UC Davis

UC Davis Previously Published Works

Title

Toward an Integrative Framework for Studying Human Evaluation: Attitudes Toward Objects and Attributes.

Permalink

<https://escholarship.org/uc/item/0pf1k3kp>

Journal

Personality and social psychology review : an official journal of the Society for Personality and Social Psychology, Inc, 22(4)

ISSN

1088-8683

Authors

Ledgerwood, Alison
Eastwick, Paul W
Smith, Leigh K

Publication Date

2018-11-01

DOI

10.1177/1088868318790718

Peer reviewed

Toward an Integrative Framework for Studying Human Evaluation:

Attitudes towards Objects and Attributes

Alison Ledgerwood*

Paul W. Eastwick*

Leigh K. Smith

University of California, Davis

*Authors contributed equally

In Press, Personality and Social Psychology Review

Word Count: 11,595

Abstract

Evaluation is central to human experience, and multiple literatures have studied it. This article pulls from research on attitudes, human and non-human mating preferences, consumer behavior, and beyond to build a more comprehensive framework for studying evaluation. First, we distinguish between evaluations of objects (persons, places, things) and evaluations of attributes (dimensions, traits, characteristics). Then, we further distinguish between *summarized attribute preferences* (a valenced response to a direction on a dimension, such as liking sweetness in desserts) and *functional attribute preferences* (a valenced response to increasing levels of a dimension in a set of targets, such as the extent to which sweetness predicts liking for desserts). We situate these constructs with respect to existing distinctions in the attitude literature (e.g., specific/general, indirect/direct). Finally, new models address how people translate functional into summarized preferences, as well as how these attribute preferences affect (a) subsequent evaluations of objects and (b) situation selection.

Keywords: attitude, evaluation, mate preferences, drivers of liking

Toward an Integrative Framework for Studying Human Evaluation:

Attitudes towards Objects and Attributes

Humans have vast stores of knowledge about themselves. Arguably one of the most important and pervasive kinds of self-knowledge is knowledge about one's own likes and dislikes. People talk about their preferences, communicate their values, and use their ideas about what they love and loathe as the basis for countless decisions—from selecting a date on a dating website to choosing a restaurant to deciding among potential job candidates to invite for an interview. But when a person considers their fondness for dorkiness in a date, or confides their secret penchant for Chick Fil-A, or emphasizes their preference for cooperativeness in a job candidate, where do these judgments come from?

In this article, we argue that in order to answer this question, researchers need a more comprehensive picture of attitudes than any single literature currently provides. First and foremost, we draw a distinction between attitudes toward *objects* (i.e., nouns that denote persons, places, or things, like fast food or Chick Fil-A), and attitudes toward *attributes* (i.e., adjectives that denote dimensions, like dorky or cooperative). Attributes are distinct from objects in that an attribute contains its own natural contrast (i.e., higher versus lower levels of itself); thus an attitude toward an attribute signifies the extent to which a person positively or negatively evaluates moving upward on a given dimension. For this reason, attitudes towards attributes (but not objects) imply some form of a dose-response association between the attribute and liking (e.g., “I like job candidates much more when they are more rather than less cooperative”).

Interestingly, although past research has studied both attitudes toward objects and attitudes toward attributes, they rarely appear as two coequal constructs of interest in the same literature. Moreover, attitudes toward attributes can be studied in fundamentally different ways—whereas some literatures focus on people's summary judgments about a given attribute (e.g., “I

like intelligence in a mate”), others focus on the extent to which the level of an attribute in a series of objects predicts a person’s liking for those objects (e.g., to what extent does the physical attractiveness of a job candidate predict a person’s evaluation of that candidate?). Overall, then, past work on attitudes toward objects and attributes has been fragmented. The present article aims to integrate this fragmented literature into a coherent and comprehensive picture of how people evaluate both objects and attributes.

Evaluating Objects and Attributes: The View from Different Literatures

Unsurprisingly, given the centrality of evaluation to human experience (Eagly & Chaiken, 1993; Greenwald, 1989; Katz, 1960; Smith, Bruner, & White, 1956), evaluations have been studied across multiple domains and disciplines, including the literatures on attitudes and persuasion, human mate preferences, non-human mate preferences, and consumer behavior. Each literature, however, tends to draw on its own traditional conceptualizations and measures, thereby unintentionally constraining its focus to one particular segment of the broader picture that we ultimately develop here.

Evaluation from the Perspective of the Attitudes Literature

The longstanding literature on attitudes in social psychology has traditionally focused on studying evaluations of objects (Albarracin, Johnson, & Zanna, 2005; Eagly & Chaiken, 1993). One of the most common and enduring definitions of an attitude in this literature is: “a psychological tendency that is expressed by evaluating a particular *entity* with some degree of favor or disfavor” (Eagly & Chaiken, 1993, p. 1, emphasis added). Other common definitions reflect a similar emphasis—for example, describing attitudes as “associations between a given object and a given summary evaluation of the object” (Fazio, 2007, p. 608). Although terms like “entity” and “object” are clearly intended to encompass the broad range of things toward which

people can express evaluations, the fact is that for decades, the attitude literature has primarily focused on studying attitudes toward persons, places, and things (e.g., evaluations of social issues, products, groups, individuals, and other entities, which tend to be nouns) rather than attitudes toward a direction on a dimension (e.g., evaluations of continuous traits, characteristics, and other dimensional qualities, which tend to be adjectives).

If pressed, researchers working within this tradition would likely assert that attributes are one kind of object, and therefore what we know about attitudes toward objects should generalize to attitudes toward attributes—an attitude toward a *toaster* and an attitude toward *crispy* should operate in much the same fashion. Indeed, in most basic research on attitudes, the object is incidental to the purpose of the study—a persuasion researcher would be equally likely to study attitude change toward a product, a person, or an issue; a study on attitude strength might just as readily assess attitudes toward social groups or censorship or squirrels (Chaiken & Ledgerwood, 2012; Chaiken, Pomerantz, & Giner-Sorolla, 1995; Fazio, 1995, 2007; Petty & Cacioppo, 1986; Roskos-Ewoldsen & Fazio, 1992; Thomsen, Borgida, & Lavine, 1995). But interestingly, even the most diverse list of attitude objects that a typical attitude researcher would consider tends to omit dimensional attributes entirely. The opening sentence of a chapter in a seminal text on attitude strength illustrates this focus well: “Attitudes are people’s evaluations of ‘objects’ as diverse as capital punishment, equality, Japanese, essay exams, me, and writing a chapter for Rich and Jon’s book on attitude strength” (Chaiken et al., 1995, p. 387).

This emphasis does not imply that attributes are absent from the attitude literature, but rather that their role is quite different from the role of attitudes toward objects. Most notably, attributes make an appearance in classic expectancy-value models of attitude formation and change and associated information integration models (Anderson, 1971; Fishbein & Ajzen, 1975;

Lampel & Anderson, 1968). According to these perspectives, a person's attitude toward an object is a function of (a) his or her subjective beliefs about which attributes characterize that object, and (b) his or her evaluations of each of these attributes.¹ Importantly, in research deriving from the expectancy-value tradition, evaluations of attributes were conceptualized as *inputs*—relevant and interesting only insofar as they could be used to predict an attitude toward an object (e.g., Anderson, 1965; Anderson & Barrios, 1961; Fishbein, 1963; Fishbein & Coombs, 1974; Kaplan, 1971, 1973). At a theoretical level, when scholars did occasionally speculate on the possible origins or processes underlying attitudes toward attributes, they assumed that these processes were identical to those involved in the formation of attitudes toward objects and therefore warranted no special consideration. In other words, evaluations of attributes (often drawn from prior norming data; Anderson, 1968) were considered a “sufficient” starting point for most studies, rather than an outcome worthy of investigation in its own right (Fishbein & Ajzen, 1975, p. 217). At a practical level, researchers seemed to abandon attempts to use experimental manipulations to change people's attitudes toward attributes, perhaps because “evaluations of attributes are often well anchored in extensive prior learning” and proved “difficult to alter” (Eagly & Chaiken, 1993, p. 237; see also Lutz, 1975). Thus, this early literature offers little insight about what factors influence an attitude toward an attribute (e.g., what factors lead someone to positively evaluate cooperativeness in job candidates or spontaneity in romantic partners).

¹ Just as attitude researchers have typically conceptualized the term “object” more broadly than we do here, Fishbein and Ajzen (1975) actually used the term “attribute” much more broadly than we do here, allowing it to include “any characteristic, quality, object, concept, value, or goal associated with the object” (Fishbein & Ajzen, 1975, p. 223). However, other literatures (e.g., the impression formation literature; Fiske, Neuberg, Beattie, & Milberg, 1987) tend to use the term “attribute” to refer more narrowly to characteristics or traits. In this manuscript, we adopt this narrow use of the term “attribute” to refer to dimensions (i.e., traits, characteristics, and other continuous qualities typically denoted by adjectives), just as we use a narrower definition of the term “object” to refer to things (i.e., persons, places, events, policies, and other bounded entities typically denoted by nouns).

In summary, although attributes were obviously important to early theories of attitude formation, the substantive research questions almost never concerned the processes that give rise to the evaluation of the attribute in the first place. As a result, researchers within the attitude tradition have rarely thought to ask how attitudes toward traits or other attributes form and change (see Salancik & Conway, 1975, for one notable exception).

Evaluation from the Perspective of the Mate Preferences Literature

Meanwhile, other literatures have focused instead on attitudes toward attributes in and of themselves. The literature on mate preferences (also called ideal partner preferences) provides one clear example—here, researchers have long been interested in answering questions about what characteristics people like or want in a partner. Beginning with sociological studies in the 1940s (Hill, 1945), researchers studying human mating have assessed people's evaluations of a range of attribute dimensions (e.g., physically attractive, intelligent) that people might find desirable.

In a typical study, participants are asked to provide a summary judgment of the extent to which they desire certain attributes in a partner by evaluating each on a rating scale. Enormous literatures across the fields of family studies (Christensen, 1947; Hill, 1945; Hudson & Henze, 1969), evolutionary psychology (Buss, 1989; Kenrick, Groth, Trost, & Sadalla, 1993; Li, Bailey, Kenrick, & Linsenmeier, 2002), and close relationships (Fletcher, Simpson, & Thomas, 2000; Fletcher, Simpson, Thomas, & Giles, 1999) have studied attitudes towards attributes in this manner, and straightforward extensions of this work have examined preferences for attributes of friends (Goodwin & Tang, 1991; Sprecher & Regan, 2002) and in-laws (Apostolou, 2007). Experimental manipulations of mate preferences are very rare (for three exceptions, see Eagly, Eastwick, & Johannesen-Schmidt, 2009; Kille, Forest, & Wood, 2013; Nelson & Morrison,

2005)—perhaps because mate preferences have proven (like other attribute preferences) to be “difficult to alter” —so this literature is largely dominated by correlational methods (e.g., correlations between mate preferences and a participant’s sex or self-views; Campbell & Wilbur, 2009). A similar approach to studying attitudes toward attributes appears occasionally in other topic areas as well, including studies of leadership (Pew Research Center Survey, Nov. 12-21, 2014), surgical training (Nisar & Scott, 2011), teaching effectiveness (Delaney, Johnson, Johnson, & Treslan, 2010), hiring discrimination (Uhlmann & Cohen, 2005), and legal procedures (Shestowsky, 2014).

