Journal of Experimental Social Psychology 85 (2019) 103837

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

How do people translate their experiences into abstract attribute preferences?

Paul W. Eastwick*, Leigh K. Smith, Alison Ledgerwood

University of California, Davis, United States of America

ARTICLE INFO

This paper has been recommended for acceptance by Tara MacDonald.

Keywords:
Preferences
Mating
Attitudes
Covariation detection
Attributes

ABSTRACT

In many literatures, scholars study *summarized attribute preferences*: overall evaluative summaries of an attribute (e.g., a person's liking for the attribute "attractive" in a mate). But we know little about how people form these ideas about their likes and dislikes in the first place, in part because of a dearth of paradigms that enable researchers to experimentally change people's attribute preferences. Drawing on theory and methods in covariation detection and social cognition, we developed a paradigm that examines how people infer summarized preferences for novel attributes from *functional attribute preferences*: the extent to which the attribute predicts an individual's evaluations across multiple targets (e.g., a person's tendency to positively evaluate mates who are more vs. less attractive). In three studies, participants encountered manipulated information about their own functional preference for a novel attribute in a set of targets. They then inferred a summarized preference for the attribute. Summarized preferences corresponded strongly to the functional preference manipulation when targets varied on only one attribute. But additional complexity (in the form of a second novel attribute) caused summarized and functional preferences to diverge, and biases emerged: Participants reported stronger summarized preferences for the attribute when the population of targets possessed more of the attribute on average (regardless of functional preference strength). We also documented some support for a standard-of-comparison mechanism to explain this inferential bias. These studies elucidate factors that may warp the translation process from people's experienced evaluative responses in the world to their overall, summary judgments about their attribute preferences.

1. Introduction

In the course of their everyday lives, people experience various preferences for attributes: Someone might evaluate potential romantic partners more positively to the extent that they are geeky, or experience curries as more negative to the extent that they are spicy. Such preferences have been termed *functional attribute preferences* (Ledgerwood, Eastwick, & Smith, 2018) or *drivers of liking* (Lawless & Heymann, 2010), and they reflect the extent to which a given attribute (e.g., spiciness) drives an individual's liking for a series of targets (e.g., curries) that vary in their level of the attribute.

But humans—perhaps uniquely among animals—can also construct and express an overall, summary evaluative judgment of an attribute, abstracted from any one specific target. In other words, people have *knowledge* about the attributes they like and dislike. A person might think about her overall penchant for geekiness in romantic partners, or exclaim to her friend: "Ugh, I hate spiciness in curries!" Such preferences have been termed *summarized attribute preferences* (Ledgerwood et al., 2018), and they reflect a person's overall evaluation of an attribute with respect to a given set of targets (e.g., how positively or negatively a person feels about the attribute *spiciness*).

It turns out that we know surprisingly little, as scientists, about how these two kinds of attitudes toward attributes are related to each other. When people think about or express a summarized attribute preference (e.g., how much they like geekiness in a partner), do they draw on a corresponding functional attribute preference (e.g., the extent to which geekiness has driven their liking of previously encountered partners), and if so, how? Logically, people should learn that they like the attributes that have actually driven their liking in the past, and many literatures assume that functional and summarized attribute preferences are linked (see Ledgerwood et al., 2018, for a review). Yet the process by which people might translate functional into summarized preferences remains opaque. One impediment to shedding light on this issue is that scholars have rarely been able to manipulate participants' attribute preferences; relative to attitudes toward objects, attitudes toward attributes have proven far more "difficult to alter" (Eagly &

* Corresponding author.
E-mail address: eastwick@ucdavis.edu (P.W. Eastwick).

https://doi.org/10.1016/j.jesp.2019.103837
Received 28 November 2018; Received in revised form 21 May 2019; Accepted 4 July 2019
0022-1031/ © 2019 Elsevier Inc. All rights reserved.

Chaiken, 1993, p. 237). If there is truth to Lewin's aphorism "If you want truly to understand something, try to change it" (Stam, 1996), then scholars do not understand attribute preferences particularly well.

The present manuscript examines—and elucidates factors that may warp or bias—the process by which people translate their experienced evaluative responses in the world (specifically, their functional attribute preferences) into overall, summary judgments (specifically, their summarized preferences). To this end, we generated an experimental framework that can illuminate the basic processes by which people form summarized attribute preferences for novel traits. We begin by identifying whether people can observe a functional preference for a continuously varying attribute and translate it into a summarized preference judgment—a domain-general form of inference that relies on people's ability to detect the covariation between targets' attributes and their valenced responses to those targets. Then, we consider whether certain factors might bias this functional-to-summarized inference process: For example, does the complexity of the task hinder this inferential process, and does complexity cause people to rely on heuristics or biases?

We ground our investigation of these issues in the extant literature on covariation detection (Alloy & Tabachnik, 1984; Klayman & Ha, 1987), which has examined a structurally similar question of how people make inferences about simple *binary* predictors and outcomes. We build on this literature to examine people's attitudes toward *continuous* attributes (see Chow, Colagiuri, & Livesey, 2019; Marsh & Ahn, 2009); this rare design feature allows us to ask new questions about biases that may emerge in contexts where some attributes are prominent relative to other attributes.

1.1. Summarized and functional attribute preferences

A recent theoretical paper distinguished between two ways of conceptualizing an attitude toward an attribute (Ledgerwood et al., 2018). *Summarized attribute preferences* are evaluative summaries (i.e., overall attitudes) of an attribute, trait, or quality—any dimension that can describe an attitude object to a greater or lesser extent. This construct appears in many literatures, including research on preferences for personality traits in other people (Anderson, 1968; Huang, Ledgerwood, & Eastwick, in press), preferences for attributes in a mate (e.g., Buss, 1989; Fletcher, Simpson, Thomas, & Giles, 1999; Hill, 1945), preferences for attributes in pets (Cohen and Todd, in press), preferences for features of workplaces (e.g., Kristof, 1996; Wood, Lowman, Harms, & Roberts, 2019), and the "value" component of classic expectancy-value models of attitudes (Fishbein & Ajzen, 1975). Commonly, summarized preferences are measured as a single numerical value on an explicit self-report rating scale (e.g., a participant might rate the desirability of *informality* in a workplace or *attractiveness* in a mate as an "8" on a 9-point scale from 1 = *not at all desirable* to 9 = *extremely desirable*; Buss, 1989; Wood et al., 2019).

Functional attribute preferences refer to the extent to which (a) the amount of an attribute, trait, or quality in each of a series of attitude objects (e.g., people, mates, pets, workplaces) drives (b) an individual's evaluation of each of those objects (Ledgerwood et al., 2018).[1] An individual would exhibit a strong functional preference for fruitiness in olive oils if the olive oils he likes tend to be fruitier than the olive oils he dislikes (Delgado & Guinard, 2011); an individual would exhibit a strong functional attribute preference for attractiveness if the partners she likes tended to be more attractive than the partners she dislikes (Eastwick & Finkel, 2008). Researchers who study nonhuman animals can *only* examine functional preferences, as animals (of course) cannot report summarized preferences. For instance, female satin bowerbirds are more attracted to males who exhibit more (vs. less) intense courtship displays (Patricelli, Uy, Walsh, & Borgia, 2002), and female swallows are more attracted to males with longer (vs. shorter) tail ornaments (Møller, 1988). Ledgerwood et al. (2018) outlines the conceptual difference between summarized and functional preferences and distinguishes it from other important demarcations in the attitude literature (e.g., indirect vs. direct measurement strategies, general vs. specific attitudes, and attitudes vs. behaviors).

People presumably experience functional attribute preferences across a wide variety of domains in the course of their everyday activities (e.g., when evaluating teammates that vary in intelligence; Caruso, Rahnev, & Banaji, 2009). But when people need to make a summarized attribute preference judgment (e.g., when describing their preferences to a friend), do they use their functional preferences to inform their summarized preferences, and if so, how? If someone has experienced greater positivity toward increasing intelligent teammates in the past, then we might expect him to translate this functional preference into a summarized preference such as "Intelligence is very important to me in a teammate." Indeed, humans often use past experiences to inform their predictions about the future (Schacter & Addis, 2007; Suddendorf & Corballis, 2007); in this way, summarized preferences might ultimately be useful in guiding subsequent decisions (Wang, Eastwick, & Ledgerwood, 2019).

Yet, the modest existing literature on the extent to which summarized attribute preferences reflect functional attribute preferences suggests that the process of inferring summarized from functional preferences might not be straightforward: The correspondence between the two seems to vary considerably depending on the nature of the judgment task. For example, participants' summarized preferences for attributes correspond moderately strongly with their functional preferences ($r = \sim 0.20$) in controlled contexts (e.g., evaluating photographs/descriptions of other people; Caruso et al., 2009; Eastwick & Smith, 2018; Wood & Brumbaugh, 2009). But in contexts where people have relatively immersive experiences with partners (e.g., face-to-face interactions), this correlation is $r < 0.05$ and not reliably significant (for a review, see Ledgerwood et al., 2018). This pattern of findings is consistent with the possibility that people are better able to translate their experienced evaluations into abstract judgments (i.e., functional preferences into summarized preferences) to the extent that the evaluative context is simple rather than complex. In the current manuscript, we investigate the possibility that complexity might hinder people's ability to infer a summarized preference from a functional preference. To do so, we create a paradigm that strips down a complicated interpersonal setting to its core components, enabling us to manipulate and examine the effect of complexity on the very basic social-cognitive process of translating functional preferences into summarized preferences.

1.2. Summarized preference formation as a covariation detection task

There is little research that directly addresses how people translate a functional attribute preference into a summarized preference (Ledgerwood et al., 2018), but several studies on covariation detection (also called contingency judgments) are conceptually analogous despite not typically involving evaluative judgments. *Covariation detection* refers to the process of determining the extent to which two variables (often called the *cue* and the *outcome*) are related to one another (Alloy & Tabachnik, 1984; Crocker, 1981, 1982; Klayman & Ha, 1987; Schaller & O'Brien, 1992; Trolier & Hamilton, 1986; Vadillo, Blanco, Yarritu, &

[1] This construct has been called a "revealed preference" in some prior work (Wood & Brumbaugh, 2009). We instead use the term "functional attribute preference" because, in the behavioral economics and judgment and decision-making literatures, the meaning of the term revealed preference is quite broad (Beshears, Choi, Laibson, & Madrian, 2008; Samuelson, 1948). A revealed preference may refer to any observable behavior, and studies that examine the correspondence between stated and revealed preferences are typically examining the attitude-behavior relationship—a well-studied topic in social psychology (Ajzen & Fishbein, 1977; Fazio, 1990). The focus of the present manuscript is the correspondence between two kinds of evaluations, not the attitude-behavior relationship.

Matute, 2016). Classic covariation detection examples include identifying whether bakery products will or will not rise depending whether yeast has or has not been added to dough (Shaklee & Mims, 1981) and whether a patient will or will not develop a disease depending on whether she has or has not displayed a particular symptom (Smedslund, 1963).

This process of covariation detection might mirror the process that takes place when people form and report their summarized preferences. That is, when prompted to disclose a summarized preference, people should presumably consult and consolidate their personal experiences with the relevant objects that possessed the attribute to varying degrees (Fazio, 1987; Fazio, Lenn, & Effrein, 1984). For example, when prompted with a preference rating scale for the item *intelligent*, a person may observe the extent to which the intelligence of several potential teammates covaries with her desire to have those teammates on her team (i.e., a functional attribute preference) and use this information to infer how much she values intelligence (i.e., a summarized attribute preference; Caruso et al., 2009).

Covariation detection is not, however, always accurate (Vadillo et al., 2016). For example, increasing the demands of the covariation task (e.g., via informational complexity, cognitive load, or demands on working memory) decreases participants' ability to estimate the correlation between the cue and the outcome (Arkes & Harkness, 1983; Shaklee & Mims, 1982; Shaklee & Tucker, 1980; Ward & Jenkins, 1965). Furthermore, of special relevance to the present studies, the addition of a third piece of information (above and beyond the cue and the outcome) can cause people to shift their covariation judgments away from the true contingency contained in the data: Sometimes participants perceive a covariation that is not actually present (as in the case of pseudocorrelations), and sometimes they overestimate or underestimate the strength of the true covariation (as in the case of overshadowing; Fiedler, Freytag, & Meiser, 2009; Price & Yates, 1993; Schaller & O'Brien, 1992; see also Kelley, 1973).

Thus, it seems possible that people may form summarized preferences by drawing on basic processes identified in the covariation detection literature—although to our knowledge, this possibility has never been empirically tested. On the one hand, the recently delineated distinction between functional and summarized preferences for attributes (Ledgerwood et al., 2018) certainly suggests that covariation detection approaches could be fruitfully applied to the domain of attitudes toward attributes, perhaps along with the moderating role of cognitive demands. On the other hand, most of the covariation detection literature examines binary and/or discrete cues and outcomes (Allan & Jenkins, 1983; Vadillo et al., 2016) rather than continuous attributes that vary as a dimension (for two exceptions, see Chow et al., 2019; Freytag, 2003), and research suggests that people have difficulty extracting summarized judgments of continuous data (Fisher & Keil, 2018). The studies reported in this manuscript are the first to systematically test whether basic principles of covariation detection can be applied to understanding the formation of summarized attribute preferences.

