decline include difficulties in defining what constitutes a relevant relation, what exactly constitutes inconsistency/mismatch, and doubts whether people are motivated to resolve inconsistency. Perhaps the most important reason was the growing interest in the "what" of cognition—the actual content, rather than just formal relations between cognitive elements. These problems remain. Interestingly, however, recently there has been a revival of consistency concepts in psychology in general, and social psychology in particular (e.g., this very book). One example of this revival is the development of constraint satisfaction models (e.g., neural networks), which have become a major tool for understanding the operation of cognition. These models are now widely used for the modeling of social phenomena in dynamical psychology (Nowak & Vallacher, 1998; Read & Simon, Chapter 4, this volume). Similarly, principles of consistency, mutual adjustment, constraint satisfaction, and synchronization are used in the modeling of social network effects (e.g., contagion effects on obesity, loneliness; Lazer et al., 2009). With all this in mind, here are some essential points about the relation of consistency to fluency.

Consistency with Fluency: Regularity

Note that fluency is basically a description of a cognitive and behavioral effect—processing is easier, smoother, faster, more accurate, and so forth. Nothing more, nothing less. To the extent that fluency is an effect, one needs to ask whether consistency creates fluency effects. To this the answer is certainly "yes." After all, consistency is about elements being tied together in a relation. Thus, via the basic mechanism of *priming, in which a presentation of one element can preactivate a related element*, consistency breeds coactivation, and thus fluency. Another way to put it is that consistency is about regularity and, as we discuss shortly, priming is essentially about exploiting regularities.

Some of these regularities are purely structural, others are associative, semantic, based on predictability, and still others are based on high-order relations. For example, priming that produces perceptual fluency can involve a simple structural "identity" relation: Presenting a picture of a polygon facilitates processing of a subsequent polygon, and presenting the word *chair* facilitates processing of subsequent word *chair*. A semantic structure underlies a similarity relation between *chair* followed by *furniture*. An experiential structure (co-occurrence in language) underlies an associative relation *chair* followed by *person*, and so forth. A predictability relation underlies the fluency of "In a waiting room, a helpful man offered an elderly woman a *chair*." We complicate this picture later when discussing phenomena such as overpriming and semantic saturation. But for now we propose that priming—one basis of fluency—is based on basic forms of consistency.

If consistency breeds fluency, the two of them should be confusable to some extent. This leads to an interesting prediction that people will have difficulty noticing conceptually inconsistent information if it is presented smoothly. Song and Schwarz (2008) found just this in their research on the so-called Moses illusion. When people are asked how many animals of each kind Moses took on the ark, they often say "Two." This is an error, since Moses comes from the biblical story of tablets, not the flood. Interestingly, participants are less likely to detect the error when the question is written in an easy-to-read font.

Consistency without Fluency: Connection and Coherence

There are, however, consistency relations that fail to produce fluency, and other consistency relations that produce fluency but not via simple priming-like mechanisms. Several issues are worth mentioning here.

When two elements are said to be "consistent," it does not mean that they are mechanistically connected in a person's mind. After all, consistency is about having a relation in some "abstract" way. Thus, the word *physician* may be semantically connected to *syringe*, but these words may nevertheless fail to prime each other (Nelson, McEvoy, & Schreiber, 1998), or they will prime each other in one context, such as *hospital*, but not another, such as *heroin* (Kellas, Paul, Martin, & Simpson, 1991). As a consequence, it is difficult for consistency theories to use the abstract notion of a "match" to a priori predict the propositions that will be subjectively represented as consistent and mechanistically work as such. This is true even when one specifies consistency as *psychological consistency* and defines it in an idiosyncratic, individual-specific way. Anecdotally, in everyday life, there are often situations in which people fail to "put two and two together"—connect two obvious cognitions—either because they do not think about them at the same time, or because they "fail to see the connection." In short, specifying consistency on the abstract level must always be accompanied by some mechanistically realistic theory of how the corresponding cognitions are realized with a capacity-constrained and motivation-limited individual mind.

A perhaps more serious issue is that consistency may refer to complex semantic relations. This is easiest to explain with a related concept of *coherence*, which is a type of higher-order consistency, including relations such as logical deduction, compatibility in explanation, causality, and so forth (Thagard, 2000). Note, for example, that logical coherence does not imply fluency, as it is not directly based on any empirically encountered regularity. This can be illustrated using syllogistic reasoning. Take the premises: "All fish live in the water" and "Salmon live in the water." Now evaluate whether this conclusion logically follows from the premises: "Therefore, salmon are fish." This conclusion is incoherent (false), because it does not conform to our logical schema of proper reasoning. However, it is not disfluent. After all, processing the premises does not necessarily make processing the sentence "Salmon are fish" slower or more difficult. In fact, it has been empirically demonstrated that people sometimes find it easier to process logically incoherent syllogisms that have surface similarity compared to logically true syllogisms that lack surface similarity (Evans, Barston, & Pollard, 1983). In short, *logical coherence* refers to a particular type of rule-based relation that can cause behavioral fluency in some circumstance (and certainly does over time), although fluency is not a necessary consequence of logical coherence.

Such phenomena are not restricted to complicated propositional logic and can occur with simpler stimuli. Reading 6×6 makes the number 36 fluent (due to associative priming). But reading $323 - 322 = 232$ involves priming, yet it is rendered incoherent by violating the simple rules of subtraction. To get at these issues, some studies placed rule-based coherence and priming in direct opposition. For example, the words *smoke* and *fire* are associatively related. This associative relation, which leads to simple priming, makes it difficult to notice the falsehood of the statement "Smoke causes fire" (Fenker, Waldmann, & Holyoak, 2005). This distinction between rule-based processes and associative processes is further explored in various "two-system" models (e.g., Strack & Deutsch, 2004). Of course, the issue of distinguishing between associative and "rule-based" processes is complex (Kruglanski & Gigerenzer, 2011). However, for our purposes here, it is only important to highlight that fluency (the "how") based on association can often run counter to coherence based on content (the "what").