Other Perspectives on Evaluation

Several other literatures on evaluation likewise focus on attitudes toward attributes, but conceptualize and measure these evaluations in a strikingly different way. Researchers studying mating preferences in non-human animals—constrained by the inability of the typical bird or fish or to offer cogent responses to questions like “How desirable is display intensity in a potential mate?”—have devised various creative ways of assessing the extent to which the level of a given attribute in a potential mate predicts the positivity of another animal’s evaluative response (Moller, 1988; Patricelli, Uy, Walsh, & Borgia, 2002; Thornhill, 1983). For instance, researchers interested in assessing female satin bowerbirds’ preferences for vocal mimicry abilities in a mate might first measure the accuracy and size of male birds’ vocal mimicry repertoires, and then use the strength of the association between these features and the males’ courtship success as a measure of the females’ preference for those attributes (Coleman, Patricelli, Coyle, Siani, & Borgia, 2007).

A similar approach to studying attitudes can be found in a number of other literatures. Researchers who study consumer preferences are often interested in assessing what they term

drivers of liking (Lawless & Heymann, 2010). Participants in such studies might be asked to evaluate a series of products that vary in the convenience of their packaging by ranking how likely they are to buy them, or to evaluate a series of olive oils that vary in bitterness using a nine-point Likert-type scale (Delgado & Guinard, 2011; Silayoi & Speece, 2005). Here, attitudes toward attributes are inferred from how strongly participants' liking for a series of products (e.g., olive oils) tracks the level of an attribute in that product (e.g., bitterness). Using a similar approach, research in organizational behavior examines attitudes toward attributes of organizations or job candidates (Heilman & Saruwatari, 1979; Turban & Keon, 1993), and research in political science examines attitudes towards attributes of immigrants or election candidates (Carnes & Lupu, 2016; Hainmueller & Hopkins, 2015). A handful of recent studies in the human mating literature have started to conceptualize attribute preferences using this approach as well (Brumbaugh & Wood, 2013; Eastwick & Finkel, 2008; Eastwick, Luchies, Finkel, & Hunt, 2014; Wood & Brumbaugh, 2009). Across these diverse domains, attitudes toward attributes are conceptualized not as a summary evaluation of the attribute in question (e.g., how much a person likes attribute X), but rather as the extent to which an attribute guides people's evaluations of a range of objects (e.g., how strongly the level of attribute X in a series of objects predicts a person's liking for each of those objects).

When viewed side by side, each of these literatures appears to be studying an important but fragmented aspect of how evaluative processes operate. By focusing almost exclusively on attitudes toward objects, attitude researchers may have inadvertently missed a number of questions that arise only when considering attitudes toward attributes as a topic of study in and of itself. Meanwhile, researchers studying attribute preferences in specific content domains such as human mating and consumer behavior may have overlooked the extent to which their research

questions could be informed by basic research on attitudes and social cognition (e.g., the role of self-perception in attitude formation, humans' susceptibility to biases; see section "New Questions and Future Directions"). Given that scholars working in these disparate literatures are fundamentally concerned with the psychology of evaluation—whether they realize it or not—there exists tremendous potential for integration.

Toward a Comprehensive Framework for Studying Attitudes

Figure 1 brings together different elements from the literatures reviewed above to depict a fuller picture of the constructs and measures that are relevant to the way that humans evaluate the world. Importantly, it depicts attitudes toward objects and attitudes toward attributes as of equivalent scientific interest; neither is subsidiary to the other.

The framework describes three evaluative constructs: Attitudes toward objects, and two kinds of attitudes toward attributes (i.e., summarized preferences and functional preferences). Notably, these constructs refer to distinct types of evaluation rather than distinct types of cognitive processes. That is, we do not assume that these three types of evaluation do or do not arise from distinct cognitive processes or mental representations; our framework intentionally leaves open the question of what cognitive processes give rise to these three types of evaluations and the extent to which these processes are similar or different. (For an in-depth discussion of the importance of separating levels of analysis in attitudinal research, see De Houwer, Gawronski, and Barnes-Holmes, 2013).

Attitudes toward Objects

An evaluation of an object is a valenced response to an entity—a person, place, or thing (typically denoted by a noun). This type of attitudinal construct is by far the most familiar to psychologists because it reflects the most common conceptualization of attitude in the

psychological literature (e.g., Eagly & Chaiken, 1993; Fazio, 2007; Zanna & Rempel, 1988); it needs little additional explication here. Much of what we know about attitudes towards objects—such as basic principles of persuasion (Chaiken & Ledgerwood, 2012; Petty & Wegener, 1998), the effects of direct experience and attitude accessibility (Glasman & Albarracin, 2006; Roskos-Ewoldsen & Fazio, 1992), the importance of subjective representation (Ledgerwood, 2014; Lord & Lepper, 1999), etc.—may generalize in some ways to attitudes toward attributes, although most of these extensions remain untested. Importantly, attitudes toward attributes also exhibit unique features that merit consideration. We turn now to examine these less familiar constructs in more detail.

Attitudes toward Attributes (Summarized)

A summarized evaluation of an attribute is a valenced response to a quality or dimension (typically denoted by an adjective). We will refer to this construct as a *summarized attribute preference*. The term “summarized” connects to the attitudes literature, where such overall judgments of positivity or negativity are typically conceptualized as valenced summaries of evaluation-relevant information (Eagly & Chaiken, 1993, 1998; Fazio, 1986; Ledgerwood, Trope, & Chaiken, 2010; Wilson, Lindsey, & Schooler, 2000), and the term “preference” connects to the mate preferences and ideal partner preferences literature, where these judgments are often assessed (Buss, 1989; Fletcher et al., 1999; Hill, 1945). Researchers examining summarized attribute preferences are often interested in assessing them within a particular content domain (e.g., romantic relationships), and therefore tend to assess summarized preferences by attaching them to the class of entities that are of interest to that domain (e.g., “to what extent do you desire the following qualities in a romantic partner?”). Yet even when a class of entities is not specified by a researcher, a person is likely to have one in mind (e.g., a person

evaluating the attribute “salty” may be implicitly considering “foods” as the relevant class of targets). A summarized preference for an attribute thus involves an evaluation of a *direction on a dimension*—how positively or negatively a person feels about moving up the scale of a given trait or characteristic—with respect to a given class of targets.

At first glance, summarized preferences may seem to share surface similarities with stereotypes in that both involve judgments about traits. But whereas summarized preferences are evaluative (to what extent do you *like* this trait in this group of people?), stereotypes are descriptive (to what extent do you think this trait *describes* this group of people?). For example, the extent to which people desire intelligence in a graduate student is a summarized preference, but the extent to which people believe that intelligence generally characterizes graduate students is a stereotype. Of course, these variables might correlate: The attributes that people believe are normative may tend to be attributes that people believe are desirable (Wood & Furr, 2016). Nevertheless, summarized preferences and stereotypes are conceptually distinct.

Attitudes toward Attributes (Functional)

A functional evaluation of an attribute is a valenced response to increasing levels of a quality or dimension in a set of targets. In other words, it is a predictive relationship between an attribute and liking—the association of (a) the level of an attribute in each of a series of targets with (b) liking for each of those targets. We will refer to this construct as a *functional attribute preference*. We use the term “functional” to acknowledge the centrality of such preferences in evolutionary (i.e., functional) approaches to the study of animal behavior. Indeed, this type of attribute evaluation would have served the clearest adaptive function in humans’ ancestral past as they encountered different possible mates, coalition partners, food sources, or environments. A person’s functional attribute preference reflects the extent to which an attribute actually predicts

his or her evaluative responses in practice when confronting targets that possess varying levels of the attribute.²

Functional preferences are most clearly depicted as within-person associations or slopes: For example, a given participant might like her various coworkers more to the extent that they are loyal, or she might like her dates more to the extent that they are attractive. Stronger within-person correlations (i.e., steeper slopes) imply stronger functional preferences. In most of the consumer preference and organizational behavior studies that assess them, functional preferences are measured in precisely this way, which allows researchers to assess individual differences in people's functional preferences (just as researchers typically assess individual differences in people's attitudes toward objects or summarized preferences for attributes). Importantly, it generally takes 2-3 times as many observations to reliably assess a slope than it does to reliably assess a mean (Cohen, 1992; Wood & Brumbaugh, 2009), so studies assessing individual differences in functional preferences often ask participants to evaluate many different targets that vary on the attribute of interest.

However, studies do not need to examine within-person associations in order to study functional preferences; it is possible for studies to capture the average functional preference in a group of participants (in essence, studying whether an attribute tends to be functionally liked in a given population, rather than measuring individual differences in functional preferences for the

² Our distinction between summarized and functional preferences has sometimes been described using the terms "stated preferences" and "revealed preferences" (Caruso, Rahnev, & Banaji, 2009; Eastwick et al., 2014; Wood & Brumbaugh, 2009). However, we prefer to avoid these terms for two reasons. First, the term "stated" refers to a particular, explicit measurement strategy, and we want to be careful to define summarized preferences to refer to a construct that may be measured either explicitly (i.e., stated) or implicitly. Second, behavioral economists frequently use the term "revealed preference" to mean simply "behavior" (Ariely & Norton, 2008; Beshears, Choi, Laibson & Madrian, 2008); in that field, the study of stated versus revealed preference correspondence is in fact the study of attitude-behavior correspondence. As described here, what we call functional preferences are evaluative; they are no more (or less) behavioral than evaluations of objects and summarized preferences. Therefore, we eschew the term "revealed preference" to avoid unintended parallels to classic attitude-behavior correspondence issues. We return to discuss more fully the question of how our distinction between summarized and functional preferences maps onto direct/indirect measurement and attitude-behavior correspondence later in the paper.

attribute). For example, the animal mating studies reviewed above frequently focus on the average functional preferences of a group of animals (e.g., the extent to which a male's vocal mimicry abilities predicts the willingness of a group of females to mate with him; Coleman et al., 2007, which is conceptually analogous to assessing the extent to which coworkers who vary in loyalty tend to be liked by their peers). Similarly, some studies in humans examine whether men and women differ substantially in their average functional preference for an attribute such as attractiveness in a romantic partner (they do not; for a meta-analysis see Eastwick et al., 2014).