1.3. The current research

In the present research, we adapted the covariation detection paradigms described above in order to investigate the social-cognitive process of forming a summarized attribute preference—that is, how people learn about a summarized preference in the first place when they encounter a given attribute for the first time. We situated our investigation in scenarios that involved a simplified version of mate preferences—the literature in which summarized preferences have been studied most extensively—to provide a launchpad for future research in this domain. We conceptualized a *functional attribute preference* (our independent variable) as the covariation between the presence of a relevant attribute (i.e., the cue) in a group of targets and how positively each target was evaluated (whether the target was likeable or not; i.e., the outcome). We conceptualized a *summarized attribute preference* (our dependent variable) as an overall summary evaluation of the relevant attribute, just like in typical studies of mate preferences.

We began in Study 1 by testing whether participants were able to translate a functional attribute preference into a corresponding summarized attribute preference. Insofar as this process mirrors how people approach a classic covariation detection task, stronger functional preferences should lead to higher summarized preference judgments—at least when the task is simple and involves only a single trait. However, since real people are complex—they have multiple traits—we also included more complex conditions in which the targets had two (uncorrelated) traits to see if increasing the complexity of the stimuli would interfere with participants' performance (as proposed in Ledgerwood et al., 2018; see Model 1).

In Studies 2 and 3, we continued our investigation of how people infer summarized from functional attribute preferences under differing conditions of complexity by considering what kinds of factors might produce biases in such contexts. In particular, we zeroed in on one potentially important source of bias in this process in the mating domain: the overall amount of an attribute in a population of targets (i.e., whether a pool of potential mates has a low or high level of an attribute on average, as proposed in Ledgerwood et al., 2018; see Model 3). In these studies, we equated functional preferences across conditions but manipulated the amount of the attribute present in the population of targets to see if participants' summarized preferences would be affected nonetheless. Study 3 identified a standard-of-comparison mechanism that potentially explains the bias that participants exhibited in Study 2.

1.3.1. Common details across study designs

These studies used an imaginary setting—just like many classic and contemporary psychological studies that investigate basic social processes (e.g., Almaraz, Hugenberg, & Young, 2018; Callahan & Ledgerwood, 2016; Fiedler, Walther, Freytag, & Stryczek, 2002; Gauthier & Tarr, 1997)—to maintain high experimental control and to limit the likelihood that participants' prior expectations would make their summarized attribute preferences difficult to alter (see Eagly & Chaiken, 1993). Specifically, these studies used a novel adaptation of a paradigm developed in the attitudes literature to generate new attitudes, called Beanfest (Fazio, Eiser, & Shook, 2004; Fazio, Pietri, Rocklage, & Shook, 2015). The typical Beanfest paradigm explores how participants learn to distinguish which beans provide positive versus negative outcomes (typically gaining versus losing points) given the various attributes (e.g., roundness) of the beans. Essentially, our new adaptation of Beanfest substituted potential dating partners for beans and an imaginary attribute "*Melb*" for bean roundness. Participants encountered functional preference information by experiencing the covariation between (a) the amount of *Melb* 24 potential mates had and (b) whether going on a date with each potential mate resulted in positive outcomes (gaining 10 points) or negative outcomes (losing 10 points). Similar to the original Beanfest paradigm, the point allocation was designed to elicit attitude formation toward novel objects—in this case, more liking for the dates who possessed more (versus less) *Melb* (i.e., a functional preference for *Melb*). After experiencing this functional attribute preference information, participants then reported their own summarized preference for *Melb* (e.g., "How much do you value *Melb* in a romantic partner?") as the dependent measure.

For all studies, we report all conditions, all exclusions, and all dependent measures relevant to our a priori hypotheses, which are described below. (All measures in the study are available at https://osf.io/5kpbj).

1.3.2. Recruitment, screening, and sample size criteria

All participants were recruited from Amazon Mechanical Turk (MTurk) and paid $0.55 to complete an online survey that lasted approximately 10 min. MTurk workers were only eligible to participate if they had not completed a previous study in this line of work, if they

lived in the United States, and if they had a HIT approval rating of at least 95%. Toward the end of each study, participants answered an attention-check question: "To show us that you have been paying attention to the instructions, please select the 'Other' option below, instead of indicating your actual region of origin." Participants who failed the attention check were excluded from analyses.[2] For all studies, race/ethnicity, relationship status, and sexual orientation information is reported in the Supplemental materials (see Appendix B).

All research designs were between subjects. The three studies described in this manuscript built on a series of prior studies in which we examined how people form judgments about other people's summarized preferences (see Appendix C); all analyses for the studies described in the main text were planned a priori and precisely mimic the analyses we conducted for the three earlier studies reported in Appendix C. In the studies for which the planned analyses involved testing for 2 × 2 interactions, we aimed to recruit 100 participants per cell. Power analyses conducted in G*Power indicated that a cell size of $n = 100$ provided 80% power to detect a simple main effect size of $d = 0.40$ (approximately the average reported effect size in social-personality psychology; Richard, Bond, & Stokes-Zoota, 2003) with alpha set at 0.05. In the studies that tested 2 × 3 interactions, we recruited 200 participants per cell. The larger sample size of 200 participants per cell gave us 98% power to detect a simple main effect size of $d = 0.40$ with alpha level set at 0.05. We further maximized power by conducting a single-paper meta-analysis, which we present at the conclusion of Study 3.

2. Study 1

In Study 1, we experimentally manipulated the strength of the functional attribute preference (weak versus strong functional preference) as well as the complexity of the stimuli participants viewed (low versus high complexity). After participants experienced their functional attribute preference for a novel attribute *Melb*, they then reported their summarized preference for *Melb*. Drawing on the mating and covariation detection research described above, we hypothesized that the strength of the functional attribute preference would influence participants' judgments of the summarized preference for that attribute, but that this effect would be stronger when participants were evaluating low rather than high complexity stimuli.

2.1. Method

2.1.1. Participants

Participants were 405 workers recruited from MTurk. Forty-one participants who completed the survey were excluded (as planned a priori) from any subsequent analyses because they selected the incorrect response to the attention check item, making our final sample size 364 Mechanical Turk workers (41.5% male; aged 18–79, $M_{age} = 37.8$, $SD = 12.4$).

2.1.2. Procedure

Participants watched a 2-min, illustrated video containing the background and instructions for the study. Participants were told to imagine that they lived on another planet where people had many different powers.

2.1.2.1. Manipulating stimuli complexity. Participants in the low complexity condition learned that they lived on a planet where people had the ability to move objects with their minds. This ability was called *Melb* and people had varying levels of it. *Melb* was depicted as a glowing, red orb centered on a person's head. The more *Melb* a person had, the larger their red orb was. Participants in the high complexity condition learned that they lived on a planet where people had varying levels of *Melb* and varying levels of an additional trait called *Flobe*—the ability to float in the air. *Flobe* was depicted as a glowing, golden disk floating underneath an individual's feet. The more *Flobe* a person had, the larger the golden disk.

2.1.2.2. Covariation detection task. Participants then played a game called "DateFest" where the goal was to gain points by making rewarding dating decisions. DateFest paralleled the classic BeanFest (Fazio et al., 2004, 2015) paradigm used to assess attitude formation toward novel stimuli in a virtual world. Participants were told that they were going to a party where they would meet 24 different party guests (presented in a random order) and that they must decide whether they would "go on a date" or "not go on the date" with each one. Each party guest was represented by a stick figure with a trivial name (e.g., Person A, Person B; Fig. 1) and was depicted on a single slide. Participants were told that some of the dates they went on would be "good experiences that you'll be happy to have had" whereas others would be "bad experiences that you'll wish you hadn't had," and that "in order to gain points, you have to learn which is which." Participants were told that while playing the game that they should "try to get an idea of what makes a date good and what makes a date bad, as well as how much *Melb* (or *Melb* and *Flobe*) each person had." Stimuli properties for all studies are presented in the Supplemental materials, Appendix A.

2.1.2.3. Manipulating participants' functional preference strength. We manipulated participants' functional preference for *Melb* by varying the strength of the covariation between the amount of *Melb* each of the 24 potential mates had and whether choosing a given date led them to gain or lose points. Twelve party guests were associated with a positive valence; going on these dates earned participants 10 points. The remaining twelve party guests were associated with a negative valence, and going on these dates lost participants 10 points (Fig. 2).

In the *weak functional preference condition*, the party guests that caused participants to earn versus lose points had very similar average values of *Melb*: The average *Melb* of the good dates was 7 and the average *Melb* of the bad dates was 6 (i.e., *Melb* was a weak predictor of whether dates were rewarding; Fig. 3, top). In the *strong functional preference condition*, the guests that caused participants to gain versus lose points had very different average values of *Melb*: The average *Melb* of the good dates was 9 and the average *Melb* of the bad dates was 4 (i.e., *Melb* was a strong predictor of whether dates were rewarding; Fig. 3, bottom). The overall average *Melb* of all 24 potential mates was held constant across the weak functional preference and strong functional preference conditions (i.e., the average *Melb* was always 6.5). A successful manipulation check of functional preference strength would reveal that the difference in the *Melb* of participants' accepted/approached versus rejected/avoided dates is larger in the strong than the weak functional preference condition.

We did not manipulate the functional preference strength of *Flobe*. As stated above, *Flobe* was simply included as a manipulation of stimuli complexity. To ensure that *Flobe* was equally likeable across both the weak and strong *Melb* functional preference conditions, the good dates always had an average *Flobe* of 8 and the bad dates always had an average *Flobe* of 5. *Melb* and *Flobe* levels were chosen so that the two traits did not correlate with one another ($r = 0.03$) across targets.

If participants chose not to go on a date with a guest, they neither gained nor lost any points, although after reporting their decision, they learned whether they would have gained or lost points if they had gone on the date (thereby ensuring that all participants learned the functional preference information to a similar extent). Thus, as in the original version of Beanfest, the game was designed to motivate participants to explore and assess the stimuli in their virtual environment, leading them to form more positive evaluations of the novel objects that were associated with gaining (vs. losing) points in the process (Fazio

[2] Including the participants who provided an incorrect answer to the attention check revealed identical conclusions to the hypothesis tests reported below.

Low Complexity Condition

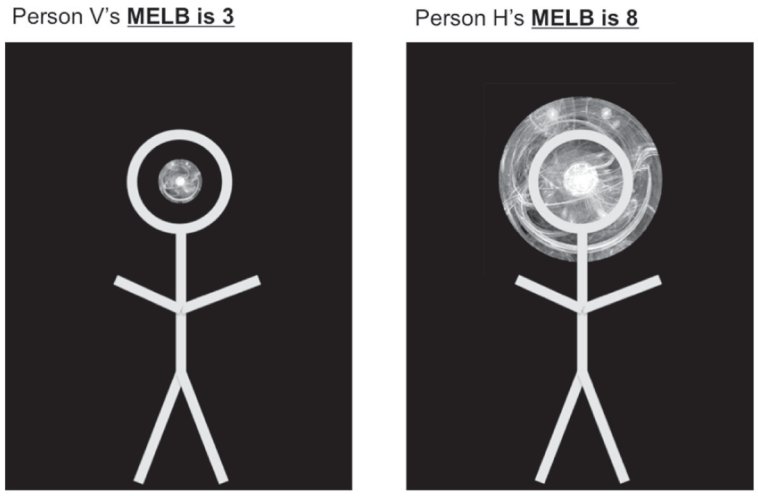

Note: Two examples of the 24 stimuli used in the low complexity (*Melb* only) condition.

High Complexity Condition

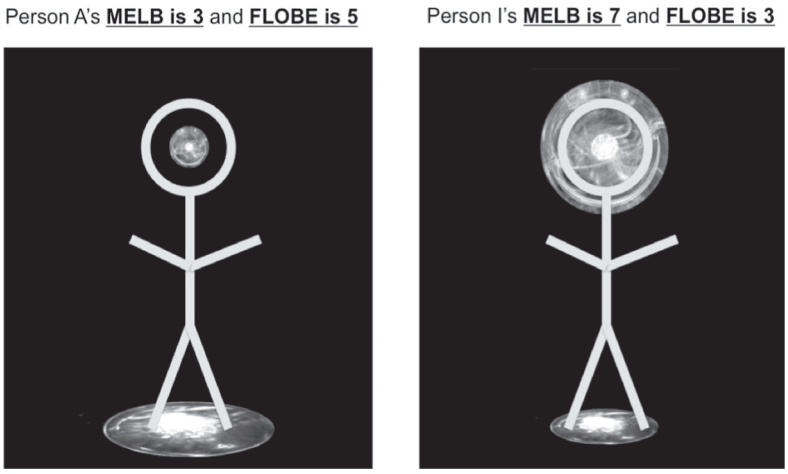

Note: Two examples of the 24 stimuli used in the high complexity (*Melb* and *Flobe*) condition.

Fig. 1. Example stimuli. Note: The top panel depicts two examples of the 24 stimuli used in the low complexity (*Melb* only) condition. The bBottom panel depicts two examples of the 24 stimuli used in the high complexity (*Melb* and *Flobe*) condition.

et al., 2015).

2.1.2.4. Participants' summarized preference for Melb. After playing DateFest, participants responded to the following four questions, which comprised the dependent measure: "How important is *Melb* to you in a romantic partner?", "How much do you value *Melb* in a romantic partner?," "How desirable is *Melb* to you in a romantic partner?," and "To what extent does *Melb* characterize your ideal romantic partner?" on scales from 1 (*not at all*) to 9 (*extremely*). These four items were highly reliable ($\alpha = 0.97$) and were thus averaged to form a scale reflecting participants' summarized preference for *Melb*.