In these examples, people occasionally make errors when faced with a structurally consistent or fluent but logically or factually incoherent situation (Song & Schwarz, 2008; Evans et al., 1983; Fenker et al., 2005). However, most of the time people get it right. Therefore, should researchers perhaps focus on coherence of content rather the more structural aspects of consistency? Not at all. We need both—structural consistency and inferential coherence. In fact, the idea that structural consistency relations are meaning-dependent is not a novel point. The balance theory literature has long noted that a balanced relation in a friendship context (Bill likes Hilary, Bill likes Monica, Hilary likes Monica) becomes unbalanced in a sexual relationship context. Thus, the mechanistic balance and fluency processes operate within a network of relations specified over a meaningful content. In some way, then, modern versions of consistency and coherence theories are extensions of this insight. We elaborate on these points when we discuss our modeling work using constraint satisfaction models, and when we go back to some more concrete examples of coherence processes in higher-order reasoning. But for now, let us return to fluency.

HOW IS IT GOING?: LINKING FLUENCY, CONSISTENCY, AND AFFECT

Fluency and consistency, or the lack thereof, have affective and motivational consequences. It is worth highlighting a few reasons why this should be so. The major explanations focus on the role of fluency and consistency as cues to the quality of the internal state of the organism and the quality of an external stimulus.

Information about Quality of Internal Processing

Affect has many functions, but one is to inform the organism about its own internal state (Carver & Scheier, 1990). When things are in physiological homeostasis we feel good, and when they are imbalanced we feel bad. Similarly, affect informs us about the current state of cognitive operations. Thus, fluent and consistent processing "feels good" because it indicates an internally coherent structure of beliefs. For example, perceptual fluency indicates good progress toward recognition of the stimulus, whereas conceptual fluency indicates progress toward a successful task solution.

Motivational Force

Besides informing the organism that processing is going well, positive affect triggered by fluency and consistency may play a motivational function and reinforce successful strategies. On the other hand, disfluency and inconsistency can serve as signals of cognitive error or incompatibility, and motivate a strategy revision (Fernandez-Duque, Baird, & Posner, 2000; Nowak & Vallacher, 1998).

Affective Implications of Incoherence

These theoretical ideas converge with classic observations from the social psychological literature that mental states characterized by low coherence can often be unpleasant. For example, several studies have documented the unpleasantness of cognitive dissonance with self-report, as well as physiological measures of affect (Elliot & Devine, 1994; for review, see Harmon-Jones, Amodio, & Harmon-Jones, 2009).

It is worth noting that some social psychological explanations for why inconsistency is unpleasant emphasize high-level social implications of holding incompatible beliefs or performing belief-inconsistent actions (e.g., writing a counterattitudinal essay). For example, some modern social psychological accounts of cognitive dissonance propose that the unpleasantness of inconsistency comes not from the mere dissonance itself, but from the implication of dissonance as a threat to self-image and self-esteem. Accordingly, self-esteem and self-affirmation manipulations help to reduce the consequences of dissonance. Note, however, that all these effects, and interventions (which are largely interchangeable), may be understood as essentially targeting negative affect. Thus, they are consistent with the idea that at the core of dissonance lies a basic affective mechanism (Tesser, 2000).

This idea is consistent with another current social psychological theory that emphasizes the finding that dissonance effects are particularly strong when they involve a mismatch between beliefs and a possible or actual action (Harmon-Jones et al., 2009). On this account, the reason why belief–action mismatch is "unpleasant" and motivating has to do with the fact that it puts the organism into a behavioral conflict (e.g., approach–avoidance conflict), which generates negative affect via basic conflict detection processes.

Finally, the idea that dissonance elicitation mechanisms are basic explains why dissonance can occur when high-level social concerns are eliminated or severely reduced. Thus, dissonance effects have been observed in completely private settings and are found in children, patients with amnesia, and even monkeys (Egan, Santos, & Bloom, 2007). As such, it seems doubtful that high-level explanations can supersede low-level explanations that focus on basic psychological mechanisms of inconsistency processing.

Information about Quality of External Stimuli

Fluency and consistency do more than just inform the perceiver about his or her own internal state. Smooth processing can also have affective consequences, because it indicates (probabilistically) whether an external stimulus is good or bad. One version of this logic is well-known from research on heuristics and biases. Fluency, or its consequences, such as familiarity and recognition, can be a learned, "fast and frugal" heuristic for identifying choices that are likely to be objectively better (Gigerenzer, 2007). Similarly, as we discuss later, fluency offers a cue that indicates the likelihood of other valued properties of external stimuli, such as symmetry, prototypicality, and so forth.

The heuristics and biases literature assumes that the fluency–judgment link is used strategically, and is subject to naive-theories and short-term environmental experiences. Later in this chapter we illustrate this point with some empirical examples. But such processes are also automatized, or even partly innate. For example, it is known, at least since Titchener (1910), that familiar (fluent) stimuli elicit a "warm glow"—an effect that can be especially powerful in infants. Accordingly, illusions of familiarity (oldness) can be produced through unobtrusive inductions of positive affect (Phaf & Roteveel, 2005). One reason for this warmth–familiarity link could be a biological predisposition for caution in encounters with novel, and thus potentially harmful, stimuli (Zajonc, 1998).

Finally, fluency can enter into complex inferences about the external information that can eventually result in affective and motivational consequences. One interesting line of research that is relevant to the question about fluency and consistency explored the effects of fluency on the perception of social truth. As Festinger (1954) observed, we often evaluate truth on the basis of social consensus, and infer that there "must be something to it" if many others believe it. Unfortunately, people are bad at tracking how often they heard what from whom, and they misread the familiarity resulting from repeated exposure as evidence for broad consensus, even if the familiarity merely results from the same, single person saying the same thing over and over again (Weaver, Garcia, Schwarz, & Miller, 2007). Experimentally, any variable that increases the fluency with which a statement can be processed also increases the likelihood that the statement is accepted as true—and merely printing the statement in higher figure–ground contrast (Reber & Schwarz, 1999) or a rhyming form (McGlone & Tofighbakhsh, 2000) is sufficient to enhance its truth value (for a review, see Schwarz, Sanna, Skurnik, & Yoon, 2007). In short, information that is presented in a fluent form comes with a higher perceived truth value; it should therefore result in more dissonance when it conflicts with one's own beliefs. This possibility awaits empirical testing but is nicely consistent with findings that dissonance depends on the subjective truth value that is assigned to related propositions (Gawronski & Strack, 2004).