Any continuous dimension that varies across a given class of targets can form the basis for a summarized or functional attribute preference. Thus, a person could express a summarized or functional preference for increasing sweetness in breakfast cereals or increasing sleekness in a car. Perhaps less intuitively, a person could also express a summarized or functional preference for increasing height in romantic partners, increasing square footage in backyards, or increasing number of coffee mugs in a collection. Notably, an attribute preference is necessarily tied to a particular class of entities: A person could express a summarized or functional preference for increasing loyalty in friends, increasing loyalty in soldiers, or increasing loyalty in coworkers, and these need not be the same. Moreover, the class of targets can range in specificity from the very specific (e.g., a preference for loyalty in yellow Labradors) to the very broad (e.g., a preference for loyalty in all creatures). In sum, then, the necessary components for an attribute preference are (1) a (summarized or functional) evaluative response elicited by (2) moving in a given direction on a dimension with respect to (3) a given class of targets.

Each of the evaluative constructs described above—evaluations of objects, summarized preferences, and functional preferences—are depicted in Figure 2 along with pictorial representations of their definitions. Each circle depicts an evaluative response (E) to each of the

following: an object (O); an attribute as a concept (A); and increasing levels of an attribute in a set of objects (a, aa, aaa, etc.).

Why Distinguish between Objects and Attributes?

As noted earlier, the distinction between objects and attributes has historical roots in classic expectancy-value models that describe how attitudes toward attributes serve as inputs for attitudes toward objects (Anderson, 1971; Fishbein & Ajzen, 1975). Perhaps the most crucial distinction between these two constructs, from the present perspective, is that an attribute contains its own natural contrast. In other words, by virtue of being a dimension, an attribute contains within itself higher *versus* lower levels; moving up the scale necessarily implies moving toward one endpoint (e.g., toward greater attractiveness in a mate) *relative to* the other (e.g., less attractiveness in a mate). Thus, an attitude toward an attribute is necessarily a *preference* for higher (vs. lower) levels of that attribute, or for lower (vs. higher) levels. For example, when a person says she likes “sweet” in wine, her statement can be taken to mean that she would rather her wines be sweeter; her summarized preference in effect represents a positive evaluation of moving upward rather than downward along the scale of sweetness.

In contrast, an object does not contain its own natural contrast. A person can of course compare their liking for two objects (e.g., “I like omelets more than pancakes”), just as a person can compare their liking for two attributes (e.g., “I like *salty* more than *sweet* in breakfast foods”). Both involve comparisons *between* two things (objects or attributes). Moreover, the second object or attribute involved in the comparison can vary; for example, a person can just as reasonably compare their liking for Coke versus Pepsi or Coke versus Sprite. A researcher assessing implicit bias could construct an Implicit Association Test that compares White versus Asian faces or White versus Black faces (Greenwald, McGhee, & Schwartz, 1998; Nosek &

Banaji, 2001; Nosek, Greenwald, & Banaji, 2005). Likewise, a researcher interested in exploring relative mate preferences (e.g., Fletcher, Tither, O'Loughlin, Friesen, & Overall, 2004) could compare liking for warmth versus status or warmth versus attractiveness. In contrast to such between-object or between-attribute comparisons, an evaluation of an attribute involves comparisons *within* that attribute—that is, an evaluation of a direction along a single attribute dimension. An attitude toward a single attribute merits the term preference all on its own because attributes contain their own natural contrast.

One can find cases where on the surface, an attitude toward an attribute and an attitude toward an object (or category of objects) would seem very similar: For example, at first glance, an evaluation of *sweet wines* (a category of objects) might seem interchangeable with an evaluation of the characteristic *sweet* in wines (an attribute). But these attitudes are not in fact the same construct. Consider an example in which three people all display equally positive evaluations toward sweet wines (Figure 3). Person A might also display a positive evaluation toward the attribute *sweet* in wines (suggesting that she likes wines more the sweeter they get). Meanwhile, Person B might display a neutral evaluation toward *sweet* in wines (suggesting that he likes sweet and not-sweet wines equally). And Person C might show a negative evaluation toward *sweet* in wines (suggesting that although she likes sweet wines, her favorite wines are not sweet at all). Thus, just because a person likes a category of attitude objects that possess a given attribute does *not* necessarily mean she has a (summarized and/or functional) preference for that attribute.

It is worth emphasizing too that an attribute preference is, by definition, specified with respect to a class of targets—not just one target. So for example, an attitude toward a particular proposal to increase student tuition (e.g., Chaiken & Eagly, 1983; Petty & Cacioppo, 1984) is an

attitude toward an object (the described proposal), not the attribute “increasing tuition.” In principle, increasing tuition could be conceptualized and measured as a summarized preference and specified with respect to a class of targets (e.g., “to what extent do you prefer greater rather than smaller tuition increases when considering various possible policies to generate revenue at our university?”). But in practice, persuasion studies have typically focused participants’ evaluations on one particular policy or issue or product, not an attribute that varies across a class of targets.

The distinction between objects and attributes is important because thinking about evaluations of attributes invites a second distinction that would not (and has not) occurred to researchers focused solely on evaluations of objects—namely, the distinction between summarized and functional preferences. We turn now to examine the importance of this second distinction in more detail.

Why Distinguish between Summarized and Functional Attribute Preferences?

Evaluations of attributes can be studied in two different ways—namely, as summarized preferences on the one hand and functional preferences on the other. There are empirical reasons to suspect that summarized and functional preferences are important to distinguish: In the few studies that have assessed both types of attribute preferences (e.g., Eastwick et al., 2014; Wood & Brumbaugh, 2009), their correspondence varies from very strong in some cases to nearly zero in others (see section “Model 1: How Do People Translate from Functional to Summarized Preferences?” below).

Furthermore, across the span of evolutionary history (i.e., phylogeny; Eastwick, 2009), these preferences likely emerged at different times. The capacity for functional preferences is ancient: Functional preferences can be observed in any organism whose approach or avoidance

responding is moderated by the presence of an attribute in a set of conspecifics, mates, predators, or prey (e.g., vocal mimicry prowess in a satin bowerbird, facial symmetry in a human face; Coleman et al., 2007; Rhodes, 2006). Summarized preferences are presumably evolutionarily more recent: They require an organism to somehow extract information about their liking for varying levels of an attribute to form an evaluation of the dimension as a concept in its own right. In fact, it is unclear whether any animals other than humans exhibit summarized preferences.

One might wonder, then, whether a summarized preference is simply an imperfect (and human-specific) measure of a functional preference. More broadly, should we consider functional versus summarized attribute preferences to be two different *measures* of the same construct, two different *behavioral constructs* (i.e., meaningfully different types of evaluative responses), or two different *mental constructs* (i.e., different kinds of mental representations or mental processes)?

The social and cognitive psychological literatures have witnessed many debates about whether two things are best considered (a) two measures, (b) two behavioral constructs, or (c) two mental constructs. For example, spontaneous and deliberate evaluations were originally cast as different measures of the same construct (i.e., option a): Spontaneous evaluation measures were (in contrast with more deliberative self-report measures) assumed to provide a “bona fide pipeline” to people’s true attitudes (Fazio, Dunton, Jackson, & Williams, 1995). Later, they were recast as different mental constructs (option c): Spontaneous (or “implicit”) attitudes and deliberate (“explicit”) attitudes were conceptualized as two distinct mental representations that were stored separately in the mind (Wilson et al., 2000). Most recently, scholars have suggested that spontaneous and deliberate evaluations should be treated as two separable behavioral

constructs (option b): distinct outcomes that arise from an unfolding set of mental processes (Gawronski & Bodenhausen, 2011). Thus, of the three possible approaches to distinguishing between spontaneous and deliberate evaluations, an outcome-level distinction (rather than a measurement or mental representation distinction) appears to afford the most theoretical traction. This approach has allowed researchers to ask questions about how spontaneous and deliberate evaluations relate, what variables predict one versus the other, and what processes they have in common versus what processes may uniquely contribute to only one of them (Gawronski & Bodenhausen, 2006; Huntsinger, 2013; Rydell, McConnell, Mackie, & Strain, 2006).

The history of this literature informs our current approach; we take as our starting point where theories of spontaneous and deliberate evaluations ultimately landed (i.e., option b). We suspect the summarized-functional distinction is more than a measurement distinction—that is, we believe it is worth studying summarized and functional preferences as distinct constructs in their own right. Doing so enables researchers to test empirically the possibility that these constructs might have different antecedents and different consequences. For example, we discuss below new results suggesting that an aspect of the social context can bias summarized preferences without influencing functional preferences, and we present new predictions suggesting that summarized and functional preferences may have different consequences for decision-making.

At the same time, we do not assume that summarized and functional attribute preferences (or evaluations of objects, for that matter) must involve domain-specific mental processes. In fact, we suspect that they are influenced by many of the same basic social-cognitive processes that play out in other domains. For instance, some of the research we describe below suggests that people may infer their summarized preferences using the same basic social-cognitive

processes that govern other forms of self-perception. Moreover, we think it would be a mistake to assume that different types of evaluative outcomes must reflect non-overlapping cognitive processes or distinct mental representations: Such an assumption conflates mental and behavioral levels of analysis (de Houwer et al., 2013) and hinders appropriate theoretical inference (Calanchini & Sherman, 2013). Instead, we recommend—at least as a starting point—treating summarized and functional preferences as distinct constructs at the level of evaluative responding (depicted as circles in Figure 2) while leaving open the question of what process or processes give rise to each one. Importantly, identifying the variables that differentially influence summarized and functional preferences will help to constrain the possible models of mental processes that can be postulated to underlie them (de Houwer et al., 2013).

Situating Our Framework in the Attitudes Literature

The asymmetry between objects and attributes described above highlights the importance of considering attitudes toward attributes as well as attitudes toward objects: Some concepts and questions about attitudes arise uniquely or at least primarily within the context of trying to understand attitudes toward attributes. But meanwhile, the attitude literature contains a number of other important distinctions that scholars have drawn over the course of the literature's long and rich history. It is therefore crucial to consider whether and how the distinctions we have introduced here map onto existing distinctions in the literature, including those that have been drawn between direct and indirect measures, general and specific attitudes, and attitudes and behaviors.

Direct and Indirect Measures

One distinction that has received considerable theoretical and empirical attention in the attitude literature is the distinction between direct and indirect measurement strategies (De

Houwer, 2006; Fazio, Jackson, Dunton, & Williams, 1995; Gawronski & Bodenhausen, 2011; Nosek, 2007). Importantly, both direct and indirect measures can be used to study attitudes toward objects, summarized attribute preferences, and functional attribute preferences (see Figure 1). For example, attitudes toward the book *Harry Potter and the Goblet of Fire* (an object) could be assessed directly (e.g., “how much do you like this book?”) or indirectly (e.g., using facial EMG to assess participants’ evaluative responses when presented with the book; Cunningham, Packer, Kesek, & van Bavel, 2009). Meanwhile, a summarized preference for *cleverness* in Harry Potter characters (an attribute) could also be assessed using either a direct measure (e.g., “how much do you like cleverness in Harry Potter characters?”) or an indirect measure (e.g., assessing facial muscle reactions to “cleverness in Harry Potter characters”).³ Likewise, a functional preference for *cleverness* in Harry Potter characters could be assessed using a direct measure (e.g., by asking participants to rate their liking for each character) or an indirect measure (e.g., measuring facial muscle reactions to each character); a researcher would then correlate (directly or indirectly measured) liking for each character with that character’s level of cleverness.