2.2. Results

2.2.1. Manipulation check

In this study, we attempted to manipulate the strength of participants' own functional preferences via the points gained versus lost by various dates. To confirm that this manipulation successfully induced participants to experience a stronger functional preference for *Melb* in the strong (vs. weak) functional preference condition, we compared participants' functional preference for *Melb* (the difference in the *Melb* of participants' accepted/approached versus rejected/avoided dates) in the two functional preference conditions. As expected, participants' average functional preference for *Melb* was stronger in the strong functional preference condition than in the weak functional preference condition, $F(1, 360) = 648.55$, $p < .001$, *partial* $\eta^2 = 0.64$, confirming that our manipulation was successful.[3] Participants in the weak functional preference condition said "yes" to dates with an average *Melb* ($M = 6.80$, $SD = 0.47$) that was somewhat higher than the dates to

[3] We also checked to make sure that the manipulation was equally effective across complexity conditions, and it was: Complexity did not moderate the effect of the functional preference manipulation on participants' accept vs. reject decisions, $F(1, 358) = 0.01$, $p = .915$, *partial* $\eta^2 = 0.00$.

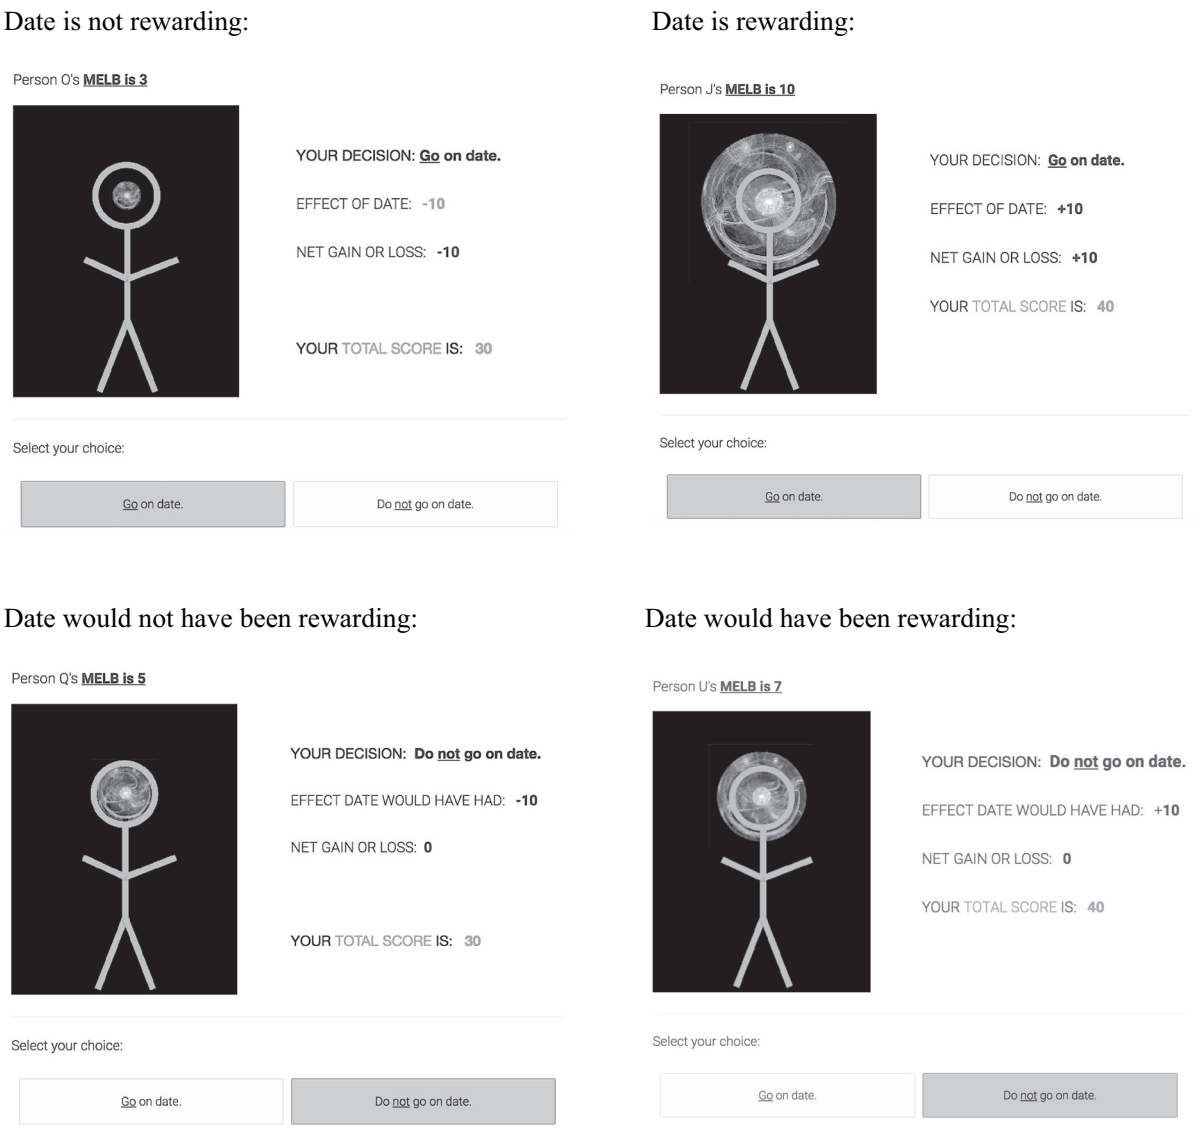

Fig. 2. Screenshots from Datefest.
Note: Participants made their decision about whether to go on the date or to not go on the date *prior* to seeing the effect the decision had on their score. Only the low complexity stimuli are presented here. High complexity stimuli included information about varying values with accompanying visual illustrations of *Flobe* (as seen in Fig. 1).

whom they said "no" ($M = 6.11$, $SD = 0.58$); in the strong functional preference condition, participants said "yes" to dates with an average *Melb* ($M = 8.48$, $SD = 0.84$) that was considerably higher than the dates to whom they said "no" ($M = 4.33$, $SD = 0.81$). In other words, our novel manipulation of functional preference strength strongly influenced participants' own functional preferences as indexed by their decisions to approach (accept) or avoid (reject) a potential date. (For details on participants' performance at the task, such as the number of points earned, see Supplemental materials Appendix G.)

2.2.2. Planned primary analyses

A 2 (functional preference strength: weak vs. strong) x 2 (stimuli complexity: low vs. high) between-subjects ANOVA revealed a significant main effect of functional preference strength, $F(1, 360) = 177.58$, $p < .001$, $\eta_p^2 = 0.33$, and a significant main effect for stimuli complexity, $F(1, 360) = 9.28$, $p = .002$, $\eta_p^2 = 0.03$. Importantly, the interaction between functional preference strength and the complexity of the stimuli was significant, $F(1, 360) = 19.23$, $p < .001$, $\eta_p^2 = 0.05$, indicating that complexity attenuated the effect of functional preference strength on summarized preference judgments (Fig. 4).

We conducted planned tests of simple main effects to further unpack this interaction. When complexity was low (*Melb* only), participants' summarized preference for *Melb* was substantially higher in the strong ($M = 7.85$, $SD = 1.65$) than the weak ($M = 4.48$, $SD = 2.33$) functional preference condition, $F(1, 360) = 171.73$, $p < .001$, $d = 1.69$. In contrast, when complexity was high (both *Melb* and *Flobe*), this effect was smaller but still significant: Participants' summarized preferences were higher in the strong ($M = 7.59$, $SD = 1.53$) than the weak ($M = 5.89$, $SD = 1.56$) functional preference condition, $F(1, 360) = 36.78$, $p < .001$, $d = 1.12$.

2.3. Discussion

This study demonstrated that participants can successfully infer their own summarized attribute preference from a corresponding functional attribute preference, and furthermore, that this ability is attenuated when evaluating more complex stimuli. Participants in the strong functional preference condition indicated that *Melb* was significantly more desirable to them than participants in the weak

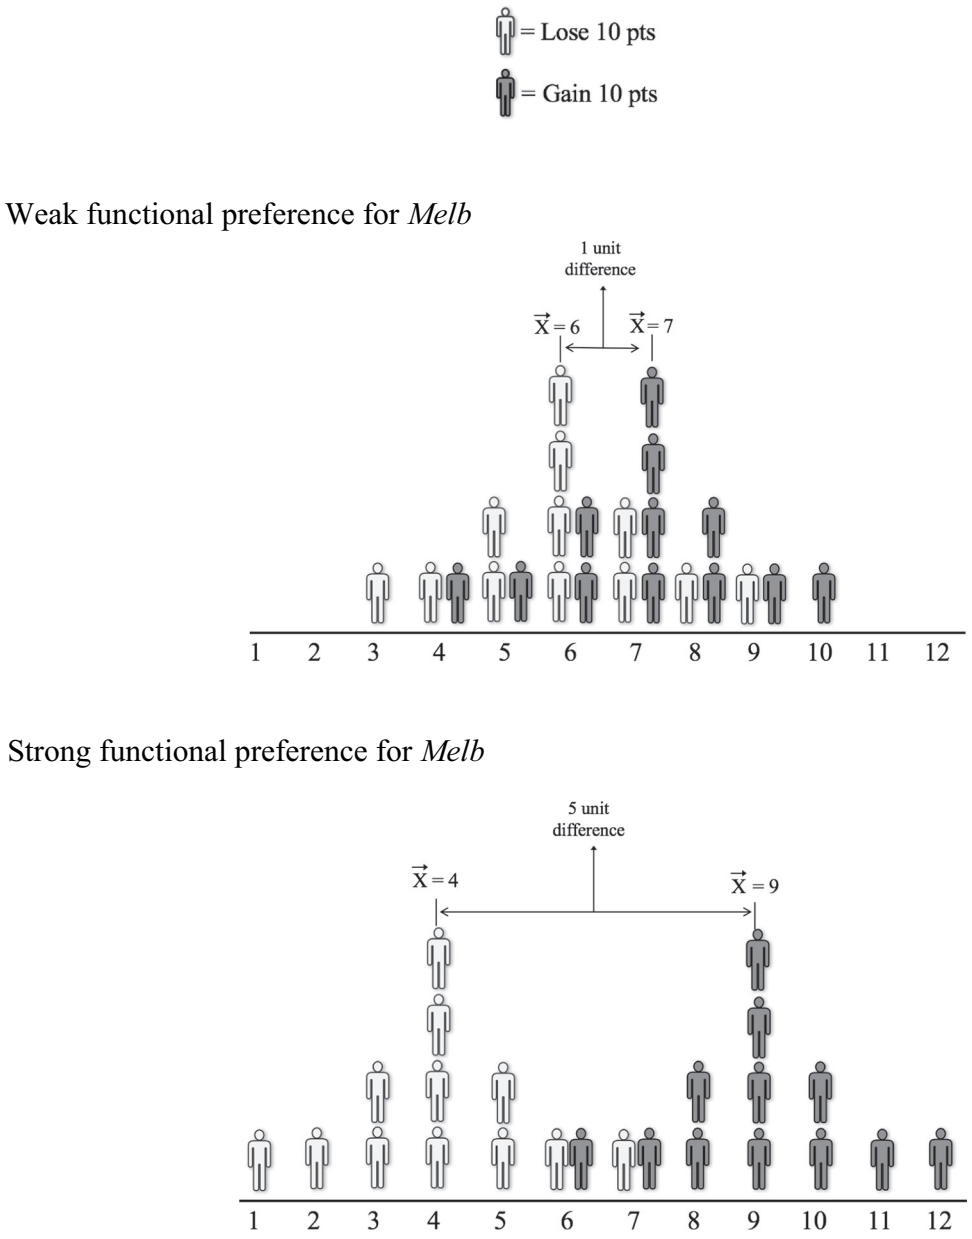

Fig. 3. Study 1 - functional preference strength distributions for *Melb*.
Note: The amount of *Melb* each person was assigned is on the horizontal axis. A small mean difference between the average *Melb* of the people pictured in white versus gray implies a weak functional preference for *Melb* (top). A large mean difference between the average *Melb* of the people pictured in white versus gray implies a strong functional preference for *Melb* (bottom).

functional preference condition, but the addition of *Flobe* reduced the size of this effect, perhaps because it interfered with the process of translating the functional preference for *Melb* into a summarized attribute preference. In this way, the addition of *Flobe* may have increased demands on working memory (Klayman & Ha, 1987; Pechmann & Ratneshwar, 1992; Schaller & O'Brien, 1992); it may have functioned analogously to mechanisms underlying the overshadowing effect (Price & Yates, 1993), a phenomenon in which the presence of one strong predictor makes it more challenging for participants to accurately learn the importance of an additional predictor.

This study is the first to experimentally demonstrate that people can infer summarized preferences from functional preferences, and that complexity can hinder the correspondence between the two. But of course, the real world is more complex still: After all, the stimuli in the current study varied on only two traits, whereas most attitude objects in real life—especially real people—vary on large number of attributes. Interestingly, we can observe a real-world analog to our experimental findings: Correlational data suggests that the association between functional and summarized preferences is larger in contexts involving photographs and descriptions (i.e., when people are considering relatively simple targets) than it is in face-to-face interactions (i.e., when people are considering relatively complex targets; Huang et al., in press; Ledgerwood et al., 2018). Thus, our Study 1 findings could represent one mechanism that explains why functional-summarized correspondence varies with the complexity of targets in the real-world.

