YU-PING CHEN, LEIF D. NELSON, and MING HSU*

Considerable attention has been given to the notion of a set of human-like characteristics associated with brands, referred to as "brand personality." The authors combine newly available machine learning techniques with functional neuroimaging data to characterize the set of processes that give rise to these associations. The authors show that brand personality traits can be captured by the weighted activity across a widely distributed set of brain regions previously implicated in reasoning, imagery, and affective processing. That is, as opposed to being constructed through reflective processes, brand personality traits seem to exist a priori inside consumers' minds, such that the authors are able to predict what brand a person is thinking about solely on the basis of the relationship between brand personality associations and brain activity. These findings represent an important advance in the application of neuroscientific methods to consumer research, moving from work focused on cataloging brain regions associated with marketing stimuli to testing and refining constructs central to theories of consumer behavior.

Keywords: consumer neuroscience, branding, brand personality, functional magnetic resonance imaging, machine learning

Online Supplement: http://dx.doi.org/10.1509/jmr.14.0606

From "Where" to "What": Distributed Representations of Brand Associations in the Human Brain

Marketers have long appreciated the role of brand positioning, the location that a brand occupies in consumers' minds relative to competing offerings, in guiding managerial decision making (Aaker 2009; Gardner and Levy 1955; Keller 1993). An understanding of how consumers feel and think about brands, for example, provides valuable guidance for developing marketing strategy in such areas as advertising, pricing, and channel strategies. Moreover, as branding has grown to focus increasingly on abstract and intangible considerations, marketers have worked to understand aspects of brand knowledge not related to the actual physical product or service specifications per se (Aaker 2012; Keller 2003).

In response, consumer researchers have expended considerable effort to decompose consumer responses to brands into their component parts (e.g., feelings, imagery, likability) (Alba and Hutchinson 1987; Bettman 1970; Keller 2003; Zaltman and Coulter 1995), which has resulted in a set of sophisticated typologies that provides rigorous scientific characterization to these complex perceptions. One canonical typology, for example, involves the characterization of the widely held notion that consumers endow brands with a set of human-like characteristics akin to personality (Aaker 1997; Levy 1959). The resulting brand personality frame-

*Yu-Ping Chen is Assistant Professor of Marketing, Department of International Business, National Taiwan University (e-mail: yupingchen@ntu.edu.tw). Leif D. Nelson is Ewald T. Grether Professor in Business Administration & Marketing, Haas School of Business, University of California, Berkeley (e-mail: leif_nelson@haas.berkeley.edu). Ming Hsu is Assistant Professor, Haas School of Business and Helen Wills Neuroscience Institute, University of California, Berkeley (e-mail: mhsu@haas.berkeley.edu). This research was supported by National Institute of Mental Health grant MH-098023 (MH) and the Hellman Family Faculty Fund (MH). The authors thank Victoria Chen for assistance with data collection. All trademarks in this article are the property of their respective owners. This article was processed by Editor in Chief Robert Meyer.

ISSN: 0022-2437 (print), 1547-7193 (electronic)

work, as proposed in Aaker's (1997) seminal work, uncovers five basic dimensions that together provide a highly robust and general account of the perceptual space underlying brands.

Despite these successes, research in consumer psychology has been largely silent on the specific processes by which intangible characteristics such as brand personality are generated and organized (Johar, Sengupta, and Aaker 2005; Keller and Lehmann 2003). More broadly, because mental constructs such as brand personality have traditionally only been measured by self-report methods, it remains challenging for researchers to probe such knowledge in cases in which consumers are unable or unwilling to fully articulate their thoughts and preferences (Ariely and Berns 2010; Haire 1950; Zaltman and Coulter 1995). Such insights are central to marketers' efforts to understand and predict the extent to which marketing actions can successfully create or affect these thoughts and feelings, which in turn influence consumer response to marketing activities (Batra, Lenk, and Wedel 2010; Van der Lans, Van den Bergh, and Dieleman 2014).