EMPIRICAL EVIDENCE FOR THE ROLE OF FLUENCY IN EVALUATION

Next we turn to a brief review of the literature establishing the fluency–affect link (for a more comprehensive review, see Winkielman, Schwarz, Fazendeiro, & Reber, 2003; Reber, Schwarz, & Winkielman, 2004). A large number of studies generally support the idea that processing fluency enhances evaluation. Here we discuss a subset of this literature focusing on five related variables: (1) repetition/mere exposure; (2) priming; (3) contrast, clarity, and duration; (4) symmetry; and (5) prototypicality. One reason we focus on these five variables is because we want to draw a connection between the empirical work and computer modeling approaches, as we discuss shortly. Importantly, many of these variables have also been found to cause other judgmental consequences of fluency (e.g., familiarity, truth, sense of temporal distance; see Schwarz, 2010).

Repetition (Mere Exposure)

The simplest way to facilitate stimulus processing is repetition. Accordingly, repetition should enhance liking (Bornstein & D'Agostino, 1994; Jacoby et al., 1989; Mandler, Nakamura, & Van Zandt, 1987). This is, of course, the famous *mere exposure effect* (MEE)—the observation that simple repetition enhances liking for an initially neutral stimulus (Zajonc, 1968). The MEE has been demonstrated behaviorally, using liking and mood judgments (e.g., Monahan, Murphy, & Zajonc, 2000). It has also been demonstrated physiologically, using facial electromyography (EMG), which relies on the observation that positive affect triggers incipient smiles (EMG activity over the cheek region), whereas negative affective trigger incipient frowns (EMG activity over the brow region). Harmon-Jones and Allen (2001) observed that repeatedly presented stimuli elicited stronger EMG activity over the "smiling," without changing the activity over the "frowning" region.

Priming

Priming also facilitates processing. Thus, it should result in increased liking, even under conditions of a single exposure. Many studies have now confirmed this. In an early investigation, participants identified and evaluated pictures of everyday objects (e.g., desk, bird, or plane). Picture processing was facilitated or inhibited by prior subliminal presentation of matching or mismatching visual contours. Results showed that pictures preceded by matched contours were identified faster, indicating higher fluency, and were liked more than pictures preceded by mismatched contours (Reber, Winkielman & Schwarz, 1998, Study 1).

Critically, there is physiological evidence for the positivity of reactions caused by priming, gathered using the earlier-mentioned EMG technique (Winkielman & Cacioppo, 2001). High fluency triggers stronger activity over the "smiling" region but not the "frowning" region. This effect occurs quickly after stimulus presentation, suggesting an automatic link between fluency and positivity.

Contrast, Clarity, and Duration

High contrast, clarity, and duration facilitate identification. Thus, these features should trigger liking. Indeed, participants like the same stimulus more when it is presented with higher contrast and longer presentation (Reber et al., 1998). Again, these manipulations trigger EMG activity specific to the cheek region, suggesting that high fluency elicits positive affect on the physiological level (Winkielman & Cacioppo, 2001).

Symmetry

Humans and nonhuman animals show a widespread preference for symmetry (Rhodes, 2006). This is often attributed to the biological value of symmetry as a signal of mate quality (e.g., Thornhill & Gangstead, 1993). However, symmetry could be appealing partly because it is fluent. After all, symmetrical stimuli are structurally simpler than nonsymmetrical stimuli. This simplicity, and redundancy of information, facilitates identification of stimuli (Palmer, 1990). Support for this comes from studies on preference and fluency of abstract shapes (Reber & Schwarz, 2006). Participants made preference judgments and also same–different judgments for symmetrical and asymmetrical shapes. The results showed that symmetrical shapes are not only more appealing but also easier to identify than comparable asymmetrical shapes.

Prototypicality

People prefer prototypicality or "averageness"—in the sense of a stimulus fitting the central tendency of a category (Rhodes, 2006). This applies to living objects, such as faces, fish, dogs and birds, and also to nonliving objects, such as color patches, furniture, wristwatches and automobiles (Halberstadt, 2006). This preference has often been explained as an evolved predisposition to interpret prototypicality as a cue to mate quality (Symons, 1979). However, there is a more straightforward fluency explanation. Given that prototypes are the most representative members of their categories, they are also fluent, as reflected in accuracy and speed of classification (Posner & Keele, 1968). This raises the possibility that prototypes are liked *because* they are fluent. Winkielman, Halberstadt, Fazendeiro, and Catty (2006) tested this idea. In some studies, participants first learned a category of visual patterns. Then they saw novel patterns varying across different levels of prototypicality. Participants classified these patterns (measure of fluency) and also rated their attractiveness. A close relationship among fluency, attractiveness, and the level of prototypicality was observed. Both fluency and attractiveness increased with prototypicality. Importantly, when fluency was statistically controlled, the relation between prototypicality and attractiveness dropped by half (though it remained significant). This suggests that processing facilitation is an important contributor to the "beauty-in-averageness" effect. Interestingly, prototypical rather than nonprototypical patterns elicited significantly greater EMG activity, suggesting that categorically fluent stimuli trigger genuine affective reactions.

CONTEXTUAL MODERATION: THE ROLE OF MOOD AND NAIVE THEORIES IN FLUENCY–AFFECT LINK

Is the reaction to fluent stimuli always positive? Based on the logic presented earlier, this link appears fixed. However, note that besides affect, fluency can also lead to other subjective experiences, such as familiarity. Furthermore, fluency itself may be subject to different interpretations, depending on the judgmental context. Both processes may have affective and judgmental consequences that occasionally flip the fluency–positivity link.