Importantly, this within-person correlation element of assessing functional preferences means that a functional preference measure requires one more layer of “indirectness” than the corresponding measure of a summarized preference. For example, imagine that a summarized preference for *cleverness* in a Harry Potter character is measured indirectly. Regardless of whether the assessment of the corresponding functional preference for *cleverness* in a Harry Potter character involves direct or indirect measures of liking, it will necessarily involve an

³ Summarized attribute preferences have been assessed primarily using direct self-reports. In fact, as far as we know, there is only a single published example of an indirect measure being used to assess summarized attribute preferences (i.e., a reaction time measure assessing positivity toward the attribute *physical attractiveness* in a romantic partner; Eastwick, Eagly, Finkel, & Johnson, 2011).

additional indirect component: namely, the correlation between liking for a series of targets (characters) and the level of the attribute (*cleverness*) in each of those targets.

Nevertheless, this connection between functional preferences and indirect measurement may be less useful than it initially appears: The processes that researchers typically map onto the distinction between direct and indirect measures do not map easily onto the summarized/functional preference distinction. For example, a functional preference for *intelligence* in a job candidate could be based on a person's careful, controlled, and deliberate judgments about a series of job candidates, whereas a summarized preference for *intelligence* in a job candidate could be quite spontaneous. Thus, although functional preference measurement has an inextricable layer of indirectness, the distinction between summarized and functional preferences does not reduce easily to the classic distinction between direct and indirect measures.⁴ Indeed, whereas indirect measures often were developed to circumvent social desirability concerns and other elements of control and awareness (e.g., Fazio, Jackson, Dunton, & Williams, 1995; Greenwald, McGee, & Schwarz, 1998; Hammond, 1948), our goal in this article is not to distinguish between evaluations that vary in how automatic versus controlled they are. Instead, we seek to distinguish between people's (directly or indirectly measured) beliefs about their attribute preferences—what they *think* they like—and the extent to which an attribute drives their (directly or indirectly measured) evaluation of various targets.

General and Specific Attitudes

Another classic distinction in the attitudes literature is the distinction between general and specific attitudes as detailed by Ajzen and Fishbein (1977, 2005). According to their *compatibility principle*, attitudes may seem to have weak predictive validity in cases where

⁴ Likewise, the distinction between summarized and functional preferences does not reduce easily to the distinction between explicit and implicit measures.

researchers specify the attitude object in mismatching ways—for instance, a general attitude (e.g., attitudes toward environmentalism) might fail to predict a specific behavior (e.g., voting on a city ordinance that would require composting) because the two attitude objects are not specified at the same level (see also Ledgerwood & Trope, 2010). To properly assess the predictive validity of attitudes, Ajzen and Fishbein suggested that researchers use general attitudes to predict general behaviors and specific attitudes to predict specific behaviors.

As with direct and indirect measures, the summarized versus functional distinction is not isomorphic with the general versus specific distinction: Preferences can be defined with respect to general or specific attributes (e.g., conscientious versus punctual) or with respect to general or specific classes of entities (e.g., Americans versus next-door neighbors). However, there are (at least) two useful connections one can draw between Ajzen and Fishbein's discussion of the compatibility principle and the present framework.

First, the compatibility principle offers an important lesson for researchers interested in measuring summarized and functional preferences: The correspondence between functional and summarized preferences may vary depending on whether the classes of targets are specified in the same way. For example, a functional preference for loyalty in Labrador retrievers (a specific class of targets) would presumably correlate more strongly with a summarized preference for loyalty in Labrador retrievers (the same specific class of targets) than with a summarized preference for loyalty in all creatures (a more general class of targets). Relatedly, a summarized preference for sweetness in breakfast cereals would probably correlate more strongly with a functional preference for sweetness in breakfast cereals that has been measured across 30

different breakfast cereals (a more general measure of functional preferences) rather than only three breakfast cereals (a more specific measure of functional preferences).⁵

Second, the closest analog of the summarized-functional preference distinction in the attitude literature is probably the distinction between two different measures of general attitudes that can be found in Ajzen and Fishbein's work: namely, an overall evaluation of a general attitude object (e.g., a person's favorability toward environmentalism) versus an average of evaluations of a set of specific attitude objects (e.g., a person's average favorability toward a city ordinance that would require composting, a new law that protects wildlife sanctuaries, and a policy to promote solar energy use). Ajzen and Fishbein (1977) actually treated these two kinds of measures as interchangeable methods for assessing the same general attitude construct. In contrast, our perspective might suggest that the distinction between these two forms of a general attitude goes deeper, at least in the case of attitudes toward attributes—an overall, summary evaluation of an attribute is not the same as the extent to which the attribute predicts evaluations of a series of specific targets. Importantly, by distinguishing between these constructs, we can begin to ask novel and interesting questions about how summarized and functional preferences might reciprocally influence each other, as we discuss in more detail below.

Attitude-Behavior Correspondence

Finally, one might be tempted to map the distinction between summarized and functional preferences for attributes onto the familiar distinction between attitudes and behavior in the attitudes literature (Ajzen & Fishbein, 1977; Fazio, 1990; Wicker, 1969). In decades past, both

⁵ Later in the paper, we will discuss studies that have observed a discrepancy between functional and summarized preferences; it is worth emphasizing here that such discrepancies emerge even when summarized and functional preferences are similarly specified (Eastwick & Finkel, 2008; Eastwick, Smith, & Ledgerwood, 2018). In other words, the compatibility principle is important to keep in mind when assessing summarized and functional preferences, but it is unlikely to account for much of the discrepancy between summarized and functional preferences that has been observed in past studies.

summarized and functional preferences would have been considered *attitudinal* rather than behavioral in the classic sense of those terms—both reflect evaluations rather than overt actions (Eagly & Chaiken, 1993). At the operational level, both constructs can be assessed using measures classically associated with attitude measurement, as described above (e.g., rating one's liking for *crispy* in breakfast foods vs. rating one's liking for a series of breakfast foods that vary in crispiness).

In some contemporary frameworks, however, summarized and functional preferences both would be considered behavioral (de Houwer et al., 2013; Sherman et al., 2008). That is, an explicit measurement strategy for both preferences might require participants to perform the behavior of using a pen to circle values on a rating scale, and an implicit measurement strategy for both preferences might require participants to perform the behavior of pressing a key on a keyboard as quickly as possible. These contemporary frameworks recast the attitude-behavior correspondence question as a question about the correspondence between different types of evaluative responses. As discussed above, it is for this reason that we represent summarized and functional preferences (as well as attitudes towards objects) as distinct forms of evaluative responding—that is, distinct latent constructs at the level of behavioral outcomes rather than cognitive representations or processes.

New Questions and Future Directions

By distinguishing between the constructs outlined in Figure 1, we can begin to ask new and interesting questions about how these variables relate to each other. Below, we highlight three interconnected models that posit a variety of psychological processes that may connect summarized preferences, functional preferences, and evaluations of objects. The first model depicts how people translate functional preferences into summarized preferences; the second

model depicts how functional and summarized preferences jointly affect evaluations of objects; the third model depicts how summarized preferences affect situation selection, which in turn feeds back to affect summarized preferences. Some elements of these models are supported by existing data, whereas other elements are novel predictions that have yet to be tested; when relevant, we describe existing studies that provide evidence consistent with a given pathway.

Model 1: How Do People Translate from Functional to Summarized Preferences?

Under many circumstances, a summarized preference and a functional preference for the same attribute in the same class of targets are logically equivalent. If the functional preference for an attribute is the extent to which a person's liking for a target depends on how much of the attribute the target possesses, then the summarized preference for that attribute should—normatively speaking—track the strength of this association, perhaps perfectly so. According to this line of reasoning, in order to form a summarized preference for an attribute, people should first observe the extent to which the attribute predicts their own likes and dislikes in the world, and then extract their summarized preference from this (and only this) information. In this sense, summarized preferences may be akin to meta-cognitive beliefs about one's own likes and dislikes—personal schemas built from real-world experiences. Although this normative argument is sound, people are not perfectly accurate self-perceivers (e.g., John & Robins, 1994; Ross, Lepper, & Hubbard, 1975), and so the extent to which people's explicit summarized preferences track their functional preferences—as well as the underlying mental processes that people use to translate functional into summarized preferences—ultimately must be addressed empirically. Model 1 depicts this translation process (Figure 4, large arrow).

As noted above, summarized and functional preferences are often studied in separate literatures. Yet a handful of studies have assessed both summarized and functional preferences

and estimated the size of the association between them. Figure 5 displays the magnitude of this association from the nine studies we could find that have reported the association between summarized and functional preferences. Included in this set of studies are (a) all relevant papers from a comprehensive review of the mating literature (Eastwick et al., 2014), (b) all relevant papers identified in reverse citation searches on Eastwick et al. (2014) and Wood and Brumbaugh (2009) for articles published after the acceptance of the Eastwick et al. (2014) article (yielding two additional papers: Brumbaugh & Wood, 2013; Eastwick & Smith, in press), and (3) two relevant papers brought to our attention by other scholars during discussions of the ideas in the present article (Caruso et al., 2009; DeBruine, 2006).

Interestingly, the magnitude of the association appears to vary considerably from one domain to the next (Figure 5). The largest effect size documented in this set of studies is the correspondence between summarized and functional preferences for the attribute *sweet* in a fairly simple class of targets: breakfast cereals. Specifically, participants' summarized preferences for sweetness in breakfast cereals correlated very strongly with their functional preferences for a series of cereals (i.e., the within-person association of cereal sweetness with liking for a set of ten cereals; Eastwick et al., 2014). In other words, when people say they like the attribute *sweet* in breakfast cereals, they do indeed tend to like cereals more to the extent those cereals are sweet.

All of the other studies in this set assessed preferences with respect to a very different kind of target—namely, other humans. In a study where participants evaluated photographs and descriptions of prospective teammates for a trivia contest, effect sizes were moderately sized; for instance, participants with strong summarized preferences for prior experience were more likely to select experienced over inexperienced teammates (Caruso, Rahnev, & Banaji, 2009). Studies

using photographs of potential dating partners have also tended to find moderately sized summarized-functional preference correlations: In two studies, participants' summarized preferences for qualities like *sexually suggestive* and *well-groomed* correlated moderately with their functional preferences when rating opposite-sex photographs that varied in these qualities (Brumbaugh & Wood, 2013; Wood & Brumbaugh, 2009). Similar studies found moderately sized summarized-functional preference correlations for *attractiveness* (Eastwick & Smith, in press) and *masculinity* (DeBruine, 2006) among participants who rated a series of photographs.