Emerging techniques in neuroscience, therefore, have been widely viewed as having the potential to overcome limitations of self-report measures by directly accessing consumers' mental contents (Ariely and Berns 2010; Plassmann, Ramsøy, and Milosavljevic 2012; Yoon et al. 2012). Perhaps most excitingly, by capturing the entire decision-making process, modern functional neuroimaging techniques have the ability to elucidate the multitude of processes engaged during consumer choice, such that the effects of marketing actions on these processes can be traced, compared, and valued.

In the context of branding, an important open question pertains to the extent to which a stable "mental map" of brand knowledge exists from which brand personality associations emerge (Keller 2003; Zaltman 1997). This is important for two reasons. First, the assumption of a stable store of knowledge underlies all existing research efforts using self-report measures to probe the intangible characteristics consumers associate with brands. Substantial research has suggested, however, that recall is often not equivalent to retrieval of information in memory but may be the construction of a plausible response (Johar, Maheswaran, and Peracchio 2006). In the extreme case, participant responses may be constructed to suit the explicit questions of consumer researchers, and these explicit measures have little to do with participants' actual thoughts about the brands. That is, it is unclear whether intangible characteristics such as brand personality traits exist a priori in consumers' minds or whether they are a product of reflective process, such that they are influenced by experimenter elicitation. Second, the existence of such a map opens the door for neuroscientific methods to address several additional important questions, such as how marketing actions affect consumers' mental representations of brand personality and the nature of the different cognitive processes that act on these representations.

Although of course still preliminary and incomplete, existing studies using functional neuroimaging techniques have already made important inroads in addressing some of these processes. For example, such research has provided evidence for inferences about the role of emotional processing in decoy effects on the basis of amygdala activation (Hedgcock and Rao 2009), in which the introduction of a third normatively irrelevant alternative was associated with significantly lower activation in areas of the brain associated with negative emotion.

"WHAT?" VERSUS "WHERE?"

Despite advances in previous research, there remain important conceptual and methodological hurdles that arise from fundamental differences between the typical goals and questions in neuroscience and marketing. In particular, localization approaches in cognitive neuroscience are inherently focused on "where"-type questions (Churchland and Sejnowski 1988; Gazzaniga 2004). For example, where in the brain does overall activation between animate and inanimate objects differ (Kriegeskorte et al. 2008)? Does the hippocampus engage more vigorously during episodic memory retrieval versus encoding (Schacter and Wagner 1999)?

Answering such questions has been invaluable in understanding how the brain organizes basic cognitive processes and how they relate to more complex constructs and representations. The finding that altruistic punishment engages brain regions known to respond to basic rewards provided early evidence that altruistic punishment may also be rewarding at a basic neurobiological level (De Quervain et al. 2004). In the context of brand personality, Yoon et al.'s (2006) pioneering study indicated important differences in processes at the neural level that are associated with trait judgments about brands and people. Specifically, compared with judgment of human traits, judgment of brand traits elicited greater engagement of the inferior prefrontal cortex, an area known to be involved in object processing, thereby challenging a strictly anthropomorphic view of brand personality.

For many (if not most) consumer researchers, however, these "where"-type questions are secondary to understanding the brain's contents and processes. That is, consumer researchers, in contrast to neuroscientists, are typically interested in "what"-type questions. For example, what is the set of associations that goes through consumers' minds when they are presented with a particular brand? How do marketing actions affect these associations?

Despite the intuitive nature of such questions, previous neuroimaging studies have not been equipped to address them. Specifically, whereas neuroscience has generally been able to deliver "where" answers, marketing continues to ask "what" questions. For example, marketers might ask, "What is going through consumers' minds when looking at a Coca-Cola advertisement?" but neuroscience has traditionally delivered, "The value of Coca-Cola can be detected in regions such as the ventromedial prefrontal cortex."