Fluency, Familiarity, and Mood

One prominent fluency-dependent experience is familiarity (Jacoby et al., 1989; Whittlesea, 1993). As we discuss shortly, fluency manipulation robustly change judgments of stimulus "oldness." But the connection between familiarity and positivity is tricky. After all, positivity of the "known and old" depends on one's current environment. Familiarity is nice in a hazardous environment but less so in a friendly one. For example, in a strange city, a familiar face or food may elicit a warm glow, whereas locally the same items prompt a yawn. Similarly, children cling to the familiar in an unsafe, shaky environment but explore in a safe, comfortable one.

So, what tells us whether the environment is dangerous? One cue is an individual's mood. Bad mood signals a problem, whereas good mood suggests a benign environment (Schwarz & Clore, 1996). Thus, the value of familiarity should depend on how people feel. In a recent study, we (De Vries, Holland, Chenier, Starr, & Winkielman, 2010) examined this reasoning by presenting participants with stimuli that were made categorically fluent and familiar via the prototypicality manipulation described earlier (Winkielman et al., 2006). Then, we measured participants' emotional and memory responses with stimulus ratings and physiological measures (facial EMG and skin conductance). Critically, we put some participants in a good mood and others in a bad mood. As predicted, sad participants showed the classic preference for familiar stimuli, including the "beauty-in-averages" effect and the MEE. These reactions were genuinely positive, as reflected by EMG measures of "smiling." However, happy mood eliminated these preferences. Even the usually highly liked familiar prototypes were rated low and did not produce smiles. This occurred, interestingly, in spite the fact that in happiness, prototypes rated as particularly familiar triggered stronger skin conductance responses, consistent with earlier research on familiarity and arousal (Morris, Cleary, & Still, 2008). This latter finding shows that these results cannot be explained by *hedonic adaptation*: loss of sensitivity to positive stimuli after the good mood induction (in addition, happy participants' responses were sensitive to objectively prettier and uglier patterns). Thus, our results suggest that happy people experience familiarity, but for them it does not "glow" warmly. Future studies may explore whether having participants focus on the fluency per se (rather than familiarity) may restore its positive effects on judgment, consistent with the observation that the meaning of *fluent processing* is malleable and context sensitive (Schwarz, 2010).

Fluency and Naive Theories

A powerful determinant of how any subjective experience, including fluency, translates into judgments is the interpretational process applied to that experience (Schwarz, 2010). After all, the experiences do not mean much by themselves. In fact, there is now plenty of evidence that under some interpretations, the experience of ease can lead to negative judgments and the experience of difficulty, to positive judgments. One early study manipulated retrieval fluency of *neutral* childhood memories via the 4/12 event paradigm (Winkielman & Schwarz, 2001). We found that the *difficulty* of retrieving neutral childhood memories (12 events) led to *more* positive childhood evaluations, then *ease* of retrieval (4 events). This occurred when participants were given a bridging naive theory that implied difficulty in recalling childhood means nothing bad had happened. The opposite result was obtained when the naive theory provided to participants implied that difficulty in recalling memories means something bad must have happened. Of course, it may not be surprising that the evaluative implications of inferences from fluency may override any hedonic tinges deriving from the fluency itself. Still, it is important to remember that experiences can be fed into various judgmental processes that can override any potential default effects.

A recent study nicely made this point using manipulations of perceptual fluency. Song and Schwarz (2009) presented to participants the names of amusement park rides that were easy or difficult to pronounce. Then participants rated the risk level of those rides, with risk framed as either a positive attribute (exciting) or a negative attribute (dangerous). Amusement park rides with difficult to pronounce names were rated as more likely to make one sick (an undesirable risk), as well as more exciting and adventurous (a desirable risk). Note that such effects involve a higher-order bridging theory that links a fluency-dependent sense of familiarity with an interpretation of what kinds of risks are good. As a final example, research in consumer behavior showed that products advertised in "easy to read" fonts are sometimes perceived as *less* valuable. This is the case when the item is a luxury product that derives some of its appeal from its rarity—when fluent processing makes such a product seem highly familiar, the product loses its special appeal (Pocheptsova, Labroo, & Dhar, 2010).

HOW DOES IT WORK?: COMPUTATIONAL MECHANISMS

Verbal descriptions of fluency are often vague. What exactly does it mean to say that one stimulus is more fluent than another? What does it mean concretely that cognitions are consistent? Answers to these questions have been provided by computational models. In this section, we first describe some basic principles underlying such models, then some actual models of fluency phenomena.

Connectionism

Much fluency-modeling work has used the neural network approach, or *connectionism*. This approach models cognition in terms of the passage of activation among simple, neuron-like units organized in large, densely interconnected networks (McClelland & Rumelhart, 1981). The individual units function as simple processors that can influence each other through connections that vary in strength and sign (facilitatory or inhibitory). This massively interconnected and parallel architecture gives the neural network approach a certain neurophysiological realism and makes it suitable for a wide variety of applications. For more biological applications, one can conceptualize the network units as actual neurons, whereas for more psychological applications, one can treat the units as blocks of neurons or functional subsystems (O'Reilly & Munakata, 2000). Many different neural network architectures have been proposed, but for the current purposes, we first focus on a simple attractor neural network (Hopfield, 1984), then address just one of the more complex models (Huber & O'Reilly, 2003).

Fluency in a Simple Attractor Network

In a typical Hopfield network, representations are encoded as attractors of the network (i.e., states into which the network dynamics converge). The processing of information with the network can be seen as a gradual, evolving process, during which each neuron adjusts to the signal coming from other neurons. Because neurons are reciprocally connected, and because many paths connect one neuron to another, activation can reverberate dynamically through the network over simulated time steps until the network settles on the identified representation. For example, when presented with a to-be-recognized pattern, the network goes through a series of adjustments and after some time approaches a stable state, an *attractor*, corresponding to the "recognition" of a particular pattern (in a strict sense, it is just a reproduction of the pattern).