In contrast to this moderately sized correspondence when participants evaluate photographs and descriptions of other people, when people evaluate potential romantic partners whom they have actually met in person, the summarized-functional preference correlation essentially drops to zero. At speed-dating events, participants' summarized preferences were unrelated to their functional preferences (Eastwick & Finkel, 2007, 2008). Similar null associations emerged when participants evaluate opposite-sex peers whom they know well (Eastwick et al., 2014).

Taken together, these studies suggest that summarized-functional preference correspondence drops as the domain changes from breakfast cereals to photographs of humans to live humans. What psychological variable(s) might underlie this trend (the x-axis in Figure 5)? One intriguing possibility, suggested by the literature on covariation detection (e.g., Schaller & O'Brien, 1992), is complexity. In particular, as the variety of dimensions on which targets vary increases, it may become more difficult for a person to infer their preferences for each dimension on the basis of their experienced evaluations for a range of targets (see also Kelley, 1973). In other words, the process of translating a functional into a summarized preference (i.e., large blue arrow in Figure 4) may be moderated by complexity (thin blue arrow).

Recent evidence provides experimental support for this hypothesis (Eastwick, Smith, & Ledgerwood, 2018), using paradigms that draw on basic principles of self-perception and attitude formation (Bem, 1967, 1972; Fazio, Sherman, & Herr, 1982). These experiments first manipulated participants' functional preferences for an unfamiliar attribute called "Melb" in a novel set of targets and then measured their summarized preferences for that attribute. For example, in the strong (vs. weak) functional preference condition, "Melb" more strongly predicted the extent to which the participant's experience of each target was positive or negative. Participants then reported their summarized preferences for this attribute as a dependent measure—allowing us to test experimentally how people would translate a functional preference into a corresponding summarized preference. Participants' summarized preferences generally tracked the strength of the functional preference manipulation, but importantly, their performance on this task worsened when they had to track two traits instead of one. Moreover, consistent with the logic that participants use a self-perception process to extract their summarized preferences that parallels the process of perceiving other people, the results were similar regardless of whether participants made inferences about their own functional preferences or someone else's functional preferences (see Bem, 1967, 1972). In other words, people generally seem to be able to observe functional preferences and translate this information to a summarized preference, but their performance is hindered when the targets are complex and they have to track multiple traits (vs. when the targets are simple and they only have to track one trait).

People's ability to translate functional into summarized preferences is likely to be moderated by other factors beyond the complexity of the target. In fact, target complexity may simply be one example of a broad class of moderators that hinder the functional-summarized

inference process by taxing working memory (Arkes & Harkness, 1983; Pechmann & Ratneshwar, 1992; Shaklee & Mims, 1982). Returning again to Figure 5: When participants interact with live potential romantic partners (vs. photographs of partners), they may be more preoccupied with the challenges of rejecting (and being rejected by) partners (Joel, Teper, & MacDonald, 2014), which could reduce available working memory for tracking their functional preferences. Other variables that reduce or constrain working memory (e.g., stress or distraction while experiencing the functional preferences relevant for inferring a summarized preference) might similarly reduce people's ability to translate functional into summarized preferences.

Together, these considerations suggest that at least part of the reason why summarized and functional preferences might start to diverge under some circumstances is because people grow less able to infer their summarized preferences from their functional preferences. When and why might this matter? From a methodological perspective, when summarized and functional preferences diverge, researchers interested in assessing preferences for attributes (e.g., traits in a romantic partner or qualities in an organization) might reach very different conclusions depending on whether they assess summarized or functional preferences (e.g., Eastwick & Finkel, 2008; Eastwick et al., 2014). From a psychological perspective, summarized-functional divergence raises the interesting possibility that people may not have unqualified insight into their own likes and dislikes, which could have interesting downstream implications—a possibility that we now turn to discuss in more detail.

Model 2: How Do Attribute Preferences Influence Evaluations of Relevant Objects?

Both functional and summarized attribute preferences could plausibly influence how people evaluate an attitude object; this process is depicted in Model 2 (in Figure 6). Functional preferences might exert a fairly direct impact on people's responses to objects that they

encounter (top arrow): A person's functional preference for adventurousness in a date might increase their likelihood of asking out a more versus less adventurous coworker; their functional preference for natural brightness in a living space might lead them to put down an offer on a brightly lit rather than a dimly lit home. Of course, the effect of any given functional preference on an evaluation of an object will be limited by the sheer number of functional preferences that contribute to the evaluation (Ahadi & Diener, 1989); a given functional preference can have a larger impact if it is one out of five rather than one out of fifty relevant attributes. Regardless of this limit, however, the functional preference should interact with the level of the attribute in a subsequently encountered object to predict participants' evaluative responses toward that object. Experimental tests of this prediction that allowed for causal inferences would be especially valuable.

The effect of summarized preferences on object evaluation (bottom arrow) is conceptually identical to the classic expectancy-value models of attitude formation described earlier: Value (i.e., the summarized preference) should interact with expectancy (i.e., the level of the attribute in the object) to predict the attitude. Although the expectancy-value literature is vast, the evidence for this particular pathway actually remains quite murky in the attitude literature because studies from the 1970s and 1980s calculated this pathway using an incorrect statistical approach (Bagozzi, 1984; Evans, 1991).⁶ This pathway is surely positive on average (see e.g., Table 3 in Bagozzi, 1984), but existing research in this literature does not offer precise tests of the summarized preference $\times$ attribute interaction on object evaluations.

Fortunately, precise tests of this pathway can be found in the human mating literature: Several relevant studies have examined the predictive impact of (a) participants' summarized

⁶ Specifically, those studies predicted the attitude dependent measure from the summarized preference $\times$ attribute interaction without using multiple regression to control for the main effects of the summarized preference and the attribute. Failure to control for those main effects leads to (dramatically) upwardly biased estimates.

preferences (i.e., value) and (b) the extent to which a potential romantic partner possesses the relevant attribute (i.e., expectancy) on (c) participants' romantic desire for the partner (Eastwick, Finkel, & Eagly, 2011; Huang et al., 2018). Moreover, these studies have begun to ask not only *whether* summarized preferences predict object evaluations, but also *when* they are more or less likely to do so. Drawing on construal level theory and related perspectives (Ledgerwood, 2014, Park, Young, & Eastwick, 2015; Trope, Ledgerwood, Liberman, & Fujita, 2018), these researchers have argued that summarized preferences—as overall evaluations of an attribute at a global or schematic level—can be conceptualized as relatively abstract regulatory tools. Therefore, construal level theory generates the prediction that summarized preferences (like other abstract tools) should more strongly guide evaluative responses to potential partners that are psychologically distant (e.g., hypothetical) than close (e.g., actually encountered face-to-face).

In studies examining this possibility (Eastwick, Finkel, & Eagly, 2011; see also Huang et al., 2018), participants indicated their summarized preferences with respect to a set of attributes. In a separate session, they first perused a dating profile (i.e., a distant, hypothetical potential partner) that contained either two of the attributes they evaluated positively or two of the attributes they evaluated negatively. When participants evaluated the dating profile, they expressed more desire for the partner whose attributes matched their positive (rather than negative) summarized preferences, consistent with the bottom pathway in Model 2 and classic expectancy-value models. Then, participants had a face-to-face interaction with the same potential partner (i.e., a real, live interaction). After this interaction, the predictive effect of summarized preferences on desire for the partner vanished. In other words, summarized preferences predicted downstream object (in this case, partner) evaluations when the object was

hypothetical and therefore psychologically distant, but not when the object was real and therefore psychologically close. Thus, the psychological distance of an attitude object may moderate the summarized preference to object evaluation pathway (thin blue line in Model 2). Additional research should continue to test this possibility and probe its generalizability across distance dimensions (e.g., temporal, spatial, and social distance as well as hypotheticality) and types of attitude objects (e.g., policies and events as well as people).

Finally, it may be fruitful to consider the possibility of a reverse causal pathway in Model 2. Indeed, some research suggests that people may actually shift their summarized attribute preferences to justify their evaluations of a relevant object. For example, when choosing between two applicants for a position (e.g., police chief), participants may elevate the desirability of ambiguously relevant attributes (e.g., education level) that uniquely characterize their preferred applicant as a way of rationalizing their decision post hoc (Uhlmann & Cohen, 2005; see also Hodson, Dovidio, & Gaertner, 2002; Norton, Vandello, & Darley, 2004, Uhlmann & Cohen, 2007). In the context of Model 2, such a process reflects a reverse causal pathway such that a salient attribute of a selected (vs. unselected) object boosts the summarized preference rating of that attribute.

Model 3: How Do Attribute Preferences Influence Situation Selection?

Another way in which attribute preferences could exhibit predictive validity is by affecting situation selection. Just as attitudes towards political issues affect people's desire to enter settings containing like-minded individuals (Motyl, Iyer, Oishi, Trawalter, Nosek, 2014; Snyder & Kendzierski, 1982), summarized preferences might guide people to place themselves in settings featuring objects that contain a higher or lower average level of an attribute. This form of situation selection typically happens at a distance, before a person has encountered or

experienced anything about the situation first-hand. Given that summarized preferences tend to predict object evaluations especially well when those objects are somewhat distant (Eastwick, Finkel, et al., 2011), summarized preferences seem like strong candidates as predictors here.

Model 3 (Figure 7) depicts a summarized preference predicting a downstream choice between two situations containing objects with a low or high average level of an attribute. For example, a person's summarized preference for adventurousness might lead them to select a dating website geared toward people who love travel and outdoor activities; their summarized preference for natural brightness in a living space might lead them to ask their realtor to show them only brightly lit houses; their summarized preference for ease of classes might lead them to attend a study abroad program known for having courses that are especially easy. In this way, summarized preferences may lead people to select themselves into particular kinds of situations, which in turn may systematically restrict the range of a given attribute dimension that the person encounters (e.g., only highly adventurous potential dates, brightly lit houses, or easy classes).

When summarized preferences correspond well to functional preferences, this type of situation selection could be very beneficial—people would tend to select themselves into situations full of targets that they especially like. But when summarized preferences and functional preferences diverge (e.g., when the targets are complex or the task constrains working memory, as discussed earlier), people might end up selecting themselves into situations comprised of the targets they *think* they like more but not the targets they would *actually* like more. For instance, a person with a strong summarized preference for adventurousness in a partner might decide to pay a premium to join an adventurous dating website, but experience no actual benefit in terms of happiness or satisfaction—that is, had she joined an alternative website

featuring less adventurous potential partners, she would have been just as satisfied with her options.