3. Study 2

In Study 2, we investigated the possibility that the process of translating a functional into a summarized attribute preference might sometimes be biased by incidental factors. In the covariation detection literature, one such factor is the frequency with which participants

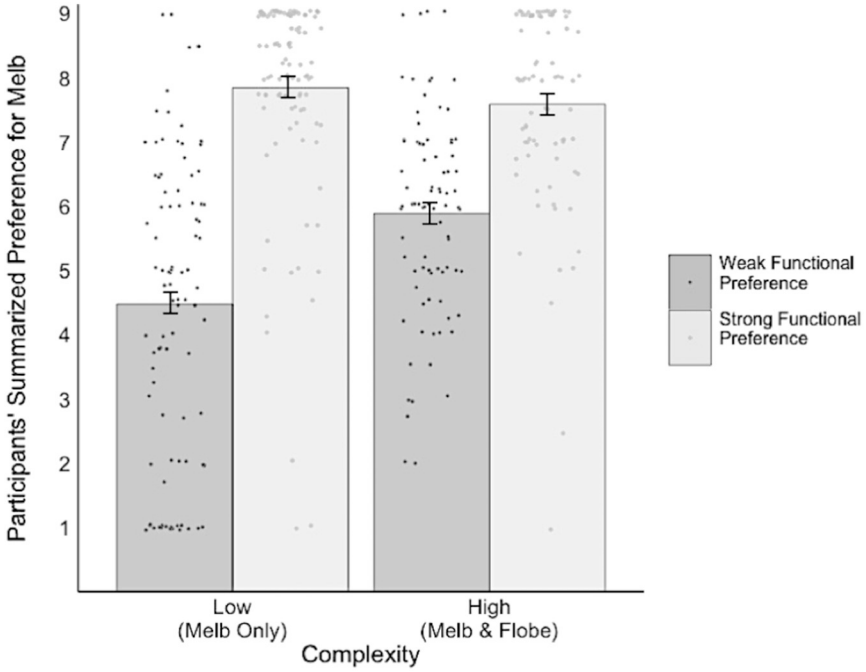

Fig. 4. Study 1 results.
Note: Error bars indicate one standard error above and below the mean.

encounter the covarying variables of interest. Specifically, the *cue-density bias* refers to the tendency for participants to judge the covariation between a cue and an outcome to be stronger when the cue appears frequently rather than infrequently (e.g., on 80% vs. 50% of trials; Allan & Jenkins, 1983; Perales, Catena, Shanks, & González, 2005; Vadillo, Musca, Blanco, & Matute, 2011; Vadillo et al., 2016). Although this bias has been studied in the context of categorical (rather than continuous) variables, we suspected that it would generalize to continuous attributes given that Study 1 provided initial evidence for the generalizability of basic covariation detection processes to summarized preference formation. Thus, we predicted that participants' summarized preferences would shift along with the average amount of an attribute in the population of targets. That is, participants should report stronger summarized preferences to the extent that the attribute is higher, on average, in the full population of (point-awarding and point-subtracting) potential partners.

In Study 2, we manipulated the average quantity (low versus high) of *Melb* present in a group of potential mates and again asked participants to make summarized preference judgments. We held the functional preference for *Melb* constant across the low and high quantity *Melb* conditions, thus ensuring that any mean difference in participants' summarized preference judgments between the low and high *Melb* conditions could be explained only by the quantity of *Melb* in the environment, not by the actual functional preference. In addition, we again manipulated stimuli complexity (low versus high); given the pattern of results we observed in Study 1, it seemed possible that the biasing effect of attribute quantity might be more likely to emerge under conditions of high complexity (i.e., when adding the attribute *Flobe*). In other words, adding an extraneous variable for participants to monitor might cause judgments of summarized preferences to be unduly influenced by incidental factors that do not actually reflect underlying functional preferences.

3.1. Method

3.1.1. Participants

Participants were 403 workers recruited from MTurk. Forty-five participants who completed the survey were excluded (as planned a priori) from any subsequent analyses because they selected the incorrect response to the attention check item, making our final sample size 358 Mechanical Turk workers (41.3% male; aged 18–77, M_{age} = 39.0, SD = 13.0).

3.1.2. Procedure

The procedure was identical to Study 1 except for the following change: Instead of manipulating functional preference strength, we manipulated attribute quantity by introducing participants to 24 targets who, on average, had either relatively *low quantities* of *Melb* (average *Melb* = 5.5) or *high quantities* of *Melb* (average *Melb* = 7.5; see Fig. 5). The functional preference for *Melb* was held constant across these two attribute quantity conditions, and was designed to be of moderate strength (i.e., the average *Melb* of the dates who earned points was always 3 units higher than the average *Melb* of the dates who subtracted points). This means that regardless of whether the average quantity of *Melb* in the population of targets was low or high, *Melb* levels predicted whether dates were rewarding to the same extent. After playing DateFest, participants answered the same four-item summarized preference dependent measure from Study 1 (α = 0.97).

3.2. Results

A planned 2 (attribute quantity: low vs. high) x 2 (stimuli complexity: low vs. high) between-subjects ANOVA indicated that there was no significant main effect of attribute quantity, $F(1, 354) = 0.26$, $p = .613$, $\eta_p^2 = 0.00$ nor stimuli complexity, $F(1, 354) = 0.004$, $p = .949$, $\eta_p^2 = 0.00$. Importantly, the analysis revealed an interaction between the quantity of *Melb* in the population and the complexity of the stimuli, $F(1, 354) = 11.57$, $p = .001$, $\eta_p^2 = 0.03$ (Fig. 6).

Planned simple main effects tests indicated that when complexity was high, participants' summarized preference for *Melb* was substantially higher in the high attribute quantity condition (M = 7.10, SD = 1.65) than in the low attribute quantity condition (M = 6.27, SD = 2.09), $F(1, 354) = 7.77$, $p = .006$, $d = 0.44$. In other words, when complexity was high, participants inferred stronger summarized preferences for *Melb* when the targets they encountered had higher (vs. lower) levels of this trait, despite the fact that functional preferences for

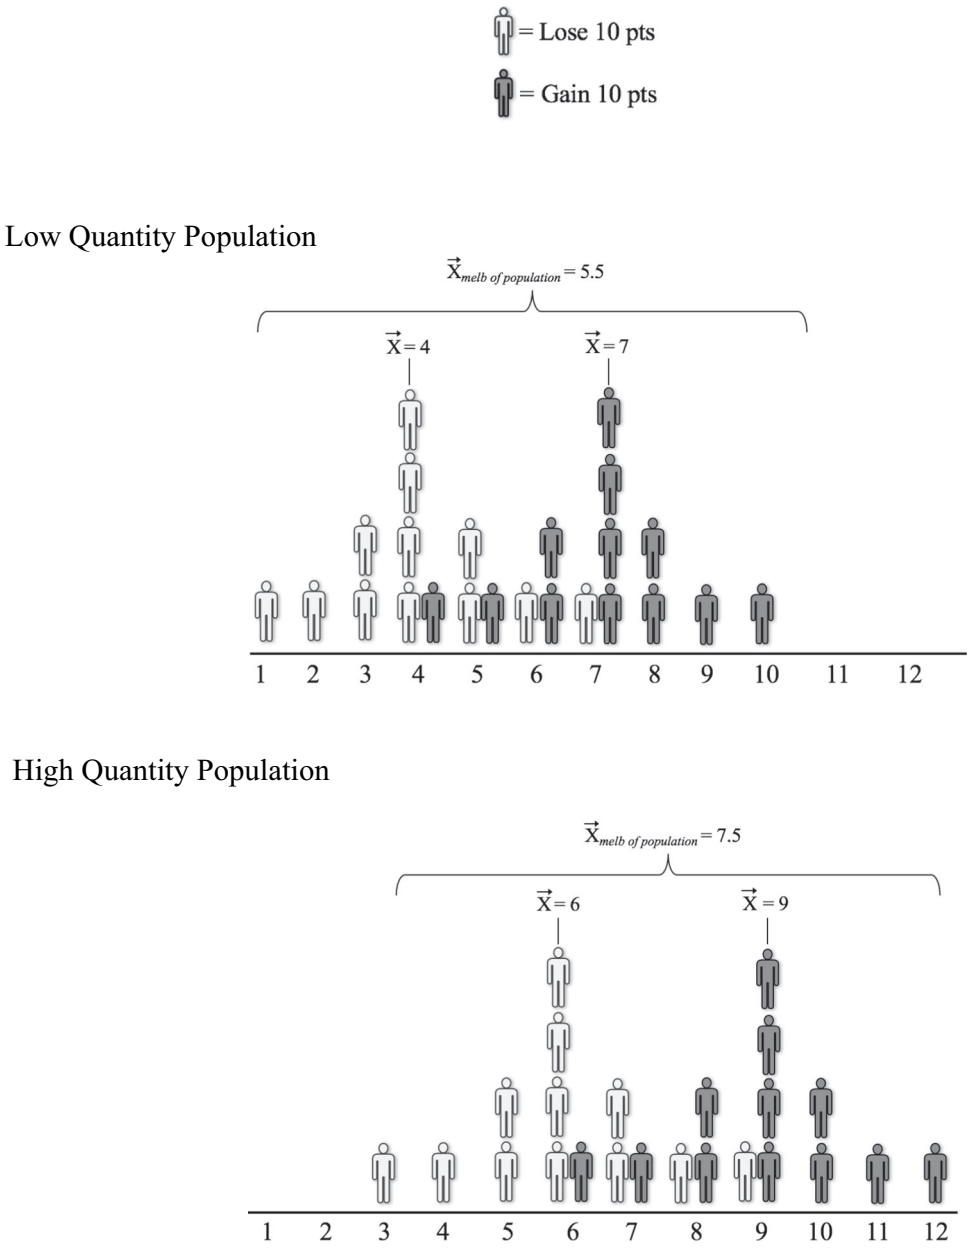

Fig. 5. Study 2 - attribute quantity distributions for *Melb*.
Note: The amount of *Melb* is on the horizontal axis. A low attribute quantity population is one in which the average *Melb* of the population is relatively low (top), and a high attribute quantity population is one is which the average *Melb* of the population is relatively high (bottom). The mean difference in the average *Melb* between the people associated with positive outcomes versus negative outcomes (the functional preference for *Melb*) is moderate (3 units) and identical in both the low and high quantity populations.

Melb were held constant. In contrast, when stimuli complexity was low, participants' summarized preference for *Melb* was slightly lower in the high (M = 6.39, SD = 2.45) than low (M = 7.00, SD = 1.79) attribute quantity conditions, $F(1, 354) = 4.12$, $p = .043$, $d = 0.28$. These findings suggest that stimuli complexity moderates the extent to which the quantity of an attribute in the population influences people's judgments of a summarized preference for that attribute.

3.3. Discussion

Study 2 suggested that the quantity of an attribute in a population of attitude objects can influence people's ability to translate a functional into summarized preference for that attribute. A pattern reminiscent of the cue-density bias (Vadillo et al., 2016) emerged in the high complexity condition: Participants inferred that they possessed a stronger summarized preference for an attribute when the mates in the environment possessed more (vs. less) of that attribute, even though functional preferences were identical across the two conditions. One mere additional trait may be sufficient to hinder people's ability to detect the extent to which an attribute predicts liking, which may in turn lead them to consider extraneous situational factors (e.g., the quantity of the attribute in the population) when inferring summarized from functional preferences.

In contrast, in the low complexity condition, we did not see a pattern reminiscent of the cue-density bias, suggesting that there may be some important differences between covariation detection with categorical variables and summarized preference formation based on continuous attributes. Indeed, in the present study, the effect of attribute quantity on summarized preferences was not only eliminated but even slightly reversed in the low complexity condition. Note, however, that

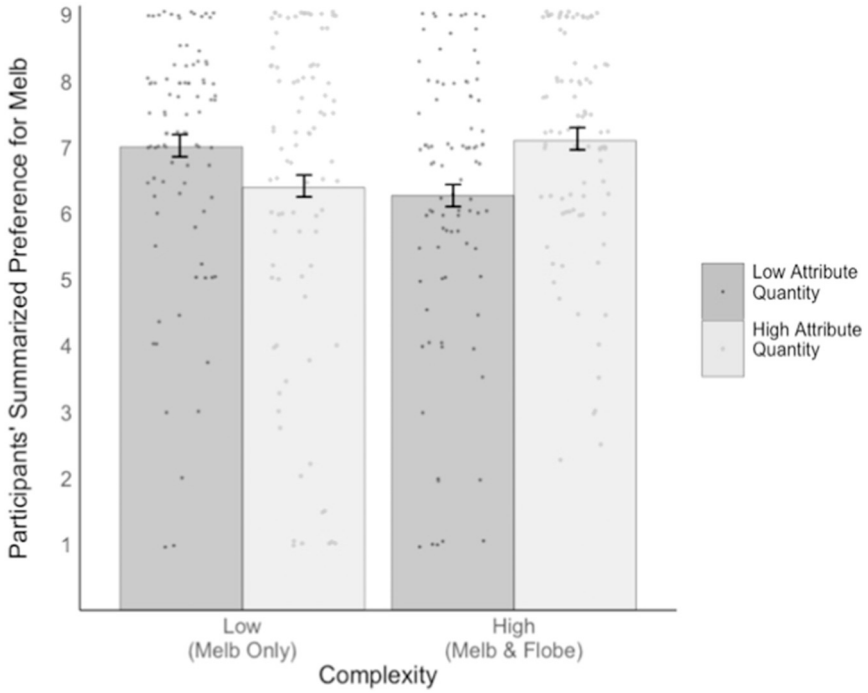

Fig. 6. Study 2 results.
Note: Error bars indicate one standard error above and below the mean.

in subsequent and supplemental studies, we see a null effect rather than a reversal when complexity is low; thus, we are most confident in the conclusion that the attribute quantity bias does not emerge under conditions of low complexity. Regardless, this finding highlights the importance of empirically testing whether and when key findings in covariation detection may or may not generalize to contexts involving continuous attributes (Chow et al., 2019).