A typical Hopfield model can be extended with a simple control mechanism, which allows the network to monitor fluency of its own processing (Lewenstein & Nowak, 1989). Specifically, one can look at the network's "volatility," or the proportion of neurons changing their state at a given point. Fewer neurons change their state when the incoming, "to-be-recognized" pattern approximates a known pattern (an attractor). Another fluency-related property is the coherence of the signals received by the neurons. In the vicinity of an attractor the signals arriving from other neurons are consistent. A related criterion is the signal-to-noise ratio. In the vicinity of the attractor, signals from other neurons typically add up, resulting in a relatively large summary signal dictating the state of a given neuron.

Extension to Graded Representations

Traditional Hopfield networks use simulated neurons that are either "on" or "off," with no graded signal between these states (but see Hopfield, 1984). More realistic simulations use a continuous range of intermediary neuron values. This allows more graded measures of the magnitude and speed of settling into attractor states (e.g., O'Reilly & Munakata, 2000). However, because many applications focus on learning and representational change, large simulated time steps are used and settling occurs in less than 10 time steps. This makes it difficult to measure relatively subtle differences in settling time. For such applications, fluency is measured rather indirectly as differentiation—the magnitude of the most active units (Norman, O'Reilly, & Huber, 2000). Providing a more direct measure of fluency based on speed of processing, we have used neural simulations with millisecond time steps. This allows measurement of the time needed to achieve peak activation (Huber, 2008) and can be used to estimate reaction times (Huber & Cousineau, 2004).

Modeling Fluency-Inducing Variables

So far, we have discussed computational models of fluency in terms of general principles. But such models can be used to specify precisely the processing dynamics that underlie affective responses in several concrete empirical phenomena we discussed earlier. Recall that experimental psychological research found that positive affect, as well as other fluency-related judgments, can be enhanced by repetition, priming, figure–ground contrast, presentation duration, symmetry, and prototypicality. How does this work computationally?

The influence of these variables can be conceptualized as reflecting a process in which fluent patterns are represented by more extreme values of activation. This leads to fewer changes of state in the network, stronger signals in the network, more differentiated states of the neurons, and faster settling. It is easy to understand how and why priming, multiple repetition, longer duration, higher contrast, and greater clarity would result in more active and better differentiated representations, stronger signals in the network, and faster settling (see Drogosz & Nowak, 2006, for a simulation of mere exposure). But what about symmetry and prototypicality? Note that the representation of symmetrical patterns is stronger because of simplicity and redundancy. For example, the left and the right sides of the symmetrical faces are identical, essentially doubling activation of a particular feature. Mental representation of symmetrical patterns is also more robustly activated because of their position independence in recognition. For example, a symmetrical face looks the same from different angles (Johnstone, 1994). Similarly, fluency of prototypicality (responsible for the "beauty-in-averages" effect) results from converging exemplars creating a strong attractor (memory) for a prototype, which then is recognized with a faster settling time (Winkielman, Hooda, & Munakata, 2004).

Priming and Habituation

Let us devote a bit more attention to the recent modeling work on priming, as this is a central variable in all fluency research. This will also allow us illustrate how computational modeling work can inspire new insights into fluency. The study of priming phenomena typically involves observing the effect of a prime stimulus upon reaction time or accuracy to a closely related or identical target stimulus. Priming is a useful tool for studying fluency, because the effect of a related prime can be viewed as "priming the pump," thus making the target more fluent. For instance, in spreading activation theory (Collins & Loftus, 1975), the prime sends activation to related concepts, which results in preactivation of the target. This pre-activation results in faster identification of the target (i.e., the target has a head start).

The original modeling work on priming mechanisms dates back to the classic interactive activation model of word recognition by McClelland and Rumelhart (1981). They implemented the idea of preactivation in a neural network, and explained how the surrounding context provided by the other letters in a word enhance processing speed of a particular target letter (e.g., identifying the *R* in SKRD vs. the *R* in WORK). Their model obviously applies to phenomena at different levels, including words, sentences, pictures, sounds, and so forth.

In recent work, we extended the McClelland and Rumelhart model (1981) by considering the effect of not only preactivation between prime and target but also habituation that can occur when the prime is processed for an extended period of time. To explore how habituation impacts fluency, we modified the interactive activation model by including habituation of synaptic resources between each sending neuron and receiving neuron (Huber & O'Reilly, 2003). This way, the associations between features are temporarily lost with excessive processing of the prime. Simulation studies demonstrated that this is an effective mechanism for optimally clearing the deleterious effects of lingering activation from recently viewed objects (e.g., the prime) so as to clear the way for perception of subsequent objects (e.g., the target). In other words, habituation serves to eliminate the effect of preactivation that might otherwise induce undue target fluency. However, this method of reducing the effect of a prime comes with a cost, making it difficult to process an identical or closely related target stimulus. In other words, with habituation, a related target is sluggish to respond, which we term *disfluency*.

Our recent studies tested behavioral and physiological predictions of this habituation model using a repetition word-priming paradigm, in which we measured identification accuracy for target words briefly flashed onscreen at the threshold of awareness. We tested different prime durations, ranging from subliminal to 2 seconds of viewing the primed word and, as predicted, there was a gradual change from positive to negative priming with increasing prime duration (Huber, 2008). Follow-up experiments used the same paradigm while measuring event-related potentials (ERPs), revealing that neural dynamics follow the same time course as the habituation model (Huber, Tian, Curran, O'Reilly, & Woroch, 2008).