Thus, summarized preferences may sometimes lead a person astray in terms of selecting themselves into situations where they would encounter targets they like more (vs. less). This problem may be especially likely to occur in the presence of variables that bias summarized preference judgments without affecting functional preferences. For instance, recent research suggests that one likely source of bias in summarized preferences is the overall amount of the attribute present in the immediate environment (e.g., whether the average adventurousness in a pool of dates is high vs. low), irrespective of a person's actual functional preference. More specifically, in some conditions of the Eastwick et al. (2018) studies described above, participants either did (or did not) encounter an environment in which one attribute tended to be especially prominent relative to other attributes (akin to a context in which potential dating partners tended to be especially adventurous on average). When participants had to track multiple traits (i.e., a condition with high complexity, the moderator depicted in Model 3), participants tended to infer stronger summarized preferences for attributes that were on average higher (vs. lower) in the set of targets they happened to be evaluating. Importantly, this effect emerged even though functional preferences were held constant across conditions: In all conditions, participants' evaluative experiences with the targets were half positive and half negative (e.g., half of the encountered targets were liked and half were disliked), and the attribute predicted positive versus negative outcomes identically. Thus, the only difference between conditions was the average level of the attribute in the encountered targets, and higher average levels of an attribute biased summarized preferences upwards. Intriguingly, this result suggests that if we surround ourselves with adventurous partners, we will infer that we have a stronger

preference for adventurous partners, regardless of the extent to which partners' adventurousness predicts our positive experiences with them (i.e., our actual functional preference for adventurousness). In this way, the model explains how summarized preferences have the potential to create a self-exacerbating feedback loop: Summarized preferences could lead people to select into situations that constrain the range of an attribute, which in turn could bias their summarized preferences to become more extreme, which could then bias situation selection yet further.

Additional Considerations

Inferences from limited experiences. Another interesting direction for future research is to consider how much information people *think* they need in order to infer a summarized preference from a functional one. Recall that objectively, a researcher needs to use a relatively large number of targets to reliably assess a participant's functional preference (or any other slope; Cohen, 1992; Wood & Brumbaugh, 2009). Do people have insight into the number of targets they would need to experience to provide enough information to accurately translate their own functional preference into a summarized preference for an attribute? It seems plausible that people might be willing to infer a summarized preference for an attribute after encountering and comparing only a few targets that vary along that attribute dimension, but that such inferences might not correspond to functional preferences very strongly.

In fact, people might even be willing to infer a summarized preference after encountering only a *single* target, simply by noticing the co-occurrence of their liking for a particular target and the presence of the attribute in question (e.g., upon trying wine for the first time, a person might notice that (a) she likes it and (b) it is sweet, and conclude that she prefers sweetness in wines). If people were in fact willing to do this, they would not be comparing their liking for

multiple targets that vary on the attribute (which earlier we noted was logically required for people to accurately know the extent to which they like an attribute). In other words, the fact that people could not actually have observed their functional preference in this situation might not prevent them from making a summarized preference judgment if asked (see also Zajonc, 1980). Future research should continue to probe the processes that underlie people's inferences about their preferences for attributes and the conditions under which they are willing to make these inferences, as well as the extent to which people's confidence in their judgments does or does not track their accuracy.

Nonmonotonic functional preferences. One complexity that we have not yet discussed is that in principle, functional preferences can have different shapes. Some functional preferences will be linear and others will be asymptotic (e.g., the attribute has diminishing returns on liking as the attribute increases); in both cases, moving one direction on the attribute dimension is consistently associated with greater liking (never less—in other words, the preference is monotonic). But some functional preferences may be nonmonotonic. Consider, for example, preferences for the attribute *chaste* in the mating domain: Functional preference investigations have suggested that people want their partners to have some, but not too much, prior sexual experience (Kenrick, Sundie, Nicastle, & Stone, 2001). In other words, functional preferences for some attributes may exhibit an optimal point, such that any deviation from that optimum results in decreased liking for a given target.

If researchers have reason to anticipate that a given functional preference is strongly nonmonotonic, they should consider alternative assessment methods designed to capture this kind of shape. For example, in the sensory research literature, scholars sometimes measure nonmonotonic functional preferences with *just-about-right* scales (Lawless & Heymann, 2010;

Rothman & Parker, 2009; van Trijp, Punter, Mickartz, & Kruithof, 2007). In designs that use the just-about-right scale, participants evaluate targets, just as with other functional preference assessment strategies. The main difference is that the participant evaluates each target with respect to a given attribute on a scale from “too low” to “too high” with a middle anchor of “just about right.” After the participant evaluates different targets, the participant’s functional preference is inferred from the targets that the participant rated as just-about-right. Just-about-right scales do have shortcomings; for example, people exhibit a centering bias such that they tend to rate the medium amount of an attribute from among the targets they happen to encounter as just-about-right (Lawless & Heymann, 2010; for a similar evaluative bias, see Lick & Johnson, 2014). Nevertheless, when people’s functional preferences for a given attribute are likely to be nonmonotonic, the application of just-about-right scales to the attitudinal domain could reveal novel insights.

Meanwhile, it is also important to consider the possibility that people might respond to summarized preference measures in a way that poorly differentiates between monotonic and nonmonotonic underlying functional preferences. In the typical case where the preference is approximately linear, participants could respond to summarized preference measures as we have assumed above (and as research that focuses on summarized preferences typically assumes as well): by evaluating the extent to which they like increasing levels of an attribute. But if the preference is nonmonotonic, people might respond to these measures by identifying the optimal level of the attribute. For example, a person rating her summarized preferences for a series of traits in a romantic partner could use a 5 on 9-point scale to indicate that she has a modest preference for partners who are more rather than less *punctual* (i.e., a moderate and monotonic preference), but she could also use a 5 on a 9-point scale to indicate that she prefers her partners

to be only moderately ambitious rather than lower or higher on the *ambitious* dimension (i.e., a nonmonotonic preference). If the underlying functional preference shape is unknown, the meaning of the summarized preference judgment becomes highly ambiguous. Thus, future research needs to more closely investigate how summarized preferences map onto functional preference shapes and, if necessary, develop new summarized preference measures that allow participants to clearly differentiate monotonic and nonmonotonic attribute preferences.

Connections to other literatures. Finally, several literatures that lie beyond the domain of attitudes may be fruitfully connected to the processes and constructs discussed above. For example, the (related) literatures on covariation detection (Alloy & Tabachnik, 1984), rule-based contingency learning (Allan, 1993), and illusory correlations (Fiedler, 2000) have examined how well people are able to discern the relation between two variables. Key moderating variables identified in these literatures (e.g., the biasing effects of expectations or rare events) may also be relevant to the way that people translate functional into summarized preferences (Model 1); indeed, these literatures inspired the experimental manipulations of functional preferences reported in Eastwick et al. (2018). In the self-perception realm, studies by Wilson and colleagues (e.g., Wilson, Laser, & Stone, 1982) examined whether people could be taught to accurately perceive the causes of their moods; similar training paradigms might boost the correspondence between summarized and functional preferences. Of course, there are substantial differences between these literatures and our framework: For instance, these paradigms typically directed people to pay attention to the relation between two variables or to make causal inferences about the effect of a stimulus on an outcome, whereas people presumably form summarized preferences even in the absence of such directives. Nevertheless, there are likely to be underdeveloped connections between the attitude literature and the literatures on covariation

detection and the perception of causes of affect—connections that are highlighted by the distinction we have drawn between summarized and functional preferences.

Conclusion

Humans often learn about the world through experience and then translate that experience into knowledge—a process captured by the concept of empiricism. Of course, empiricism forms the basis of the scientific enterprise, but people can also be empirical in developing their own self-knowledge—naïve theories about their own personal likes and dislikes (Heider, 1958; Wegener & Petty, 1998). Indeed, despite the fact that many organisms have preferences for objects and for attributes, humans are perhaps unique among animals in their ability to translate their experiences with objects and attributes into verbalizable knowledge about their preferences—to not only like and dislike, but also think about and communicate their likes and dislikes.

How well do people translate their evaluative experiences into evaluative knowledge, and under what circumstances? As the preceding review reveals, we have only just begun to address this question. Studying summarized and functional preferences for attributes could shed new light on the processes underlying evaluation, especially as future studies contrast how people learn about their attitudes toward objects versus attitudes toward attributes. Meanwhile, the study of people's ability to translate functional into summarized preferences could tell us a great deal about the impressive powers—and also limits—of humans' naïve empiricism about the self.

References

- Ahadi, S., & Diener, E. (1989). Multiple determinants and effect size. *Journal of Personality and social Psychology*, 56, 398-406.
- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, 114, 435-448.
- Ajzen, I., & Fishbein, M. (1977). Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin*, 84, 888-918.
- Ajzen, I., and Fishbein, M. (2005). The influence of attitudes on behavior. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *The Handbook of Attitudes* (pp. 173–221). Mahwah, NJ: Erlbaum.
- Albarracín, D., Johnson, B. T., & Zanna, M. P. (Eds.). (2005). *The Handbook of Attitudes*. Mahwah, NJ: Erlbaum.
- Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological review*, 91, 112-149.
- Anderson, N. H. (1965). Averaging versus adding as a stimulus-combination rule in impression formation. *Journal of Experimental Psychology*, 70, 394-400.
- Anderson, N. H. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, 9, 272-279.
- Anderson, N. H. (1971). Integration theory and attitude change. *Psychological Review*, 78, 171-206.
- Anderson, N. H., & Barrios, A. A. (1961). Primacy effects in personality impression formation. *The Journal of Abnormal and Social Psychology*, 63, 346-350.

- Apostolou, M. (2007). Sexual selection under parental choice: The role of parents in the evolution of human mating. *Evolution and Human Behavior*, 28, 403-409.
- Ariely, D., & Norton, M. I. (2008). How actions create—not just reveal—preferences. *Trends in Cognitive Sciences*, 12, 13-16.
- Arkes, H. R., & Harkness, A. R. (1983). Estimates of contingency between two dichotomous variables. *Journal of Experimental Psychology: General*, 112, 117-135.
- Bagozzi, R. P. (1984). Expectancy-value attitude models an analysis of critical measurement issues. *International Journal of Research in Marketing*, 1, 295-310.
- Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, 74, 183-200.
- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology*. New York: Academic Press, 1972.
- Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics*, 92, 1787-1794.
- Brumbaugh, C. C., & Wood, D. (2013). Mate preferences across life and across the world. *Social Psychological and Personality Science*, 4, 100-107.
- Buss, D. M. (1989). Sex differences in human mate preferences: Evolutionary hypotheses tested in 37 cultures. *Behavioral and Brain Sciences*, 12, 1-14.
- Calanchini, J., & Sherman, J. W. (2013). Implicit attitudes reflect associative, non-associative, and non-attitudinal processes. *Social and Personality Psychology Compass*, 7, 654-667.
- Campbell, L., & Wilbur, C. J. (2009). Are the traits we prefer in potential mates the traits they value in themselves? An analysis of sex differences in the self-concept. *Self and Identity*, 8, 418-446.