4. Study 3

Given that Study 2 identified a novel interaction between complexity and attribute quantity, the next question becomes *why* complexity causes participants to erroneously use information about the amount of the attribute present in the population when making summarized preference judgments. What information do participants attend to as they learn about attribute preferences, and why does adding a second trait allow the average level of *Melb* to bias the process of inferring summarized preferences? We drew from classic studies in social cognition to posit that this pattern of results might arise from a standard-of-comparison mechanism, which refers to the tendency for people to spontaneously compare a focal target against a currently salient reference point or standard, and then to contrast the target of judgment away from the salient standard (Herr, Sherman, & Fazio, 1983; Moskowitz, 2004; Schwarz & Bless, 1992, 2007; Sherif & Hovland, 1961). In a classic demonstration of this effect, participants judged an ambiguous target Donald to be less hostile after thinking about a very hostile vs. nonhostile standard (e.g., Hitler vs. Santa Claus; Herr, 1986). Participants make these comparisons even when the standard should not be relevant to the task at hand, as illustrated by cases where participants' judgments are influenced by salient others' attitudes (Ledgerwood & Chaiken, 2007), labels on a response scale (Schwarz, 1999), and trait words flashed on a screen (Moskowitz & Skurnik, 1999).

When applied to the present studies, the standard-of-comparison logic suggests that the addition of the second trait *Flobe* in the high complexity condition may have provided participants with a reference standard against which they spontaneously compared *Melb*. In other words, people may tend to compare spontaneously the level of one attribute against other attributes when information about more than one attribute is available (e.g., "the good dates stood out as especially high on *Melb* relative to *Flobe*, therefore I really value *Melb* in a partner").

In Study 2, the level of *Flobe* was kept constant across quantity conditions to parallel a (common) real-world circumstance in which two groups differs on one trait but not a second, uncorrelated trait. Of course, this feature of our study design meant that in the low quantity condition, the average amount of *Melb* was always less than the average amount of *Flobe*, and in the high quantity condition, the average amount of *Melb* was always greater than the average amount of *Flobe* (Fig. 7, Unequal Flobe Condition). If participants spontaneously compared *Melb* with *Flobe* across targets, *Melb* would have seemed particularly high (relative to *Flobe*) in the high (vs. low) quantity condition, and thus *Melb* may have become particularly salient as a trait guiding evaluations of potential mates. In other words, participants may have perceived the positively valenced targets in the high quantity condition as having particularly high levels of *Melb* because they were spontaneously using *Flobe* as a standard.

To test this account, Study 3 investigated whether participants use the amount of *Flobe* present in the environment as a reference standard when inferring summarized preferences for *Melb*. Specifically, we created an additional pair of conditions in which the average levels of *Melb* equaled the average levels of *Flobe* in the population (Fig. 7, Equal Flobe Condition), so that *Melb* would no longer seem especially high (in the high attribute quantity condition) or low (in the low attribute quantity condition) relative to *Flobe*. If participants were using *Flobe* as a referent in the prior studies, then participants' summarized preference reports should not be affected by the amount of *Melb* present in the population in the Equal Flobe conditions. The design of Study 3 thus effectively disentangles our manipulation of absolute *Melb* quantity (high vs. low) from a manipulation of relative *Melb* compared to *Flobe*, allowing us to test whether it is absolute or relative levels of *Melb* that bias summarized preference judgments. Study 3 also encompasses a full, direct replication of Study 2, allowing us to assess the replicability of those findings across studies.

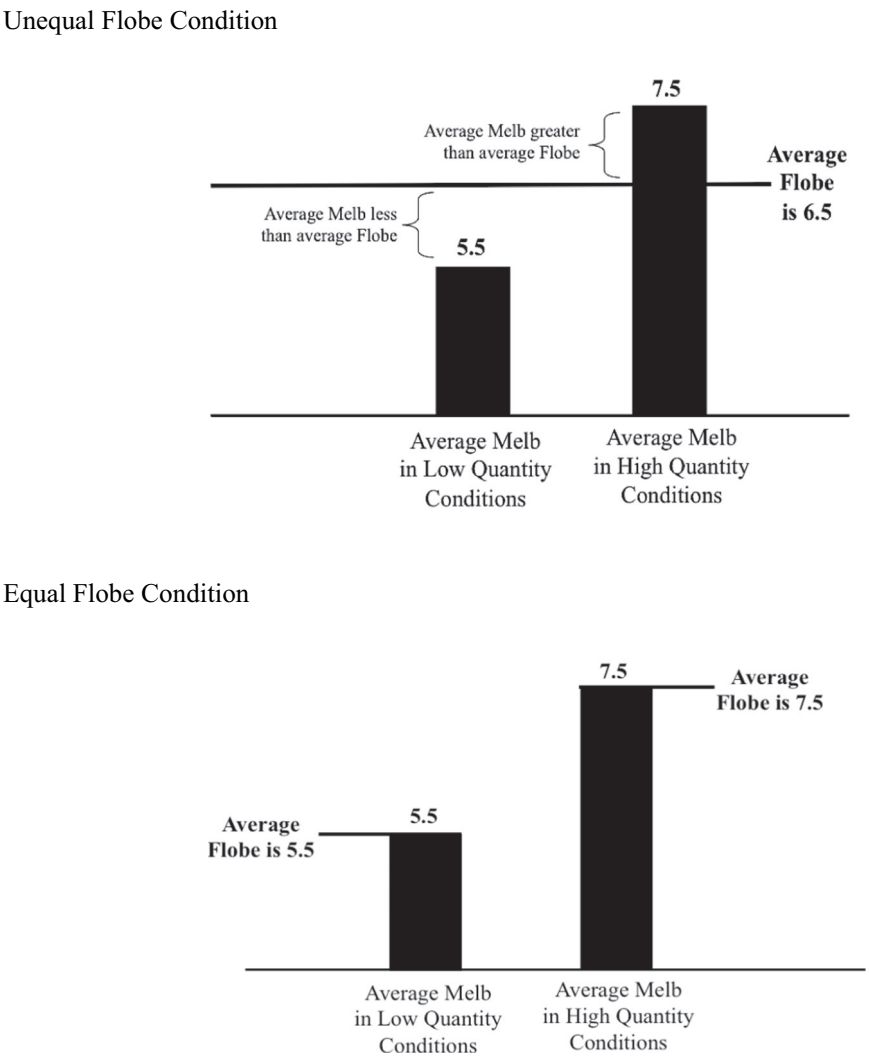

Fig. 7. Study 3 - graphical depiction of *Melb* vs. *Flobe* of stimuli.
Note: In the unequal *Flobe* condition (top), the average amount of *Melb* in the population is not equal to the average about of *Flobe* in the population (as in the Study 2, high complexity conditions). In the equal *Flobe* condition (bottom), we shift the average amount of *Flobe* in the population to be equal to the average amount of *Melb* in the population.

4.1. Method

4.1.1. Participants

Participants were 1227 workers recruited from MTurk. Following our a priori exclusion criteria, one hundred and twenty-six participants who completed the survey were excluded for failing the attention check, making our final sample size 1101 (40.0% male; aged 18–88, M_{age} = 37.3, SD = 12.7).

4.1.2. Procedure

Study 3 involved a 2 (attribute quantity: low vs. high) x 3 (reference standard: *no Flobe* vs. unequal *Flobe* vs. equal *Flobe*) between-subjects design. The Datefest procedure was largely identical to Studies 1 and 2: Participants were asked to imagine that they lived on a planet where people had many different powers and then chose whether or not to go on dates with 24 potential mates. The attribute quantity factor was manipulated just as in Study 2, so that the average amount of *Melb* in the pool of potential mates was either low or high.

For the reference standard factor, we manipulated the extent to which participants were able to use the amount of *Flobe* in the population as a standard against which to judge the relative value of *Melb*. This factor had three conditions: (a) the *no Flobe* condition, (b) the *unequal Flobe* condition, and (c) the *equal Flobe* condition. The *no Flobe* condition was identical to the low complexity condition in Study 2 wherein potential mates had, on average, relatively low average amounts of *Melb* (average = 5.5) or relatively high average amounts of *Melb* (average = 7.5). Just as in previous studies, no information regarding *Flobe* was provided in this condition. The *unequal Flobe* condition was identical to the high complexity condition in Study 2. The average *Melb* of potential mates was either less than their average amount of *Flobe* (low attribute quantity condition: *Melb* is 5.5 and *Flobe* is 6.5) or greater than their average amount of *Flobe* (high attribute quantity condition: *Melb* is 7.5 and *Flobe* is 6.5; see Fig. 7, Unequal Flobe Condition). In the *equal Flobe* condition, we adjusted the average amount of *Flobe* to be equal to the average amount of *Melb*. In the low attribute quantity condition, the average *Melb* of potential mates was 5.5, so the average amount of *Flobe* was adjusted to be 5.5 as well (we accomplished this by subtracting 1 unit of *Flobe* from each of the 24 potential mates that appeared in the *unequal Flobe* condition). In the high attribute quantity condition, the average *Melb* of potential mates was 7.5, so the average amount of *Flobe* was adjusted to be 7.5 as well (we accomplished this by adding 1 unit of *Flobe* to each of the 24 potential mates; see Fig. 7, Equal Flobe Condition).

After playing DateFest, participants completed the same four-item

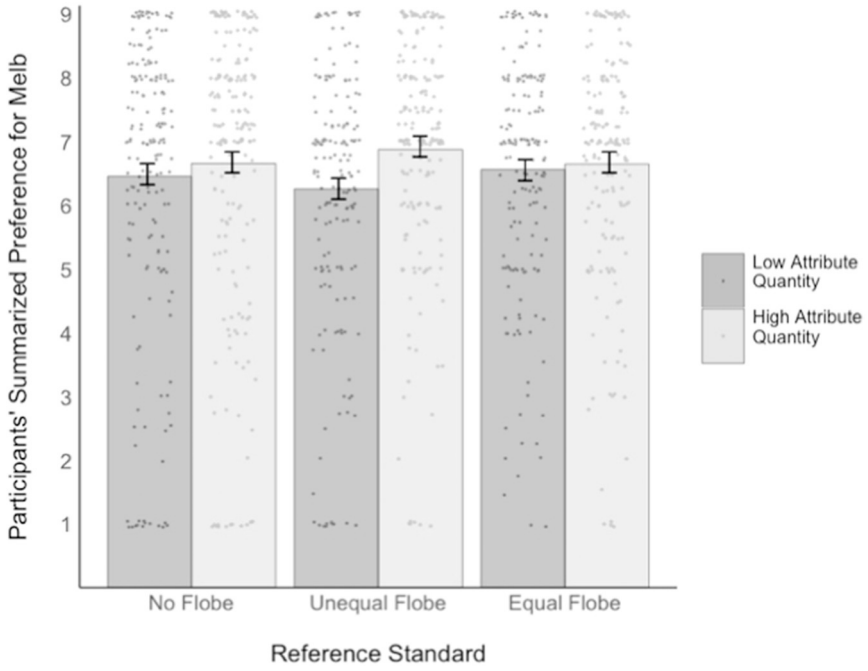

Fig. 8. Study 3 results.
Note: Error bars indicate one standard error above and below the mean.

summarized preference measure as in Studies 1 and 2 ($\alpha = 0.97$).

4.2. Results

A planned 2 (attribute quantity: low vs. high *Melb*) x 3 (reference standard: *no Flobe* vs. unequal *Flobe* vs. equal *Flobe*) factorial ANOVA revealed a significant main effect of attribute quantity, $F(1, 1095) = 6.12$, $p = .014$, $\eta_p^2 = 0.006$ and no main effect of reference standard, $F(2, 1095) = 0.056$, $p = .946$, $\eta_p^2 = 0.00$. The attribute quantity × reference standard interaction did not reach significance, $F(2, 1095) = 1.79$, $p = .167$, $\eta_p^2 = 0.003$, although the pattern of results looked strikingly similar to those observed in Appendix C, Supplemental Study 3 (see Fig. 8). Given that the significance test for this interaction was likely underpowered,[4] we continued on to conduct our planned simple main effects analyses, focusing especially on the effect size estimates in each condition.