Note that our habituation model is quite broad in scope and predicts that similar effects should exist not just for reading but more generally for perception of any stimulus. Demonstrating this generality, we found the same crossover pattern from positive to negative priming as a function of increasing prime duration with repetition priming of social stimuli—faces (Rieth & Huber, 2010). Furthermore, the habituation model predicted that similar effects should exist at multiple levels of representation, including those that underlie memory retrieval. In a series of experiments we examined the manner in which short-term repetition priming can bias recognition of previously studied words. These experiments once again produced a gradual crossover pattern from positive to negative priming as a function of increasing prime duration, even though the task was to indicate whether an easily seen target word had been previously studied (Huber, Clark, Curran, & Winkielman, 2008).

All of these fluency–disfluency experiments used repetition priming. However, the habituation model predicts that similar effects should occur for similarity priming. To test this claim, we recently performed an experiment involving a task of speeded evaluations of highly valenced positive or negative pictures (e.g., a picture of puppies vs. a picture of a mutilated body). The primes were also highly-valenced pictures, and the manipulation was consistency between prime and target (both positive or both negative vs. mixed trials). As with the earlier experiments, prime presentation varied across a range of durations. Once again there was a gradual crossover pattern from positive to negative priming, even though nothing repeated except the implied valence of the picture (Irwin, Huber, & Winkielman, 2010). Similarly, we have demonstrated habituation effects for *semantic satiation* (Tian & Huber, 2010), revealing that the apparent loss of meaning for a repeated word is due not to habituation in the semantic representation but rather to loss of association (i.e., loss of synaptic resources) between the lexical form of a word and its meaning (i.e., the meaning is still there, but it cannot be reached). Based on this result, and the associated model, this suggests that conceptual disfluency (and perhaps conceptual fluency) is due to difficulty or ease of access rather than direct inhibition or activation of the concept.

Summary and Implications

This section has shown that experimental work on fluency can be enriched by computational modeling. We have reported a variety of successfully simulated experiments with relatively simple models of processing. These models allow us to link the seemingly different fluency variables in a coherent, mechanistically based framework. One interesting aspect of connectionist modeling is that its principles relate to, and mechanistically explain, old gestalt notions, such as the aforementioned *prägnanz*, or the idea that perception and cognition "prefer" and are facilitated by "good," concise, structures (Palmer, 1990).

In addition, these models revealed interesting new counterintuitive properties of fluency, such as a gradual transition from positive to negative priming. One interesting implication of such mechanical models is that they allow for a reconceptualization of some effects that were previously conceptualized as reflecting "strategic" mechanisms. For example, it is well known that it is harder to obtain the MEE when stimuli are presented consciously, for a long duration, or when they are presented in homogenous fashion (Bornstein, 1989). Similarly, priming effects get weaker or reverse when participants are aware of and attend to the priming stimulus (Jacoby et al., 1989). These effects are typically explained as reflecting a judgmental correction—drawing an analogy to situations in which subjects adjust their judgments because they are aware of the stimulus's undue influence (e.g., Strack, Schwarz, Bless, Kuebler, & Waenke, 1993). However, as described earlier, the satiation model views awareness, attention, or homogenous repetition as conditions that promote automatic habituation, thus eliminating or reversing the typically positive impact of primes.

MODELING COMPLEX COGNITION: FLUENCY AND COHERENCE OF ARGUMENTS AND STORIES

In our discussion so far, fluency and consistency have been mostly modeled with simple, isolated stimuli—a picture, a word, or a sentence. However, a few researchers bring a consistency-oriented modeling perspective to complex cognition (see Read & Simon, Chapter 4, and Topolinski, Chapter 6, this volume). The idea here is that consistency processes operate even when individuals are processing a series of complex, related elements, such as persuasive messages, logical arguments, stories, or legal evidence. Though these messages might be so large that they, as a whole, cannot be fit into working memory at one time, the overall ease of integrating the whole set of these messages (e.g., the whole legal argument) differs in fluency. And, as a consequence, this fluency influences the evaluation of the plausibility, favorability, and truthfulness of the entire argument. Below, we briefly outline how fluency can function in two such models: the coherentist account of reasoning

(Thagard, 2000) and the story model of jury decision making (Pennington & Hastie, 1992).

Coherentist accounts of thought propose that people evaluate conflicting propositions based on not only logical inferences but also the extent to which they are coherent with a larger network of facts and propositions (Thagard, 2000). One advantage of coherentist accounts is their ability to show how people handle information that seems to support different, irreconcilable conclusions. For example, consider a claim by your coworker that he did not receive your text message telling him you would like your book returned. If your text messages do sometimes disappear, then you cannot rule out the possibility that the text actually never reached him. But if, for example, you knew that your coworker had his own urgent use for the book, this information would be coherent with the belief that your coworker received and ignored the text message. Of course, the ideas that it never arrived or that it was ignored are incoherent (mutually inhibiting of each other's activation/acceptance). Additional evidence in favor of one of these propositions increases that proposition's activation, thus indirectly inhibiting the other proposition. Factors decreasing the acceptability of one proposition also make the other more acceptable. For example, if a norm that people should be taken at their word were present and activated, then this would inhibit the belief that the text was ignored and bolster the belief that it was never seen. Coherence theories thus provide a framework for seeing how people might make judgments about plausibility or truthfulness.

Mechanistically, Thagard proposes that one arrives at the choice between two conflicting propositions by first summoning relevant cognitive elements to mind. In the O. J. Simpson case (analyzed fully in Thagard 2000) elements that might be evidence include "Nicole Simpson was murdered," and "a bloody glove was found in O. J.'s home," and propositions such as "O. J. killed Nicole," "O. J. is innocent," and "The glove was planted." These cognitive elements can be thought of as nodes in a network. *Constraints*, or coherence relations between the nodes, can be either positive (for concepts coherent with one another) or negative (for concepts incoherent with one another). *Coherence relations* include explanation, deduction, facilitation, association, and a positive emotional connection. *Incoherence relations* include inconsistency and incompatibility, and negative association. *Positive coherence relations* between two elements mean that the activation of one element increases the activation of another, while *negative coherence constraints* mean that the activation of one decreases activation of the other. In computer models of coherence, the nodes in this network are then given an initial activation level and started on an updating cycle. The network eventually settles on an equilibrium activation state. Coherence with important nodes and membership in a group of mutually coherent nodes lead to a node being favored for acceptance. Those nodes with activation above a certain threshold are "accepted" and those below it are rejected, resulting in an accepted network of concepts, consistent with the more "preferred" solution.