- Carnes, N., & Lupu, N. (2016). Do voters dislike working-class candidates? Voter biases and the descriptive underrepresentation of the working class. *American Political Science Review*, 110, 832-844.
- Caruso, E. M., Rahnev, D. A., & Banaji, M. R. (2009). Using conjoint analysis to detect discrimination: Revealing covert preferences from overt choices. *Social Cognition*, 27, 128-137.
- Chaiken, S., & Eagly, A. H. (1983). Communication modality as a determinant of persuasion: The role of communicator salience. *Journal of personality and social psychology*, 45, 241-256.
- Chaiken, S., & Ledgerwood, A. (2012). A theory of heuristic and systematic information processing. In P. A. M. Van Lange, A. W. Kruglanski, & E. T. Higgins (Eds.), *Handbook of Theories of Social Psychology* (pp. 246-266). London, UK: Sage Publications Ltd.
- Chaiken, S., Pomerantz, E.M., & Giner-Sorolla, R. (2005). Structural consistency and attitude strength. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude Strength: Antecedents and Consequences* (pp. 387-412). Mahwah, NJ: Erlbaum.
- Christensen, H. T. (1947). Student views on mate selection. *Marriage and Family Living*, 9, 85-88.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112, 155.
- Coleman, S. W., Patricelli, G. L., Coyle, B., Siani, J., & Borgia, G. (2007). Female preferences drive the evolution of mimetic accuracy in male sexual displays. *Biology Letters*, 3, 463-466.
- Cunningham, W. A., Packer, D. J., Kesek, A., & Van Bavel, J. J. (2009). Implicit measurement of attitudes: A physiological approach. *Insights from the new implicit measures*, 485-512.

- DeBruine, L. M., Jones, B. C., Little, A. C., Boothroyd, L. G., Perrett, D. I., Penton-Voak, I. S., ... & Tiddeman, B. P. (2006). Correlated preferences for facial masculinity and ideal or actual partner's masculinity. *Proceedings of the Royal Society of London B: Biological Sciences*, 273, 1355-1360.
- De Houwer, J. (2006). What are implicit measures and why are we using them? In R. W. Wiers & A. W. Stacy (Eds.), *The handbook of implicit cognition and addiction* (pp. 11-28). Thousand Oaks, CA: Sage Publishers.
- De Houwer, J., Gawronski, B., & Barnes-Holmes, D. (2013). A functional-cognitive framework for attitude research. *European Review of Social Psychology*, 24, 252-287
- Delaney, J., Johnson, A. N., Johnson, T. D., & Treslan, D. L. (2010). *Students' Perceptions of Effective Teaching in Higher Education*. Memorial University of Newfoundland, Distance Education and Learning Technologies.
- Delgado, C., & Guinard, J. X. (2011). How do consumer hedonic ratings for extra virgin olive oil relate to quality ratings by experts and descriptive analysis ratings? *Food Quality and Preference*, 22, 213-225.
- Eagly, A. H., & Chaiken, S. (1993). *The Psychology of Attitudes*. Harcourt Brace Jovanovich College Publishers.
- Eagly, A. H., & Chaiken, S. (1998). Attitude Structure and Function. In D. T. Gilbert, S. T. Fiske & G. Lindzey (Eds.), *The Handbook of Social Psychology, Vols. 1 and 2 (4th ed.)* (pp. 269-322). New York, NY, US: McGraw-Hill.
- Eagly, A. H., Eastwick, P. W., & Johannesen-Schmidt, M. (2009). Possible selves in marital roles: The impact of the anticipated division of labor on the mate preferences of women and men. *Personality and Social Psychology Bulletin*, 35, 403-414.

- Eastwick, P. W. (2009). Beyond the pleistocene: Using phylogeny and constraint to inform the evolutionary psychology of human mating. *Psychological Bulletin*, 135, 794-821.
- Eastwick, P. W., Eagly, A. H., Finkel, E. J., & Johnson, S. E. (2011). Implicit and explicit preferences for physical attractiveness in a romantic partner: A double dissociation in predictive validity. *Journal of Personality and Social Psychology*, 101, 993-1011.
- Eastwick, P. W., & Finkel, E. J. (2007). Northwestern speed-dating study II. *Unpublished Raw Data*.
- Eastwick, P. W., & Finkel, E. J. (2008). Sex differences in mate preferences revisited: do people know what they initially desire in a romantic partner? *Journal of Personality and Social Psychology*, 94, 245-264.
- Eastwick, P. W., Finkel, E. J., & Eagly, A. H. (2011). When and why do ideal partner preferences affect the process of initiating and maintaining romantic relationships? *Journal of Personality and Social Psychology*, 101, 1012-1032. doi: 10.1037/a0024062
- Eastwick, P. W., Luchies, L. B., Finkel, E. J., & Hunt, L. L. (2014). The predictive validity of ideal partner preferences: A review and meta-analysis. *Psychological Bulletin*, 140, 623-665.
- Eastwick, P. W., & Smith, L. K. (in press). Sex-differentiated effects of physical attractiveness on romantic desire: A highly powered, preregistered study in a photograph evaluation context. *Comprehensive Results in Social Psychology*.
- Eastwick, P. W., Smith, L. K., & Ledgerwood, A. (2018). How do people infer preferences for attributes? *Unpublished Manuscript, University of California, Davis, CA*.
- Evans, M. G. (1991). The problem of analyzing multiplicative composites: Interactions revisited. *American Psychologist*, 46, 6-15.

- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 23, pp. 75-109). San Diego, CA: Academic Press.
- Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude Strength: Antecedents and Consequences* (pp. 247-282). Mahwah, NJ: Erlbaum.
- Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition*, 25, 603-637.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, 69, 1013-1027.
- Fazio, R. H., Sherman, S. J., & Herr, P. M. (1982). The feature-positive effect in the self-perception process: Does not doing matter as much as doing?. *Journal of Personality and Social Psychology*, 42, 404-411.
- Fiedler, K. (2000). Illusory correlations: A simple associative algorithm provides a convergent account of seemingly divergent paradigms. *Review of General Psychology*, 4, 25-58.
- Fishbein, M. (1963). An investigation of the relationships between beliefs about an object and the attitude toward that object. *Human Relations*, 16, 233-239.
- Fishbein, M., & Ajzen, I. (1975). *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley.
- Fishbein, M., & Coombs, F. S. (1974). Basis for decision: An attitudinal analysis of voting behavior. *Journal of Applied Social Psychology*, 4, 95-124.

- Fiske, S. T., Neuberg, S. L., Beattie, A. E., & Milberg, S. J. (1987). Category-based and attribute-based reactions to others: Some informational conditions of stereotyping and individuating processes. *Journal of Experimental Social Psychology*, 23, 399-427.
- Fletcher, G. J. O., Simpson, J. A., & Thomas, G. (2000). Ideals, perceptions, and evaluations in early relationship development. *Journal of Personality and Social Psychology*, 79, 933-940.
- Fletcher, G. J., Simpson, J. A., Thomas, G., & Giles, L. (1999). Ideals in intimate relationships. *Journal of Personality and Social Psychology*, 76, 72-89.
- Fletcher, G. J. O., Tither, J. M., O'Loughlin, C., Friesen, M., & Overall, N. (2004). Warm and homely or cold and beautiful? Sex differences in trading off traits in mate selection. *Personality and Social Psychology Bulletin*, 30, 659-672.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132, 692-731.
- Gawronski, B., & Bodenhausen, G. V. (2011). The associative-propositional evaluation model: Theory, evidence, and open questions. In M. P. Zanna & J. M. Olson (Eds.), *Advances in Experimental Social Psychology* (Vol. 44, pp. 59–127). Amsterdam, The Netherlands: Elsevier.
- Glasman, L. R., & Albarracín, D. (2006). Forming attitudes that predict future behavior: a meta-analysis of the attitude-behavior relation. *Psychological Bulletin*, 132, 778-822.
- Goodwin, R., & Tang, D. (1991). Preferences for friends and close relationships partners: A cross-cultural comparison. *The Journal of Social Psychology*, 131, 579-581.

- Greenwald, A. G. (1989). Why attitudes are important: Defining attitude and attitude theory 20 years later. In A. R. Pratkanis, S. J. Breckler & A. G. Greenwald (Eds.), *Attitude Structure and Function* (pp. 429-440). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of Personality and Social Psychology*, 74, 1464-1480.
- Hainmueller, J., & Hopkins, D. J. (2015). The hidden american immigration consensus: A conjoint analysis of attitudes toward immigrants. *American Journal of Political Science*, 59, 529-548.
- Hammond, K. R. (1948). Measuring attitudes by error-choice: an indirect method. *The Journal of Abnormal and Social Psychology*, 43, 38-48.
- Heider, F. (1958). *The Psychology of Interpersonal Relations*. New York: Wiley.
- Heilman, M. E., & Saruwatari, L. R. (1979). When beauty is beastly: The effects of appearance and sex on evaluations of job applicants for managerial and nonmanagerial jobs. *Organizational Behavior and Human Performance*, 23, 360-372.
- Hill, R. (1945). Campus values in mate-selection. *Journal of Home Economics*, 37, 554-558.
- Hodson, G., Dovidio, J. F., & Gaertner, S. L. (2002). Processes in racial discrimination: Differential weighting of conflicting information. *Personality and Social Psychology Bulletin*, 28, 460-471.
- Hudson, J. W., & Henze, L. F. (1969). Campus values in mate selection: A replication. *Journal of Marriage and the Family*, 31, 772-775.

- Huntsinger, J. R. (2013). Anger enhances correspondence between implicit and explicit attitudes. *Emotion, 13*, 350-357.
- Joel, S., Teper, R., & MacDonald, G. (2014). People overestimate their willingness to reject potential romantic partners by overlooking their concern for other people. *Psychological science, 25*, 2233-2240.
- John, O. P., & Robins, R. W. (1994). Accuracy and bias in self-perception: individual differences in self-enhancement and the role of narcissism. *Journal of Personality and Social Psychology, 66*, 206-219.
- Kaplan, M. F. (1971). Dispositional effects and weight of information in impression formation. *Journal of Personality and Social Psychology, 18*, 279-284.
- Kaplan, M. F. (1973). Stimulus inconsistency and response dispositions in forming judgments of other persons. *Journal of Personality and Social Psychology, 25*, 58-64.
- Katz, D. (1960). The functional approach to the study of attitudes. *Public Opinion Quarterly, 24*, 163-204. doi:10.1086/266945
- Kelley, H. H. (1973). The processes of causal attribution. *American psychologist, 28*, 107-128.
- Kenrick, D. T., Groth, G. E., Trost, M. R., & Sadalla, E. K. (1993). Integrating evolutionary and social exchange perspectives on relationships: Effects of gender, self-appraisal, and Involvement level on mate selection criteria. *Journal of Personality and Social Psychology, 64*, 951-969.
- Kenrick, D. T., Sundie, J. M., Nicastle, L. D., & Stone, G. O. (2001). Can one ever be too wealthy or too chaste? Searching for nonlinearities in mate judgment. *Journal of Personality and Social Psychology, 80*, 462-471.