Planned simple main effects tests revealed that (replicating Study 2), participants in the *unequal Flobe* condition expressed substantially stronger summarized preferences when the quantity of *Melb* was high ($M = 6.88$, $SD = 1.83$) rather than low ($M = 6.26$, $SD = 2.06$), $F(1, 1095) = 8.76$, $p = .003$, $d = 0.32$. In the *no Flobe* condition, this effect disappeared: participants' summarized preferences did not differ between the high ($M = 6.66$, $SD = 2.23$) and low ($M = 6.46$, $SD = 2.27$) attribute quantity conditions, $F(1, 1095) = 0.918$, $p = .338$, $d = 0.08$. Of special importance to the present study, in the *equal Flobe* condition, the effect of attribute quantity on summarized preference judgments also disappeared: Participants' summarized preferences did not differ between the high ($M = 6.65$, $SD = 1.86$) and low ($M = 6.56$, $SD = 1.84$) attribute quantity conditions, $F(1, 1095) = 0.152$, $p = .697$, $d = 0.05$. That is, when *Flobe* could not be used as a reference standard by which to judge the relative value of *Melb*, participants' summarized preferences for *Melb* were not biased by the absolute quantity of *Melb* in the population of potential mates.

4.3. Discussion

The results of Study 3 provide insight into the mechanism underlying the pattern of bias revealed in Study 2. Specifically, the biasing effect of attribute level on summarized preferences observed in Study 2 (in the high complexity condition) *only* emerged in the *unequal Flobe* condition, when the relative level of *Melb* was high or low compared to *Flobe*. When *Flobe* could no longer be used as a standard of comparison against which to judge *Melb* as high or low, the biasing effect of attribute level disappeared. In other words, participants inferred stronger summarized preferences for *Melb* when there was more *Melb* in the overall population relative to another, less abundant trait (*Flobe*). To our knowledge, this mechanism has not been identified previously in the covariation detection literature, perhaps because this literature has tended to examine binary rather than continuous attributes as cues. That is, the possibility that a cue may become more salient as a driver of an outcome to the extent that it is high *relative* to another cue may apply uniquely to the study of continuous cues like attributes, highlighting the new insights that can be gleaned by integrating the covariation detection literature with the study of summarized attribute preferences.

5. Supplemental studies on others' attribute preferences

To the extent that the findings documented here reflect the basic social-cognitive processes we have proposed above, we would expect them to apply similarly across people's inferences about others' preferences for attributes as well as their own preferences for attributes. In other words, if summarized preferences can be inferred from functional preferences through a combination of covariation detection and standard of comparison judgments, then we would expect people to show a similar pattern of responding when making inferences about another person's summarized preferences (e.g., how much does Casey like *Melb* in a partner?) as they do when making inferences about their own summarized preferences. Indeed, we found very similar results in a series of studies that asked participants to make inferences about another person's preferences after observing the extent to which *Melb* and

[4] We recognize that our study design may have been underpowered to detect the two-way interaction here, given that attenuation interactions require very large sample sizes to detect (Simonsohn, 2014). We address this issue by conducting a within-paper meta-analysis across all of the quantity manipulation studies we conducted in the next section.

Flobe were associated with that person's liking for potential partners (see Supplemental materials Appendices C–F for full details). These parallel sets of findings evoke the basic self-perception theory principle that the types of observational and inferential processes that inform how people come to understand themselves similarly inform how they come to understand other people (Bem, 1967, 1972; Fazio, 1987). Thus, taken together, the results of these studies may delineate a domain-general inference process that people use to construct judgments about their own and others' attribute preferences.

6. The big picture: meta-analyses of the *Melb* quantity effect across studies

The programmatic nature of our experiments allowed us to shed light on basic psychological mechanisms while also incorporating direct replications of several critical effects. Following current best-practice recommendations for multi-study papers, we conducted single-paper meta-analyses of all conducted studies that tested the relevant effects to get a better sense of the overall story conveyed by these data. Assuming that all studies are included, single-paper meta-analyses help researchers avoid being misled by natural sampling fluctuations across studies (Braver, Thoemmes, & Rosenthal, 2014; Goh, Hall, & Rosenthal, 2016; McShane & Böckenholt, 2017; Lakens & Etz, 2017; Ledgerwood, Soderberg, & Sparks, 2017; cf. Vosgerau, Simonsohn, Nelson, & Simmons, 2018).

First, in order to provide a cumulative picture of the effect of *Melb* quantity on judgments of participants' own summarized preferences, we meta-analyzed the effect of *Melb* quantity from all the relevant "own-preference" (i.e., Datefest) studies that we conducted (i.e., Study 2 and Study 3). The effect sizes were amalgamated to parallel the organization of Study 3 (i.e., a no *Flobe* condition, an unequal *Flobe* condition, and an equal *Flobe* condition). These three cumulative effect size estimates (d) are presented in Fig. 9. The effect of quantity in the no *Flobe* condition across studies was very small with a 95% CI overlapping with zero ($d = -0.04$), the effect of quantity in the unequal *Flobe* condition was medium-sized ($d = 0.36$), and the effect of quantity in the equal *Flobe* condition was very small ($d = 0.05$) with a 95% CI overlapping with zero. A test of the meta-analytic 2 × 3 interaction (i.e., testing whether these three effect sizes significantly differed from each other) was significant, $\chi^2 = 7.49$, $p = .024$ (B. McShane, personal communication, March, 27, 2017).

Second, in order to provide a cumulative picture of the effect of *Melb* quantity on judgments of others' summarized preferences (the supplemental studies described in Appendices C–F), we meta-analyzed the effect of *Melb* quantity from all the relevant "other-preference" studies that we conducted (i.e., Supplemental Studies 2a, 2b, 2c, 3, and the liking conditions of Supplemental Study 4). These three effect sizes are also presented in Fig. 9. The effect of quantity in the no *Flobe* condition across studies was small with a 95% CI nearly overlapping with zero ($d = 0.16$), the effect of quantity in the unequal *Flobe* condition was medium-sized ($d = 0.49$), and the effect of quantity in the equal *Flobe* condition was very small with a 95% CI overlapping with zero ($d = 0.06$). A test of the meta-analytic 2 × 3 interaction was also significant, $\chi^2 = 10.04$, $p = .007$.

These meta-analyses—which included every manipulation of attribute quantity that we ever conducted using the Datefest (i.e., own-preference) and the other-preference paradigms—bolster two conclusions. First, the average level of *Melb* in a pool of mates clearly biased summarized preference judgments, but only when there was a second quantity *Flobe* that participants could use as a reference standard. Second, this conclusion was true for both the own-preference and other-preference studies; these effects appear to reflect a domain-general mechanism that applies to the way people form inferences about preferences for attributes.

7. General discussion

People experience and think about attribute preferences in virtually all domains of life, but the basic inference processes that inform how people come to know these preferences in the first place are not well understood. The present research sheds important new light on how people form summarized preference judgments from functional preferences—the first experimental paradigm to systematically manipulate preferences for attributes. Participants in these studies (a) experienced the extent to which an attribute predicted positive responses for a series of targets (i.e., a functional attribute preference) and then (b) generated an overall, summary judgment of the preference for that attribute (i.e., a summarized preference). Following in the tradition of classic studies

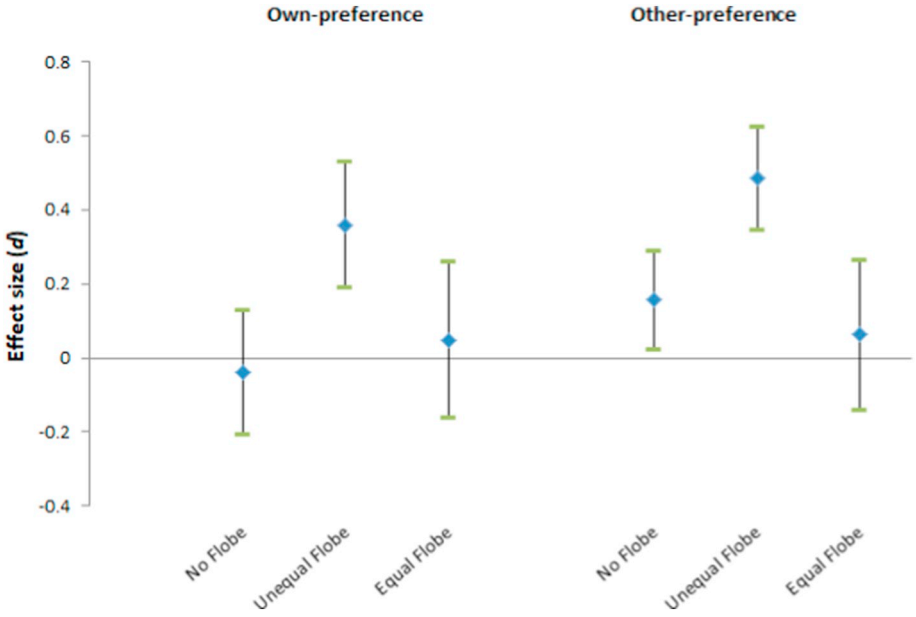

Fig. 9. Meta-analytic effect of quantity manipulation on summarized preferences.
Note: Bars depict 95% confidence interval. Own-preference studies are reported in Studies 1, 2, and 3 of the current article. Other-preference studies are reported in the Supplemental materials.

that used imaginary stimuli to illustrate how people initially form stereotypes (e.g., the illusory correlation; Hamilton & Gifford, 1976), the present studies used imaginary attributes to illuminate how people may initially infer attribute preferences.

Results revealed that, under certain circumstances (i.e., when participants only had to keep track of a single trait, *Melb*), participants' summarized attribute preferences strongly tracked their functional preferences. However, in Study 1, the inclusion of merely one additional attribute (*Flobe*) above and beyond the focal attribute (*Melb*) added sufficient complexity that summarized and functional preferences began to diverge. This finding provides an intriguing possible explanation for why the correspondence between functional and summarized preferences may be stronger in studies that ask participants to evaluate photographs and descriptions of other people (relatively simple stimuli) rather than live targets in face-to-face interactions (relatively complex stimuli; for reviews, see Eastwick, Finkel, & Simpson, 2019; Eastwick, Luchies, Finkel, & Hunt, 2014; Ledgerwood et al., 2018).

In Study 2, we saw that in the wake of this divergence between summarized and functional preferences, a bias reminiscent of the cue-density bias in the covariation detection literature emerged. Specifically, the average quantity of a trait in a class of targets influenced summarized preference judgments, independently from the underlying functional preference: Participants in the high complexity condition inferred stronger summarized preferences for *Melb* in a target when the *population* of targets had a high rather than low amount of *Melb* on average. Study 3 suggested that a standard-of-comparison mechanism provides a plausible explanation for this biasing effect of attribute quantity (Schwarz & Bless, 2007). Specifically, when participants were tracking multiple traits in a population of targets, they seemed to spontaneously use one trait as a reference standard against which to compare the relative importance or weight of another, orthogonal trait. That is, participants inferred stronger summarized preferences for *Melb* more in a partner when the population of partners had relatively more *Melb* relative to *Flobe* compared to when the population had less *Melb* relative to Flobe.

7.1. Implications for research on covariation detection and preferences

The present article primarily drew from covariation detection perspectives to illuminate the formation of attitudes toward novel attributes. At the same time, however, our results can also offer new contributions to the covariation detection literature. First, our meta-analytic results suggested that the biasing effect of attribute quantity (akin to the cue-density bias in the covariation detection literature; Vadillo et al., 2016) disappeared when participants needed to track only a single attribute, whereas the cue-density bias can emerge for categorical cues even under such simple conditions (e.g., Perales et al., 2005; Vadillo et al., 2011). This finding therefore suggests that there may be important but understudied differences in how people make inferences from categorical versus continuous cues. Second, to our knowledge, the comparative mechanism documented in Studies 2 and 3 has not been studied with respect to the cue-density bias in covariation detection. It is therefore possible that, in a symptom-disease covariation detection task (Smedslund, 1963) for example, participants may overestimate the diagnosticity of a continuous symptom (e.g., severity of a rash) to the extent that it is more salient than a second continuous symptom (e.g., amount of coughing), even if both symptoms have identical predictive power. In other words, participants may make errors in their attempts to simplify the complex world of continuous attributes that they would not make if presented with binary attributes (see also Fisher & Keil, 2018). Future covariation detection research might consider examining continuous attributes in greater detail, especially given the many life domains where continuous attributes are common (e.g., health; Chow et al., 2019; person perception, Freytag, 2003; consumer preferences; Lawless & Heymann, 2010; friendship formation; Huang et al., in press).

Another promising avenue for future research is to investigate whether functional and summarized attribute preferences have distinct downstream consequences. For example, a summarized preference such as "I like spaciousness in an apartment" may subsequently lead a person to tell his realtor to show him only relatively spacious apartments, thereby systematically restricting his own range of experiences (Ledgerwood et al., 2018; Wang et al., 2019). In other words, summarized preferences may be critical in determining how people select into novel situations (e.g., which dating website to sign up for, what kinds of schools to visit when choosing colleges; Kurzban & Weeden, 2007; Motyl, Iyer, Oishi, Trawalter, & Nosek, 2014). Conversely, functional preferences may be critical in determining people's choices after they have already had some experience in the relevant situation (e.g., whether they ask for a second date, what college they choose to attend after a visiting weekend; Wang et al., 2019).

7.2. Implications for mate preferences

Our studies were originally motivated by a desire to understand where attribute preferences such as mate preferences come from. Experimental approaches can prove central to answering such a research question (Spencer, Zanna, & Fong, 2005), and yet experimental examinations of mate preferences are surprisingly uncommon: Apart from the studies reported here, there are only three published articles reporting experiments in which researchers have attempted to change participants' mate preferences (Eagly, Eastwick, & Johannesen-Schmidt, 2009; Kille, Forest, & Wood, 2013; Nelson & Morrison, 2005). We hope that our paradigm can serve as a template for future studies to unpack how people come to arrive at the abstract judgment that they prefer particular attributes (e.g., "I like intelligence in a partner" or "I desire attractiveness in a mate").