In particular, localization approaches may fail to capture representations and processes that are not contained in any single set of brain regions but rather emerge from the correlated activity across a network of brain areas (Kriegeskorte, Goebel, and Bandettini 2006; Mitchell et al. 2008). That complex constructs such as conceptual knowledge emerge out of a distributed system has a long and distinguished history dating back at least to Lashley's (1950) search for engrams and connectionist models of learning systems (Hinton, McClelland, and Rumelhart 1986; McClelland and Rogers 2003).

At the extreme, an inability to address "what"-type questions leaves open the possibility that brain regions believed to underlie a specific process are actually involved in some completely unrelated process. For example, amygdala activation in the decoy effects may instead be related to some other aspect of the task that has nothing to do with decoy effects (Huettel et al. 2009; Poldrack 2011). This possibility is particularly salient in the case of consumer neuroscience, given the complexity of marketing stimuli. One way to address this concern is to show that the information content in question is actually contained within the set of identified brain regions.

CONNECTING "WHAT" AND "WHERE"

In this article, we take an important step toward enabling consumer researchers to address both "what" and "where" questions using brain imaging data (Kay et al. 2008; Kriegeskorte, Goebel, and Bandettini 2006; Mitchell et al. 2008). In more basic cognitive processes such as vision and memory, these methods have revolutionized researchers' abilities to ask questions about how information is encoded, maintained, and retrieved at various stages of processing in ways that test and inform psychological theories of memory and perception (Kay et al. 2008; Rissman and Wagner 2012). The central insight of this approach is to use cross-validation techniques to consider whether a distributed set or "pattern" of brain activity contains some set of information predicted by cognitive and behavioral theories (Kriegeskorte, Goebel, and Bandettini 2006; Poldrack 2011).

First, to address the "what" question, we attempt to recover the set of thoughts and feelings that consumers associate with brands in a passive viewing task. Importantly, the participant in our experiment is not prompted to make any specific judgment but, rather, is asked to freely think about the brand. If brand personality traits associated with brands exist in the mind of the consumer a priori, we should in principle be able to "read out" these contents on the basis of brain activity alone; however, this would not be possible if traits are solely the consequence of ratings prompted by the researcher.

This approach is based on two key assumptions. First, we assume that the mental representation of brand personality is contained in the responses of a stable and possibly distributed network of regions (Kriegeskorte, Goebel, and Bandettini 2006; Mitchell et al. 2008). That is, there exists a stable mapping between brain and mind such that the mental representation of brand personality is reflected in the activity levels of a network of brain regions. Second, we assume that the psychological architecture provides a reasonable first-order approximation of the mental representation (Mitchell et al. 2008; Poldrack 2011). In the case of brand personality, this is equivalent to assuming that each brand is located within a five-dimensional representation space (captured by, e.g., sincerity, competence), where the specific location is given as a five-tuple within the space.

Assumption 1: There exists a neural representation, consisting of a widely distributed network, of mental representation of brand personality.

Assumption 2: The brand personality framework captures mental representations of a set of intangible brand characteristics.

Importantly, our second assumption makes clear the distinction between our approach and those of previous studies aimed at predicting consumer choice (Deppe et al. 2005; Murawski et al. 2012; Tusche, Bode, and Haynes 2010; Van der Laan et al. 2012). In this latter set of studies, the authors conducted decoding based on observable choice behavior and made no attempt to test the plausibility of models of the underlying psychological processes. In the same way that early decoding studies of visual systems (e.g., Haxby et al. 2001; Haynes and Rees 2005) were conducted with no reference to the intermediate psychological features underlying observable *inputs* (e.g., faces, houses), these studies make no references to intermediate psychological processes underlying observable *outputs*. In contrast, our approach is referred to as model-based decoding, which is distinct from those that do not assume some underlying model of the representational space (for details, see Haynes and Rees 2006; Poldrack 2011).