- Kille, D. R., Forest, A. L., & Wood, J. V. (2013). Tall, dark, and stable: Embodiment motivates mate selection preferences. *Psychological Science*, 24, 112-114.
- Lampel, A. K., & Anderson, N. H. (1968). Combining visual and verbal information in an impression-formation task. *Journal of Personality and Social Psychology*, 9, 1-6.
- Lawless, H. T., & Heymann, H. (2010). *Sensory Evaluation of Food: Principles and Practices*. Springer Science & Business Media.
- Ledgerwood, A. (2014). Evaluations in their social context: distance regulates consistency and context dependence. *Social and Personality Psychology Compass*, 8, 436-447.
- Ledgerwood, A., Trope, Y., & Chaiken, S. (2010). Flexibility now, consistency later: psychological distance and construal shape evaluative responding. *Journal of Personality and Social Psychology*, 99, 32-51.
- Li, N. P., Balley, J. M., Kenrick, D. T., & Linsenmeier, J. A. (2002). The necessities and luxuries of mate preferences: Testing the tradeoffs. *Journal of Personality and Social Psychology*, 82, 947-955.
- Lick, D. J., & Johnson, K. L. (2014). Recalibrating gender perception: Face aftereffects and the perceptual underpinnings of gender-related biases. *Journal of Experimental Psychology: General*, 143, 1259-1276.
- Lord, C. G., & Lepper, M. R. (1999). Attitude representation theory. In M. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 31, pp. 265–343). San Diego, CA: Academic Press.
- Lutz, R. J. (1975). Changing brand attitudes through modification of cognitive structure. *Journal of Consumer Research*, 1, 49-59.

- Møller, A. P. (1988). Female choice selects for male sexual tail ornaments in the monogamous swallow. *Nature*, 332, 640-642.
- Motyl, M., Iyer, R., Oishi, S., Trawalter, S., & Nosek, B. A. (2014). How ideological migration geographically segregates groups. *Journal of Experimental Social Psychology*, 51, 1-14.
- Nelson, L. D., & Morrison, E. L. (2005). The symptoms of resource scarcity judgments of food and finances influence preferences for potential partners. *Psychological Science*, 16, 167-173.
- Nisar, P. J., & Scott, H. J. (2011). Key attributes of a modern surgical trainer: perspectives from consultants and trainees in the United kingdom. *Journal of Surgical Education*, 68, 202-208.
- Norton, M. I., Vandello, J. A., & Darley, J. M. (2004). Casuistry and social category bias. *Journal of personality and social psychology*, 87, 817-831.
- Nosek, B. A., & Banaji, M. R. (2001). The go/no-go association task. *Social Cognition*, 19, 625-666.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, 31, 166-180
- Park, L. E., Young, A. F., & Eastwick, P. W. (2015). (Psychological) Distance Makes the Heart Grow Fonder: Effects of Psychological Distance and Relative Intelligence on Men's Attraction to Women. *Personality and Social Psychology Bulletin*, 41, 1459-1473.
- Patricelli, G. L., Uy, J. A. C., Walsh, G., & Borgia, G. (2002). Sexual selection: male displays adjusted to female's response. *Nature*, 415, 279-280.

- Pechmann, C., & Ratneshwar, S. (1992). Consumer covariation judgments: Theory or data driven? *Journal of Consumer Research*, 19, 373-386.
- Petty, R. E., & Cacioppo, J. T. (1984). The effects of involvement on responses to argument quantity and quality: Central and peripheral routes to persuasion. *Journal of personality and social psychology*, 46, 69-81.
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (pp. 123–205). New York: Academic Press.
- Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annual Review of Psychology*, 57, 199-226.
- Roskos-Ewoldsen, D. R., & Fazio, R. H. (1992). On the orienting value of attitudes: attitude accessibility as a determinant of an object's attraction of visual attention. *Journal of Personality and Social Psychology*, 62, 198-211.
- Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology*, 32, 880-892.
- Rothman, L., & Parker, M. J. (2009). Just about right (JAR) scales: Design, usage, benefits, and risks. ASTM International.
- Rydell, R. J., McConnell, A. R., Mackie, D. M., & Strain, L. M. (2006). Of two minds: Forming and changing valence-inconsistent implicit and explicit attitudes. *Psychological Science*, 17, 954-958.
- Salancik, G. R., & Conway, M. (1975). Attitude inferences from salient and relevant cognitive content about behavior. *Journal of Personality and Social Psychology*, 32, 829-840.

- Schaller, M., & O'Brien, M. (1992). "Intuitive Analysis of Covariance" and Group Stereotype Formation. *Personality and Social Psychology Bulletin*, 18, 776-785.
- Shaklee, H., & Mims, M. (1982). Sources of error in judging event covariations: Effects of memory demands. *Journal of Experimental Psychology: Learning, memory, and cognition*, 8, 208-224.
- Shestowsky, D. (2014). The psychology of procedural preference: How litigants evaluate legal procedures ex ante. *Iowa Law Review*, 99, 637-710.
- Sherman, J. W., Gawronski, B., Gonsalkorale, K., Hugenberg, K., Allen, T. J., & Groom, C. J. (2008). The self-regulation of automatic associations and behavioral impulses. *Psychological Review*, 115, 314.
- Smith, M. B., Bruner, J. S., & White, R. W. (1956). *Opinions and Personality*. New York: Wiley.
- Snyder, M., & Kendzierski, D. (1982). Acting on one's attitudes: Procedures for linking attitude and behavior. *Journal of Experimental Social Psychology*, 18, 165-183.
- Sprecher, S., & Regan, P. C. (2002). Liking some things (in some people) more than others: Partner preferences in romantic relationships and friendships. *Journal of Social and Personal Relationships*, 19, 463-481.
- Trope, Y., Ledgerwood, A., Liberman, N., & Fujita, K. (2018). Different levels of mental and social tools enable expansive versus contractive regulation. *Unpublished manuscript, New York University*.
- Thomsen, C. J., Borgida, E., & Lavine, H. (1995). The causes and consequences of personal involvement. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude Strength: Antecedents and Consequences* (pp. 191-214). Mahwah, NJ: Erlbaum.

- Thornhill, R. (1983). Cryptic female choice and its implications in the scorpionfly *harpobittacus nigriceps*. *The American Naturalist*, 122, 765-788.
- Turban, D. B., & Keon, T. L. (1993). Organizational attractiveness: An interactionist perspective. *Journal of Applied Psychology*, 78, 184-193.
- Uhlmann, E. L., & Cohen, G. L. (2005). Constructed criteria: Redefining merit to justify discrimination. *Psychological Science*, 16, 474-480.
- Uhlmann, E. L., & Cohen, G. L. (2007). "I think it, therefore it's true": Effects of self-perceived objectivity on hiring discrimination. *Organizational Behavior and Human Decision Processes*, 104, 207-223.
- Van Trijp, H. C., Punter, P. H., Mickartz, F., & Kruithof, L. (2007). The quest for the ideal product: Comparing different methods and approaches. *Food Quality and Preference*, 18, 729-740.
- Wegener, D. T., & Petty, R. E. (1998). The naive scientist revisited: Naive theories and social judgment. *Social Cognition*, 16, 1-7.
- Wicker, A. W. (1969). Attitudes versus actions: The relationship of verbal and overt behavioral responses to attitude objects. *Journal of Social Issues*, 25, 41-78.
- Wilson, T. D., Laser, P. S., & Stone, J. I. (1982). Judging the predictors of one's own mood: Accuracy and the use of shared theories. *Journal of Experimental Social Psychology*, 18, 537-556.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101-126.

- Wood, D., & Brumbaugh, C. C. (2009). Using revealed mate preferences to evaluate market force and differential preference explanations for mate selection. *Journal of Personality and Social Psychology*, 96, 1226-1244.
- Wood, D., & Furr, R. M. (2016). The correlates of similarity estimates are often misleadingly positive: The nature and scope of the problem, and some solutions. *Personality and Social Psychology Review*, 20, 79-99.
- Zanna, M. P., & Rempel, J. K. (1988). Attitudes: A new look at an old concept. In D. Bar-Tal & A. W. Kruglanski (Eds.), *The Social Psychology of Knowledge* (pp. 315-334). Cambridge, UK: Cambridge University Press.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151-175.

Figure Captions

Figure 1: An integrative framework for attitudes towards objects and attitudes toward attributes.

Figure 2: Three constructs: attitudes towards objects, summarized attribute preferences, and functional attribute preferences.

Note: Circles are latent behavioral constructs, O = object, E = evaluation, A = attribute (as a concept), a = attribute (as exhibited by an object), $\rightarrow$ = causes and/or predicts

Figure 3: Three people with identical attitudes towards “sweet wine” but different attitudes towards “sweetness” in wine.

Figure 4: Model 1: The process of translating functional into summarized preferences.

Note: See Figure 2 for definitions. [†] The complexity moderational pathway is supported by meta-analytic data (Figure 5); Eastwick, Smith, & Ledgerwood (2018).

Figure 5: Nine studies examining the correspondence between summarized and functional preferences.

Note: Functional preferences were all assessed as a within-person association of an attribute with evaluations across several targets. Eastwick et al. (2014) values for breakfast cereals examined the attribute *sweet*, DeBruine (2006) examined the attribute *masculinity*, Eastwick & Smith (in press) examined *attractiveness*, and the effect sizes for the remaining studies reflect average summarized-functional preference correlations across several different attributes. 95% confidence intervals are estimated based on the *N*.

Figure 6: Model 2: The process by which attitudes towards attributes influence evaluations of relevant objects.

Note: See Figure 2 for definitions. [†] The distance moderational pathway is supported by Eastwick, Finkel, et al. (2011, Studies 1 and 2); Huang, Eastwick, & Ledgerwood (2018).

Figure 7: Model 3: The process by which attitudes towards attributes influence situation selection.

Note: See Figure 2 for definitions. [†] The quantity effect is supported by Eastwick, Smith, & Ledgerwood (2018).

Figures

Figure 1

Target of
Evaluation

Object

Attribute

Attitude
Construct

**Attitudes
toward
Objects**
toaster

**Summarized
Preferences**
crispy

**Functional
Preferences**
crispy

Example
Measures

Rating scale
assessing liking
for *toaster*

Relative
reaction time to
positive vs.
negative words
after seeing a
toaster prime

Rating scale
assessing desire
for *crispy* in
breakfast foods

Relative
reaction time to
positive vs.
negative words
after seeing a
crispy prime

Within-person
association of
crispy with
liking for
specific
breakfast foods

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6

Figure 7