7.2.1. Implications of the attribute quantity bias for sex differences

The effect of the quantity manipulation provides a possible new roadmap for tackling an intriguing conundrum in the (highly) complex domain where men and women express summarized and functional preferences for physical attractiveness in real live mates. This literature contains three well replicated, highly robust findings that have striking parallels to the current results. First, men and women differ in their summarized preferences for attractiveness (Buss, 1989). Second, men and women do *not* differ in their functional preferences for attractiveness (Eastwick et al., 2014).[5] Third, women are perceived (by both sexes) to be more attractive than men, on average (Eastwick & Smith, 2018; Fletcher, Kerr, Li, & Valentine, 2014; Marcus & Miller, 2003).

These three facts, in conjunction, can be mapped onto the high-complexity findings in Studies 2 and 3 to suggest a possible social-cognitive explanation for these real-world patterns: Participants who encountered high levels of *Melb* in the pool of potential mates inferred a stronger summarized preference than participants who encountered low levels of *Melb*, even though functional preferences were identical. In parallel, people who encounter high levels of attractiveness in the pool of potential mates (i.e., men evaluating women) may infer a stronger summarized preference for attractiveness than people who encounter low levels of attractiveness (i.e., women evaluating men), even though functional preferences for attractiveness are identical. Of course, further research would be required to build a bridge from the carefully controlled, stripped-down context of the current studies to the

[5] On occasion, Eastwick et al. (2014) has been mischaracterized as documenting a disconnect between summarized attribute preferences and mate choice (e.g., Gerlach et al., 2019). In fact, the Eastwick et al. (2014) meta-analysis reviewed 95 studies on functional attribute preferences for two attributes (i.e., attractiveness and earning potential) and found that these two functional preferences were not sex differentiated.

real-world mating domain. The striking parallels between lab and real world suggest such research could prove very fruitful; indeed, the social-cognitive explanation described here is, to our knowledge, the most parsimonious explanation proposed thus far that could account for the summarized versus functional discrepancy in men's and women's preferences for attractiveness (Eastwick et al., 2014).[6]

How exactly might the standard-of-comparison mechanism documented here operate in the mating domain? Consider that men and women are similar on myriad traits, any of which could serve as reference standards in real life (Hyde, 2005). Our findings suggest that people might spontaneously use any such trait (e.g., gregariousness; $d = 0.07$, Feingold, 1994) as a standard against which to judge the value of a trait such as attractiveness. In other words, one possible reason that men say they have a stronger summarized preference for attractiveness than women is because they tend to evaluate a population of partners (i.e., women) who are more attractive *than they are gregarious*, whereas women evaluate a population of partners (i.e., men) who are less attractive than they are gregarious.

Some scholars of human mating may find it strange that mate preferences for traits would reflect such seemingly irrational influences (Fletcher et al., 2014; Gerlach, Arslan, Schultze, Reinhard, & Penke, 2019). Yet if we make the modest assertion that people learn about their own mate preferences by observing how a trait covaries with their liking judgments, then the covariation detection literature unambiguously suggests that there is potential for people to be biased by information other than the contingency of interest (Vadillo et al., 2016)—including, perhaps, arbitrary referents (Bless & Schwarz, 2010; Moskowitz & Skurnik, 1999). Importantly, the standard-of-comparison account suggests novel and empirically testable predictions in the mating domain: For example, irrespective of any true underlying functional preference sex difference, the extent to which (a) women are considered to be higher than men on a particular trait on average should be associated with (b) the size of the sex difference in the summarized preference for that trait.

7.2.2. Connecting summarized preferences for Melb to real-world mate preferences

Building on longstanding theoretical accounts suggesting that people often learn about their attitudes through their life experiences (Bem, 1972; Fazio, 1987), we examined whether a similar kind of experiential learning process could explain the formation of summarized attribute preferences. Of course, it is possible that real-life mate preferences more closely resemble an instinct, and people are predisposed to find certain attributes desirable or undesirable without the need for an elaborate inference and abstraction process like the one implicated in these studies. But if any component of human mate preferences were to fit such a description, it would be the *functional preference*: the extent to which an attribute drives a person's experienced evaluation of partners that vary along that attribute dimension. That is, people might feel more aroused in the presence of attractive (vs. unattractive) people without extensive learning, just as they experience a fear of snakes without extensive learning (Öhman & Mineka, 2003). In order to make an abstract judgment like "I like physical attractiveness in a partner" (or "I am afraid of snakes"), people should have to learn and make inferences from those experienced evaluations, and it is this psychological process that our studies examined.

[6] This mechanism might also apply to gay men and lesbian women. First, the summarized preference sex difference for attractiveness is the same: Gay men have higher preferences for attractiveness than lesbian women (West, Popp, & Kenny, 2008). Second, as gay men become immersed in gay subculture, they may encounter populations of potential partners who are more concerned about physical appearance; the reverse may be true for lesbian women (Morrison, Morrison, & Sager, 2004; Siever, 1994). Nevertheless, this explanation for the preferences of gay men and lesbian women remains speculative.

Alternatively, one could argue that the present paradigm is too artificial to tell us anything useful about the learning and abstraction process that underlies real-life reports of mate preferences. Perhaps when people vividly and viscerally experience liking for a potential partner, the process of translating a functional to a summarized preference differs from the process we documented here. Indeed, one fruitful direction for future research will be to probe the generalizability of the present findings to increasingly realistic interpersonal settings. But at the same time, the consistent application of an ecological validity criterion for accepting a study as useful for understanding mate preferences would force one to disregard most of the vignette and hypothetical studies in which summarized preferences have their strongest predictive validity (e.g., see Eastwick et al., 2014, 2019 for reviews). In our view, it is important to study the consequences of mate preferences even if those consequences are most evident under tightly controlled circumstances, just as it is important to study how people form mate preferences in the first place under tightly controlled circumstances.

7.3. Limitations and future directions

We think these studies provide important initial evidence that a domain-general covariation detection mechanism may have relevance for understanding how people learn their summarized attribute preferences, including summarized preferences for attributes in a mate. However, it is possible that our findings apply to the way people form preferences for attributes generally speaking but not mate preferences in particular; after all, classic evolutionary psychological perspectives posit that mating processes will be governed by an array of domain-specific adaptations (Buss, 1995). Our approach interrogates the theoretical assumption that the mating mind largely consists of specialized mental machinery, and we explicitly recommend that scholars test the extent to which basic, domain-general mental processes may also play a role in how people make decisions in the mating domain (e.g., Huang et al., in press). Nevertheless, it is possible that domain-specific mechanisms would qualify the findings we report here, especially given that participants had no expectation that they would need to rely on their summarized preferences to make judgments or decisions in the future, the way they would with real-world mate preferences.

It is also worth noting that for the standard-of-comparison process described here to apply in more externally valid contexts, people would need to spontaneously compare different attributes without the aid of numerical scale values (e.g., "this person's stand-out quality is really her high intelligence," or "what sticks out most about him is his humor, more so than any other trait"). People certainly engage in such comparisons when thinking about themselves—that is, it is easy for people to identify traits that characterize the self more than other traits (Mueller, Thompson, & Dugan, 1986). Given that people can and do think about traits in the context of other traits (see also Hamilton & Zanna, 1974), it seems plausible that one trait may often "stand out" relative to other traits (e.g., when a person is especially intelligent but not very funny or nice), thereby enhancing the salience of the trait just like other forms of trait distinctiveness (Nelson & Miller, 1995; Skowronski & Carlston, 1989; Taylor & Fiske, 1978). Future research could also examine the extent to which escalating complexity beyond two simultaneously encountered attributes further attenuates participants' abilities and produces effect sizes that approximate those observed when people evaluate people evaluate online dating partners and/or face-to-face partners.

7.4. Conclusion

These studies examined the process by which people translate their experienced evaluative responses in the world—their functional preferences—into overall, summary judgments. Few studies have examined the intersection of these two different kinds of preferences,

largely because they tend to occupy different literatures (e.g., functional preferences are studied in the consumer products and nonhuman mating literatures; summarized preferences are studied in the human mate preferences literature; Ledgerwood et al., 2018). We offered a fresh experimental approach to understanding functional and summarized preferences by applying perspectives and methods from several related but distinct literatures—social cognition (Schwarz & Bless, 2007; Sherif & Hovland, 1961), covariation detection (Alloy & Tabachnik, 1984), and human mating (Eastwick et al., 2014). We found that the generation of an abstract, summary judgment about the desirability of an attribute may require a somewhat involved inferential process. As a part of this process, people consider the underlying functional attribute preference (i.e., the extent to which the attribute inspires liking), but they also consider extraneous information, such as the quantity of a trait in the population. Future research on attitudes toward attributes—including the widely studied topic of mate preferences—may benefit from continuing to borrow both theory and paradigms from the corpus of work on social cognition and covariation detection.

Acknowledgements

This research was supported by UC Davis Academic Senate Small Grants. We wish to thank Jeffrey Sherman and Bertram Gawronski for providing valuable feedback on prior drafts of this manuscript.

Supplemental materials - Appendices A-G

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jesp.2019.103837.

References

Ajzen, I., & Fishbein, M. (1977). Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin, 84*, 888–918.
Allan, L. G., & Jenkins, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning and Motivation, 14*, 381–405.
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. https://doi.org/10.1037/0033-295X.91.1.112.
Allred, S. R., Crawford, L. E., Duffy, S., & Smith, J. (2016). Working memory and spatial judgments: Cognitive load increases the central tendency bias. *Psychonomic Bulletin & Review,* 1–7. https://doi.org/10.3758/s13423-016-1039-0.
Almaraz, S. M., Hugenberg, K., & Young, S. G. (2018). Perceiving sophisticated minds influences perceptual individuation. *Personality and Social Psychology Bulletin, 44*, 143–157.
Anderson, N. H. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology, 9*, 272–279.
Arkes, H. R., & Harkness, A. R. (1983). Estimates of contingency between two dichotomous variables. *Journal of Experimental Psychology: General, 112*, 117–135.
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 (Vol. Ed.), *Advances in experimental social psychology. Vol. 6. Advances in experimental social psychology* (pp. 1–62). New York, NY: Academic Press.
Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics, 92*, 1787–1794. https://doi.org/10.1016/j.jpubeco.2008.04.010.
Bless, H., & Schwarz, N. (2010). Mental construal and the emergence of assimilation and contrast effects: The inclusion/exclusion model. In M. P. Zanna, & M. P. Zanna (Vol. Eds.), *Advances in experimental social psychology. Vol. 42. Advances in experimental social psychology* (pp. 319–373). San Diego, CA, US: Academic Press. https://doi.org/10.1016/S0065-2601(10)42006-7.
Braver, S. L., Thoemmes, F. J., & Rosenthal, R. (2014). Continuously cumulating meta-analysis and replicability. *Perspectives on Psychological Science, 9*, 333–342. https://doi.org/10.1177/1745691614529796.
Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science, 6*, 3–5.
Buss, D. M. (1989). Sex differences in human mate preferences: Evolutionary hypotheses tested in 37 cultures. *Behavioral and Brain Sciences, 12*, 1–49. https://doi.org/10.1017/S0140525X00023992.
Buss, D. M. (1995). Evolutionary psychology: A new paradigm for psychological science. *Psychological Inquiry, 6*, 1–30.
Callahan, S. P., & Ledgerwood, A. (2016). On the psychological function of flags and logos: Group identity symbols increase perceived entitativity. *Journal of Personality and Social Psychology, 110*, 528–550.
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., & 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.
Chow, J. Y., Colagiuri, B., & Livesey, E. J. (2019). Bridging the divide between causal illusions in the laboratory and the real world: The effects of outcome density with a variable continuous outcome. *Cognitive research: Principles and implications. Vol. 4. Cognitive research: Principles and implications* (pp. 1–15).
Cohen, S. E., & Todd, P. M. (2019). Stated and revealed preferences in companion animal choice. *Behavior Research Methods* (in press).
Crocker, J. (1981). Judgment of covariation by social perceivers. *Psychological Bulletin, 90*, 272–292. https://doi.org/10.1037/0033-2909.90.2.272.
Crocker, J. (1982). Biased questions in judgment of covariation studies. *Personality and Social Psychology Bulletin, 8*, 214–220. https://doi.org/10.1177/0146167282082005.
De Dreu, C. K. W. (2003). Time pressure and closing of the mind in negotiation. *Organizational Behavior and Human Decision Processes, 91*, 280–295.
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.* Orlando, FL, US: Harcourt Brace Jovanovich College Publishers.
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., & 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. https://doi.org/10.1037/0022-3514.94.2.245.
Eastwick, P. W., Finkel, E. J., & Simpson, J. A. (2019). Best practices for testing the predictive validity of ideal partner preference-matching. *Personality and Social Psychology Bulletin, 45*, 167–181.
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. https://doi.org/10.1037/a0032432.
Eastwick, P. W., & Smith, L. K. (2018). Sex-differentiated effects of physical attractiveness on romantic desire: A highly powered, preregistered study in a photograph evaluation context. *Comprehensive Results in Social Psychology, 3*, 1–27.
Fazio, R. H. (1987). Self-perception theory: A current perspective. In M. P. Zanna, J. M. Olson, & C. P. Herman (Vol. Eds.), *Social influence: The Ontario symposium. Vol. 5. Social influence: The Ontario symposium* (pp. 129–150). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.
Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Vol. Ed.), *Advances in experimental social psychology. Vol. 25. Advances in experimental social psychology* (pp. 75–109). New York, NY: Academic Press.
Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology, 87*, 293–311.
Fazio, R. H., Lenn, T. M., & Effrein, E. A. (1984). Spontaneous attitude formation. *Social Cognition, 2*, 217–234. https://doi.org/10.1521/soco.1984.2.3.217.
Fazio, R. H., Pietri, E. S., Rocklage, M. D., & Shook, N. J. (2015). Chapter three-positive versus negative valence: Asymmetries in attitude formation and generalization as fundamental individual differences. In J. M. Olson, & M. P. Zanna (Vol. Eds.), *Advances in experimental social psychology. Vol. 51. Advances in experimental social psychology* (pp. 97–146). Burlington: Academic Press.
Feingold, A. (1994). Gender differences in personality: A meta-analysis. *Psychological Bulletin, 116*, 429–456.
Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: An integrative account of an intriguing cognitive illusion. *Psychological Review, 116*, 187–206.
Fiedler, K., Walther, E., Freytag, P., & Stryczek, E. (2002). Playing mating games in foreign cultures: A conceptual framework and an experimental paradigm for inductive trivariate inference. *Journal of Experimental Social Psychology, 38*, 14–30.
Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention, and behavior: An introduction to theory and research.* Reading, MA: Addison-Wesley.
Fisher, M., & Keil, F. C. (2018). The binary bias: A systematic distortion in the integration of information. *Psychological Science, 29*, 1846–1858.
Fletcher, G. J., Kerr, P. S., Li, N. P., & Valentine, K. A. (2014). Predicting romantic interest and decisions in the very early stages of mate selection: Standards, accuracy, and sex differences. *Personality and Social Psychology Bulletin, 40*, 540–550.
Fletcher, G. O., Simpson, J. A., Thomas, G., & Giles, L. (1999). Ideals in intimate relationships. *Journal of Personality and Social Psychology, 76*, 72–89.
Freytag, P. (2003). *Contextually determined typicality.* Berlin, Germany: Logos.
Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: Exploring mechanisms for face recognition. *Vision Research, 37*, 1673–1682.
Gerlach, T. M., Arslan, R. C., Schultze, T., Reinhard, S. K., & Penke, L. (2019). Predictive validity and adjustment of ideal partner preferences across the transition into romantic relationships. *Journal of Personality and Social Psychology, 116*, 313–330.
Goh, J. X., Hall, J. A., & Rosenthal, R. (2016). Mini meta-analysis of your own studies: Some arguments on why and a primer on how. *Social and Personality Psychology Compass, 10*, 535–549.
Gosling, S. D., & Johnson, J. A. (2010). *Advanced methods for conducting online behavioral research.* Washington, DC: American Psychological Association.
Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12*, 392–407.