More specifically, by identifying the particular brand a person is thinking about on the basis of the evoked brain responses, our study requires the brand personality framework to offer greater predictive power than null models that do not capture these characteristics. That is, drawing on how a person's brain differentially responds to Coca-Cola and Pepsi, we investigate whether it is possible to learn about the representational space of brand personality in the brain and use this relationship to infer whether that person is thinking about Apple or Microsoft.

H_1: Brand personality traits associated with brands exist in the mind of the consumer a priori and can be recovered from brain activity during a passive viewing task.

Next, to connect "what" to "where," we characterize the set of brain regions that contain brand personality information. This enables us to address the extent to which brand personality contents are distributed in the brain. In previous decoding studies, contents related to more basic perceptual processes have been found to be contained in relatively circumscribed regions of the occipital and temporal lobes (Kriegeskorte et al. 2008; Naselaris et al. 2009). This is the case even for relatively abstract constructs such as objects and faces, which are largely restricted to regions within the inferior temporal cortex or biological motion in the superior temporal sulcus (Haynes and Rees 2005; Kriegeskorte et al. 2008). In contrast, higher-order constructs such as conceptual knowledge have been shown to have a much more distributed neural basis, drawing on a wide set of brain regions, including those involved in sensory processing as well as higher-order cognitive regions (Mitchell et al. 2008; Tyler and Moss 2001).

More importantly, the resulting map of predictive regions enables us to make inferences about the processes by which brand personality emerges. Previous neuroimaging studies have implicated a diverse array of brain regions in brand processing, including regions involved in autobiographical memory and person judgment (medial prefrontal cortex [mPFC]; Deppe et al. 2005; Schaefer et al. 2006; Schaefer and Rotte 2010), semantic memory retrieval (lateral prefrontal cortex [lPFC]; Klucharev, Smidts, and Fernández 2008; McClure et al. 2004; Yoon et al. 2006), affective processing and interoception (insula; Bruce et al. 2013), and episodic and spatial memory (hippocampus, Esch et al. 2012; McClure et al. 2004), among others. Although these findings are typically discussed in isolation, it is possible that they all reflect a shared set of cognitive and affective processes from which brand personality representation emerges.

H_2: Consistent with connectionist models of learning and memory, brand personality contents are distributed widely across the brain.

METHODS

Participants

A total of 17 participants (6 women; mean age = 34.2 years, SD = 6.5) from the San Francisco Bay area were recruited from Craigslist to participate in the functional magnetic resonance imaging (fMRI) study. Although this sample size is on the lower end of standard functional neuroimaging studies based on univariate approaches, it is on par with or exceeds those of comparable multivariate decoding studies (Formisano et al. 2008; Mitchell et al. 2008). The total time for the whole experiment was approximately three hours, including the instruction, the scanning session, and the postexperiment questionnaires. Each participant was paid $70 in cash upon completion of the experiment. An additional 25 undergraduate students were recruited for a behavior-only study in exchange for course credit. These participants completed an online questionnaire on the same set of brands and traits of the brand association scale. All informed consent was obtained as approved by the Internal Review Board at University of California, Berkeley.

Procedure

Participants in the fMRI study underwent scanning in a passive viewing task involving logos of 44 well-known brands (Figure 1, Panel A). We selected the set of brands from Interbrand's list of 100 Best Global Brands (www.interbrand.com) to ensure diversity in brand associations and represented industries. Each of the 44 stimulus items was presented four times in a pseudo random sequence on a gray background (Figure 1, Panel B), and each presentation lasted for four to eight seconds. Before the scanning session, participants were instructed to think about the characteristics or traits associated with the brand but told that they were free to think about any characteristic or trait such that no attempt was made to obtain consistency of the associations either across participants or across repetition times. Following the scanning session, participants completed a survey that included the 42-item brand association scale (Aaker 1997), familiarity, and preference for each of the 44 brands. The brand association scale involved judgment of the descriptiveness of 42 traits for each brand (see Table S1 in the Web Appendix), with a five-point scale (1 = "not at all descriptive," and 5 = "extremely descriptive").