The connection between fluency and coherence in such models is straightforward. Processing of one stimulus increases the activation of related concepts. This activated store of concepts allows one to understand quickly (fluently) the meaning of related sentences, and to integrate them more easily with related beliefs. One implication of this is that, in the coherentist model, the underlying mechanism for judging the believability of a statement, or for judging between two competing arguments or stories, might be our subjectively experienced fluency, and the positive–negative valence resulting from it. If we review a few competing explanations of an event, then the more "coherent" account creates a more positive feeling than the less coherent account and is thus accepted.

This idea fits nicely with data from the experiments on the story model of jury decision making (Pennington & Hastie, 1992). It has been shown that the order in which the evidence for jury trials is presented strongly influences whether the accused is found innocent or guilty. Presumably, jurors expect a story that smoothly connects motive, opportunity, and evidence as part of a narrative. A successful prosecution team takes advantage of this, hoping perhaps that the resulting fluency of integration gets misattributed to judgments of plausibility and truth, or simple liking of "their side." Though it is often recognized that the story model and coherence theories have a connection (Thagard, 2000), a fluency approach suggests a particular synthesis and makes specific empirical predictions. Elements of the story should be arranged in such an order that the subjective processing ease of the overall story is maximized, and that fluency, with its resulting affective consequences, gets properly attributed. In short, the fluency account provides the actual underlying mechanism for explaining why a set of coherent concepts is favored as an explanation.

CONCLUDING SUMMARY

We have discussed the relation among fluency, consistency, and coherence. We emphasized their connection to evaluation, though we also made clear that a host of other social cognitive judgments are influenced via similar principles. We covered behavioral and physiological experiments, and showed how computer models can capture the underlying processes. By taking this multimethod approach, and discussing phenomena from different levels of complexity, from simple words to intricate stories, we have tried to emphasize similarities in how the mind processes information. One limitation of previous formal approaches to cognition was their silence on the role of specific stimulus features. Thus, our chapter has also highlighted how the "what" is connected to the "how" of processing. Given all this, we see a bright future for fluency and consistency theories that are closely tied to mechanistic principles of the mind.

REFERENCES

Abelson, R. P. (1983). Whatever became of consistency theory? *Personality and Social Psychology Bulletin, 9*, 37–54.

Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968–1987. *Psychological Bulletin, 106*, 265–289.

Bornstein, R. F., & D'Agostino, P. R. (1994). The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition, 12*, 103–128.

Carver, C. S., & Scheier, M. F. (1990). Origins and functions of positive and negative affect: A control-process view. *Psychological Review, 97*, 19–35.

Collins, A. M., & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review, 82*, 407–428.

De Vries, M., Holland, R. W., Chenier, T., Starr, M. J., & Winkielman, P. (2010). Happiness cools the warm glow of familiarity: Psychophysiological evidence that mood modulates the familiarity–affect link. *Psychological Science, 21*, 321–328.

Drogosz, M., & Nowak, A. (2006). A neural model of mere exposure: The EXAC mechanism. *Polish Psychological Bulletin, 37*, 7–15.

Egan, L. C., Santos, L. R., & Bloom, P. (2007). The origins of cognitive dissonance: Evidence from children and monkeys. *Psychological Science, 18*, 978–983.

Elliot, A. J., & Devine, P. G. (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. *Journal of Personality and Social Psychology, 67*, 382–394.

Evans, J. St. B. T., Barston, J. L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition, 11*, 295–306.

Fenker, D. B., Waldmann, M. R., & Holyoak, K. J. (2005). Accessing causal relations in semantic memory. *Memory and Cognition, 33*, 1036–1046.

Fernandez-Duque, D., Baird, J. A., & Posner, M. I. (2000). Executive attention and metacognitive regulation. *Consciousness and Cognition, 9*, 288–307.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations, 7*, 123–146.

Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.

Gawronski, B., & Strack, F. (2004). On the propositional nature of cognitive consistency: Dissonance changes explicit, but not implicit attitudes. *Journal of Experimental Social Psychology, 40*, 535–542.

Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. New York: Viking.

Halberstadt, J. B. (2006). The generality and ultimate origins of the attractiveness of prototypes. *Personality and Social Psychology Review, 10*, 166–183.

Harmon-Jones, E., & Allen, J. B. (2001). The role of affect in the mere exposure effect: Evidence from psychophysiological and individual differences approaches. *Personality and Social Psychology Bulletin, 27*, 889–898.

Harmon-Jones, E., Amodio, D. M., & Harmon-Jones, C. (2009). Action-based model of dissonance: A review, integration, and expansion of conceptions of cognitive conflict. *Advances in Experimental Social Psychology, 41*, 119–166.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology, 21*, 107–112.

Hopfield, J. J. (1984). Neurons with graded response have collective computational

properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, USA, 81,* 3088–3092.
Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General, 137,* 324–347.
Huber, D. E., Clark, T. F., Curran, T., & Winkielman, P. (2008). Effects of repetition priming on recognition memory: Testing a perceptual fluency-disfluency model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 1305–1324.
Huber, D. E., & Cousineau, D. (2004). A race model of perceptual forced choice reaction time. *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 687–692). Hillsdale, NJ: Erlbaum.
Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science, 27,* 403–430.
Huber, D. E., Tian, X., Curran, T., O'Reilly, C., & Woroch, B. (2008). The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects. *Journal of Experimental Psychology: Human Perception and Performance, 34,* 1389–1416.
Irwin, K. R., Huber, D. E., & Winkielman, P. (2010). Automatic affective dynamics: An activation–habituation model of affective assimilation and contrast. In T. Nishida, L. C. Jain, & C. Faucher (Eds.), *Modelling machine emotions for realizing intelligence: Foundations and applications* (pp. 17–34). Berlin: Springer-Verlag.
Jacoby, L. L., Kelley, C. M., & Dywan, J. (1989). Memory attributions. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 391–422). Hillsdale, NJ: Erlbaum.
Johnstone, R. A. (1994). Female preference for symmetrical males as a by-product of selection for mate recognition. *Nature, 372,* 172–175.
Kellas, G., Paul, S. T., Martin, M., & Simpson, G. B. (1991). Contextual feature activation and meaning access. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 47–71). New York: Elsevier.
Kruglanski, A., & Gigerenzer, G. (2011). Intuitive and deliberate judgments are based on common principles. *Psychological Review, 118,* 97–109.
Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., et al. (2009). Computational social science. *Science, 323,* 721–723.
Lewenstein, M., & Nowak, A. (1989). Recognition with self-control in neural networks. *Physical Review, 40,* 4652–4664.
Mandler, G., Nakamura, Y., & Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 646–648.
McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88,* 375–407.
McGlone, M. S., & Tofighbakhsh, J. (2000). Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms. *Psychological Science, 11,* 424–428.
Monahan, J. L., Murphy, S. T., & Zajonc, R. B. (2000). Subliminal mere exposure: Specific, general, and diffuse effects. *Psychological Science, 6,* 462–466.
Morris, A. L., Cleary, A. M., & Still, M. S. (2008). The role of autonomic arousal in feelings of familiarity. *Consciousness and Cognition, 17,* 1378–1385.
Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida

word association, rhyme, and word fragment norms. Retrieved 6/6/2011, from *www.usf.edu/freeassociation*.

Norman, K. A., O'Reilly, R. C., & Huber, D. E. (2000). *Modeling hippocampal and neocortical contributions to recognition memory*. Poster presented at the 7th (Seventh) Cognitive Neuroscience Society Meeting, San Francisco, CA.

Nowak, A., & Vallacher, R. R. (1998). *Dynamical social psychology*. New York: Guilford Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Palmer, S. E. (1990). Modern theories of gestalt perception. *Mind and Language, 5*, 289–323.

Pennington, N., & Hastie, R. (1992). Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology, 62*, 189–206.

Phaf, R. H., & Roteveel, M. (2005). Affective modulation of recognition bias. *Emotion, 5*, 309–318.

Pocheptsova, A., Labroo, A., & Dhar, R. (2010). Making products feel special: When metacognitive difficulty enhances evaluation. *Journal of Marketing Research, 47*, 1059–69.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 353–363.

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition, 8*, 338–342.

Reber, R., & Schwarz, N. (2006). Perceptual fluency, preference, and evolution. *Polish Psychological Bulletin, 37*, 16–22.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review, 8*, 364–382.

Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science, 9*, 45–48.

Rhodes, G. (2006). The evolution of facial attractiveness. *Annual Review of Psychology, 57*, 199–226.

Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. *Journal of Experimental Psychology: Human Perception and Performance, 36*, 596–618.

Schwarz, N. (2010). Meaning in context: Metacognitive experiences. In B. Mesquita, L. F. Barrett, & E. R. Smith (Eds.), *The mind in context* (pp. 105 –125). New York: Guilford Press.

Schwarz, N., & Clore, G.L. (1996). Feelings and phenomenal experiences. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 433–465). New York: Guilford Press.

Schwarz, N., Sanna, L., Skurnik, I., & Yoon, C. (2007). Metacognitive experiences and the intricacies of setting people straight: Implications for debiasing and public information campaigns. *Advances in Experimental Social Psychology, 39*, 127–161.

Song, H., & Schwarz, N. (2008). Fluency and the detection of misleading questions: Low processing fluency attenuates the Moses illusion. *Social Cognition, 26*, 791–799.

Song, H., & Schwarz, N. (2009). If it's difficult to pronounce, it must be risky: Fluency, familiarity, and risk perception. *Psychological Science, 20*, 135–138.

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review, 8*, 220–247.

Strack, F., Schwarz, N., Bless, H., Kuebler, A., & Waenke, M. (1993). Awareness of the influence as a determinant of assimilation versus contrast. *European Journal of Social Psychology, 23*, 53–62.

Symons, D. (1979). *Evolution of human sexuality*. New York: Oxford University Press.

Tesser, A. (2000). On the confluence of self-esteem maintenance mechanisms. *Personality and Social Psychology Review, 4*, 290–299.

Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press.

Thornhill, R., & Gangestad, S. W. (1993). Human facial beauty: Averageness, symmetry, and parasite resistance. *Human Nature, 4*, 237–269.

Tian, X., & Huber, D. E. (2010). Testing an associative account of semantic satiation. *Cognitive Psychology, 60*, 267–290.

Titchener, E. B. (1910). *A textbook of psychology*. New York: Macmillan.

Weaver, K., Garcia, S. M., Schwarz, N., & Miller, D. T. (2007). Inferring the popularity of an opinion from its familiarity: A repetitive voice can sound like a chorus. *Journal of Personality and Social Psychology, 92*, 821–833.

Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1235–1253.

Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation leads to positive affect. *Journal of Personality and Social Psychology, 81*, 989–1000.

Winkielman, P., Halberstadt, J., Fazendeiro, T., & Catty, S. (2006). Prototypes are attractive because they are easy on the mind. *Psychological Science, 17*, 799–806.

Winkielman, P., Hooda, P., & Munakata, Y. (2004). *Neural network model of fluency for average patterns*. Unpublished manuscript, University of Denver, Denver, CO.

Winkielman, P., & Schwarz, N. (2001). How pleasant was your childhood?: Beliefs about memory shape inferences from experienced difficulty of recall. *Psychological Science, 2*, 176–179.

Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 189–217). Mahwah, NJ: Erlbaum.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology, 9*, 1–27.

Zajonc, R. B. (1998). Emotions. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The Handbook of social psychology* (pp. 591–632). Boston: McGraw-Hill.