Hamilton, D. L., & Zanna, M. P. (1974). Context effects in impression formation: Changes in connotative meaning. *Journal of Personality and Social Psychology, 29*, 649–654.
Herr, P. M. (1986). Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology, 51*, 1106–1115. https://doi.org/10.1037/0022-3514.51.6.1106.
Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology, 19*, 323–340. https://doi.org/10.1016/0022-1031(83)90026-4.
Hill, R. (1945). Campus values in mate selection. *Journal of Home Economics, 37*, 554–558.
Huang, S. A., Ledgerwood, A., & Eastwick, P. W. (2019). How do ideal friend preferences and interaction context affect friendship formation? Evidence for a domain-general relationship initiation process. *Social Psychological and Personality Science* (in press).
Hyde, J. S. (2005). The gender similarities hypothesis. *American Psychologist, 60*, 581–592. https://doi.org/10.1037/0003-066X.60.6.581.
Kasof, J. (1993). Sex bias in the naming of stimulus persons. *Psychological Bulletin, 113*, 140–163. https://doi.org/10.1037/0033-2909.113.1.140.
Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist, 28*, 107–128.
Kille, D. R., Forest, A. L., & Wood, J. V. (2013). Tall, dark, and stable: Embodiment motivates mate selection preferences. *Psychological Science, 24*, 112–114.
Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*, 211–228. https://doi.org/10.1037/0033-295X.94.2.211.
Kraut, R., Olson, J., Banaji, M., Bruckman, A., Cohen, J., & Couper, M. (2004). Psychological research online: Report of Board of Scientific Affairs' Advisory Group on the Conduct of Research on the Internet. *American Psychologist, 59*, 105–117.
Kristof, A. L. (1996). Person-organization fit: An integrative review of its conceptualizations, measurement, and implications. *Personnel Psychology, 49*, 1–49.
Kurzban, R., & Weeden, J. (2007). Do advertised preferences predict the behavior of speed daters? *Personal Relationships, 14*, 623–632.
Lakens, D., & Etz, A. J. (2017). Too true to be bad: When sets of studies with significant and nonsignificant findings are probably true. *Social Psychological and Personality Science, 8*, 875–881.
Lawless, H. T., & Heymann, H. (2010). *Sensory evaluation of food: Principles and practices.* New York, NY: Springer Science & Business Media.
Ledgerwood, A., & Chaiken, S. (2007). Priming us and them: Automatic assimilation and contrast in group attitudes. *Journal of Personality and Social Psychology, 93*, 940–956.
Ledgerwood, A., Eastwick, P. W., & Smith, L. K. (2018). Toward an integrative framework for studying human evaluation: Attitudes towards objects and attributes. *Personality and Social Psychology Review, 22*, 378–398.
Ledgerwood, A., Soderberg, C. K., & Sparks, J. (2017). Designing a study to maximize informational value. In J. Plucker, & M. Makel (Eds.). *Toward a more perfect psychology: Improving trust, accuracy, and transparency in research*. Washington DC: American Psychological Association.
Marcus, D. K., & Miller, R. S. (2003). Sex differences in judgments of physical attractiveness: A social relations analysis. *Personality and Social Psychology Bulletin, 29*, 325–335.
Marsh, J. K., & Ahn, W. K. (2009). Spontaneous assimilation of continuous values and temporal information in causal induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*, 334–352.
McShane, B. B., & Böckenholt, U. (2017). Single-paper meta-analysis: Benefits for study summary, theory testing, and replicability. *Journal of Consumer Research, 43*, 1048–1063.
Møller, A. P. (1988). Female choice selects for male sexual tail ornaments in the monogamous swallow. *Nature, 332*, 640–642.
Morrison, M. A., Morrison, T. G., & Sager, C.-L. (2004). Does body satisfaction differ between gay men and lesbian women and heterosexual men and women?: A meta-analytic review. *Body Image, 1*, 127–138.
Moskowitz, G. B. (2004). *Social cognition: Understanding self and others.* New York, NY: Guilford Press.
Moskowitz, G. B., & Skurnik, I. W. (1999). Contrast effects as determined by the type of prime: Trait versus exemplar primes initiate processing strategies that differ in how accessible constructs are used. *Journal of Personality and Social Psychology, 76*, 911–927.
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.
Mueller, J. H., Thompson, W. B., & Dugan, K. (1986). Trait distinctiveness and accessibility in the self-schema. *Personality and Social Psychology Bulletin, 12*, 81–89.
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.
Nelson, L. J., & Miller, D. T. (1995). The distinctiveness effect in social categorization: You are what makes you unusual. *Psychological Science, 6*, 246–249.
Öhman, A., & Mineka, S. (2003). The malicious serpent: Snakes as a prototypical stimulus for an evolved module of fear. *Current Directions in Psychological Science, 12*, 5–9.
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. https://doi.org/10.2307/2489395.
Perales, J. C., Catena, A., Shanks, D. R., & González, J. A. (2005). Dissociation between judgments and outcome-expectancy measures in covariation learning: A signal detection theory approach. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 1105–1120.
Price, P. C., & Yates, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition, 21*, 561–572.
Richard, F. D., Bond, C. F., Jr., & Stokes-Zoota, J. J. (2003). One hundred years of social psychology quantitatively described. *Review of General Psychology, 7*, 331–363.
Sainsbury, R. S., & Jenkins, H. M. (1967). Feature-positive effect in discrimination learning. *Proceedings of the annual convention of the American Psychological Association. Vol. 2. Proceedings of the annual convention of the American Psychological Association* (pp. 17–18).
Samuelson, P. A. (1948). Consumption theory in terms of revealed preference. *Economica, 15*, 243–253. https://doi.org/10.2307/2549561.
Schacter, D. L., & Addis, D. R. (2007). Constructive memory: The ghosts of past and future. *Nature, 445*, 27.
Schaller, M., & O'Brien, M. (1992). Intuitive analysis of covariance and group stereotype formation. *Personality and Social Psychology Bulletin, 18*, 776–785.
Schwarz, N. (1999). Self-reports: How the questions shape the answers. *American Psychologist, 54*, 93–105. https://doi.org/10.1037/0003-066X.54.2.93.
Schwarz, N., & Bless, H. (1992). Scandals and the public's trust in politicians: Assimilation and contrast effects. *Personality and Social Psychology Bulletin, 18*, 574–579.
Schwarz, N., & Bless, H. (2007). Mental construal processes: The inclusion/exclusion model. In D. A. Stapel, & J. Suls (Eds.). *Assimilation and contrast in social psychology* (pp. 119–141). Philadelphia, PA: Psychology Press.
Shaklee, H., & Mims, M. (1981). Development of rule use in judgments of covariation between events. *Child Development, 52*, 317–325. https://doi.org/10.2307/1129245.
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. https://doi.org/10.1037/0278-7393.8.3.208.
Shaklee, H., & Tucker, D. (1980). A rule analysis of judgments of covariation between events. *Memory & Cognition, 8*, 459–467. https://doi.org/10.3758/BF03211142.
Sherif, M., & Hovland, C. I. (1961). *Social judgment: Assimilation and contrast effects in communication and attitude change.* New Haven, CT: Yale University Press.
Siever, M. D. (1994). Sexual orientation and gender as factors in socioculturally acquired vulnerability to body dissatisfaction and eating disorders. *Journal of Consulting and Clinical Psychology, 62*, 252–260.
Simonsohn, U. (2014). [17] No-way interactions. Retrieved from http://web.archive.org/web/20150206205257/http://data-colada.org/2014/03/12/17-no-way-interactions-2/.
Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin, 105*, 131–142.
Smedslund, J. (1963). The concept of correlation in adults. *Scandinavian Journal of Psychology, 4*, 165–173. https://doi.org/10.1111/j.1467-9450.1963.tb01324.x.
Spencer, S. J., Zanna, M. P., & Fong, G. T. (2005). Establishing a causal chain: Why experiments are often more effective than mediational analyses in examining psychological processes. *Journal of Personality and Social Psychology, 89*, 845–851.
Stam, H. J. (1996). Theory & practice. In C. Tolman, F. Cherry, R.v. Hezewijk, & I. Lubek (Eds.). *Problems of theoretical psychology* (pp. 24–32). Toronto: Captus Press.
Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences, 30*, 299–313.
Taylor, S. E., & Fiske, S. T. (1978). Salience, attention, and attribution: Top of the head phenomena. In L. Berkowitz (Vol. Ed.), *Advances in experimental social psychology. Vol. 11. Advances in experimental social psychology* (pp. 249–288). New York: Academic Press.
Trolier, T. K., & Hamilton, D. L. (1986). Variables influencing judgments of correlational relations. *Journal of Personality and Social Psychology, 50*, 879–888.
Tweney, R. D., Doherty, M. E., Worner, W. J., Pliske, D. B., Mynatt, C. R., Gross, K. A., & Arkkelin, D. L. (1980). Strategies of rule discovery in an inference task. *Quarterly Journal of Experimental Psychology, 32*, 109–123. https://doi.org/10.1080/00335558008248237.
Vadillo, M. A., Blanco, F., Yarritu, I., & Matute, H. (2016). Single-and dual-process models of biased contingency detection. *Experimental Psychology, 63*, 3–19.
Vadillo, M. A., Musca, S. C., Blanco, F., & Matute, H. (2011). Contrasting cue-density effects in causal and prediction judgments. *Psychonomic Bulletin & Review, 18*, 110–115.
Vosgerau, J., Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2018). *Internal meta-analysis makes false-positives easier to produce and harder to correct.* (Available at SSRN).
Wang, Y. A., Eastwick, P. W., & Ledgerwood, A. (2019). *Ideas about liking predict situation selection at a distance. Manuscript under review.* Davis, CA: University of California.
Ward, W. C., & Jenkins, H. M. (1965). The display of information and the judgment of contingency. *Canadian Journal of Psychology, 19*, 231–241. https://doi.org/10.1037/h0082908.
West, T. V., Popp, D., & Kenny, D. A. (2008). A guide for the estimation of gender and sexual orientation effects in dyadic data: An actor-partner interdependence model approach. *Personality and Social Psychology Bulletin, 34*, 321–336.
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. https://doi.org/10.1037/a0015300.
Wood, D., Lowman, G. H., Harms, P., & Roberts, B. W. (2019). Exploring the relative importance of normative and distinctive organizational preferences as predictors of work attitudes. *Journal of Applied Psychology, 104*, 270–292.