fMRI Data Acquisition

Functional images were acquired on a Siemens 3T TIM/Trio scanner at Henry H. Wheeler Jr. Brain Imaging Center at University of California, Berkeley. We used an echo planar imaging sequence to acquire the functional data: repetition time = 2,000 ms; echo time = 30 ms; voxel resolution = 3 mm × 3 mm × 3 mm; field of view read = 192 mm; field of view phase = 100%; interleaved series order. The scan sequences were axial slices approximately flipped 30 degrees to the anterior commissure–posterior commissure axis. We acquired high-resolution structural T1-weighted scans (1 mm × 1 mm × 1 mm) by using an MPRage sequence.

Behavioral Data Analysis

To characterize personality features associated with our brands using participant ratings on the set of traits outlined in the Aaker (1997) framework (Figure 1, Panel C), we used a factor-analytic approach to summarize variation in trait ratings and reduce collinearity issues. We factor-analyzed mean trait ratings using principal components analysis and Varimax rotation. We selected factors if the associated eigenvalues were greater than 1 and explained a significant portion of variance (see Table S2 in the Web Appendix). Each brand was reexpressed in terms of its personality vector, defined as the strength of association between the brand and the personality factors (e.g., excitement, competence).

fMRI Data Preprocessing

Image data were preprocessed in the following order using SPM8 (Statistical Parametric Mapping, Wellcome Trust Centre for Neuroimaging): correction for slice time artifacts, realignment, coregistration to the participant's T1 image, and normalization to Montreal Neurological Institute coordinates. Finally, consistent with previous multivoxel pattern analysis studies, data were left unsmoothed to preserve local voxel information (Clithero, Carter, and Huettel 2009; Haynes and Rees 2006).

fMRI Data Analysis

Figure 2 presents an illustration of our analytical approach. We summarize the main analytical process before describing the steps in more detail. Following extraction of a representative fMRI image for each brand, we held two brands out of the set of 44 total brands (e.g., Disney and Gucci; see Figure 2, Panel A). We then used these brain responses, together with the brand personality factors for the 42 remaining in-sample brands (Figure 2, Panel B), to obtain an fMRI map for each of the five brand personality factors (Figure 2, Panel C) so we could calculate predicted fMRI maps for each of the two holdout fMRI images for Disney and Gucci by combining the brand personality factor scores of the holdout brands with the brand personality fMRI maps (Figure 2, Panel D). Finally, we determined whether we could correctly predict whether each holdout brand is Disney or Gucci by comparing the similarity between the predicted and actual neural maps. After completion, we iterated this procedure over all possible pairwise combinations of brands and performed significance testing using a permutation procedure by shuffling over the fMRI image and brand personality pairings. Next, we provide more detailed description of the procedures.

Extracting neural responses to brands. To identify the representative fMRI image of a brand, we used the procedure outlined in Mumford et al. (2012) to account for the fact that in rapid event-related designs, the evoked blood oxygen level–dependent signal for adjacent trials will over-

Figure 1
EXPERIMENTAL PARADIGM AND BEHAVIORAL RESULTS

A: The 44 Brands and Their Associated Logos Used in the Experiment[a]

B: Experimental Protocol: fMRI Passive Viewing Task[b]

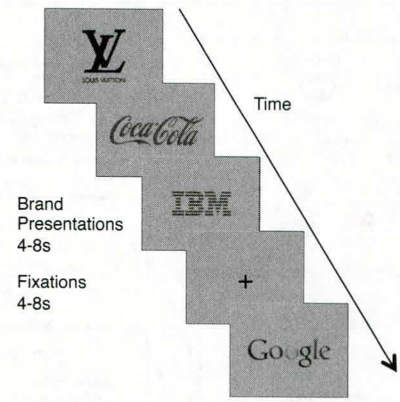

C: Quantitative Descriptions of Brand Association[c]

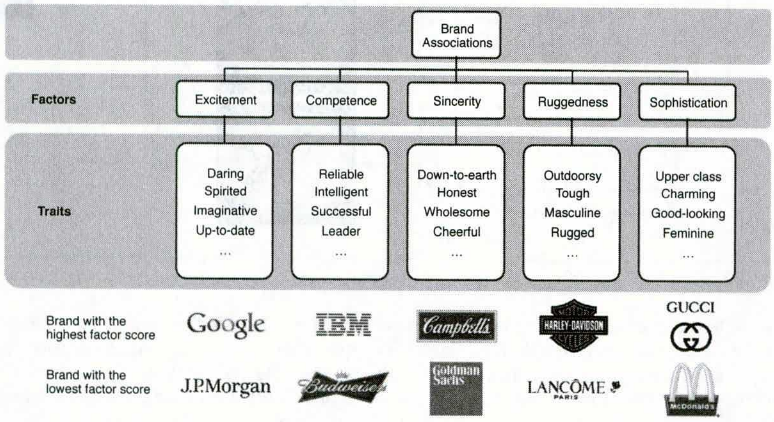

D: Radar Chart of Example Brands[d]

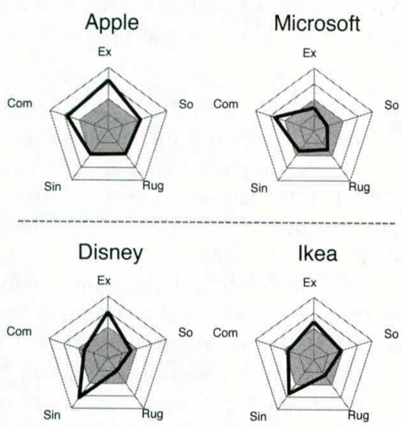

E: Robustness of Association Architecture[e]

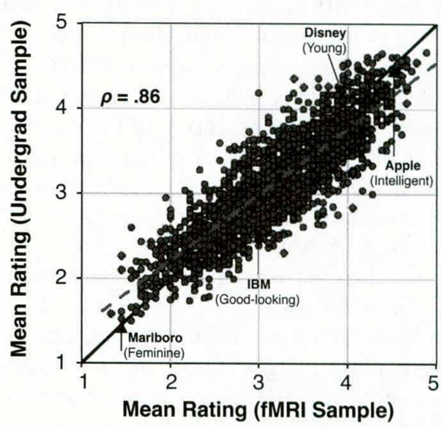

[a]Brands were chosen from Interbrand's list of top global brands.

[b]Participants were instructed to think about the characteristics or traits associated with each brand. For each trial, a brand logo was presented for four to eight seconds on a gray background.

[c]We derived the quantitative descriptions using the Aaker (1997) brand association framework. For each brand, participants rated a set of 42 traits, yielding a set of five latent features through factor analysis. Examples of the extreme brands are presented at the bottom to illustrate how brand associations were captured in this framework.

[d]The brands in this chart reside in the same industry but possess distinct associations (Apple and Microsoft) or reside in different industries but possess similar associations (Disney and IKEA). Each vertex indicates a brand personality factor (Ex = excitement, Com = competence, Sin = sincerity, Rug = ruggedness, So = sophistication). The radar chart for each brand shows the brand's factor score on each of the five dimensions. Shaded (unshaded) regions indicate negative (positive) factor scores.

[e]Participant ratings were highly correlated with those from an independent pool of undergraduate students (Pearson $r = .86$, $p < 10^{-10}$).

Figure 2
EMPIRICAL APPROACH

A: The Leave-Two-Out Procedure of the Cross-Validation Approach[a]	*B: The Five Personality Features Summarizing Brand Associations*[b]	*C: Learned Coefficients*[c]	*D: Schematic with Predicted and Observed fMRI Images for Disney and Gucci*[d]

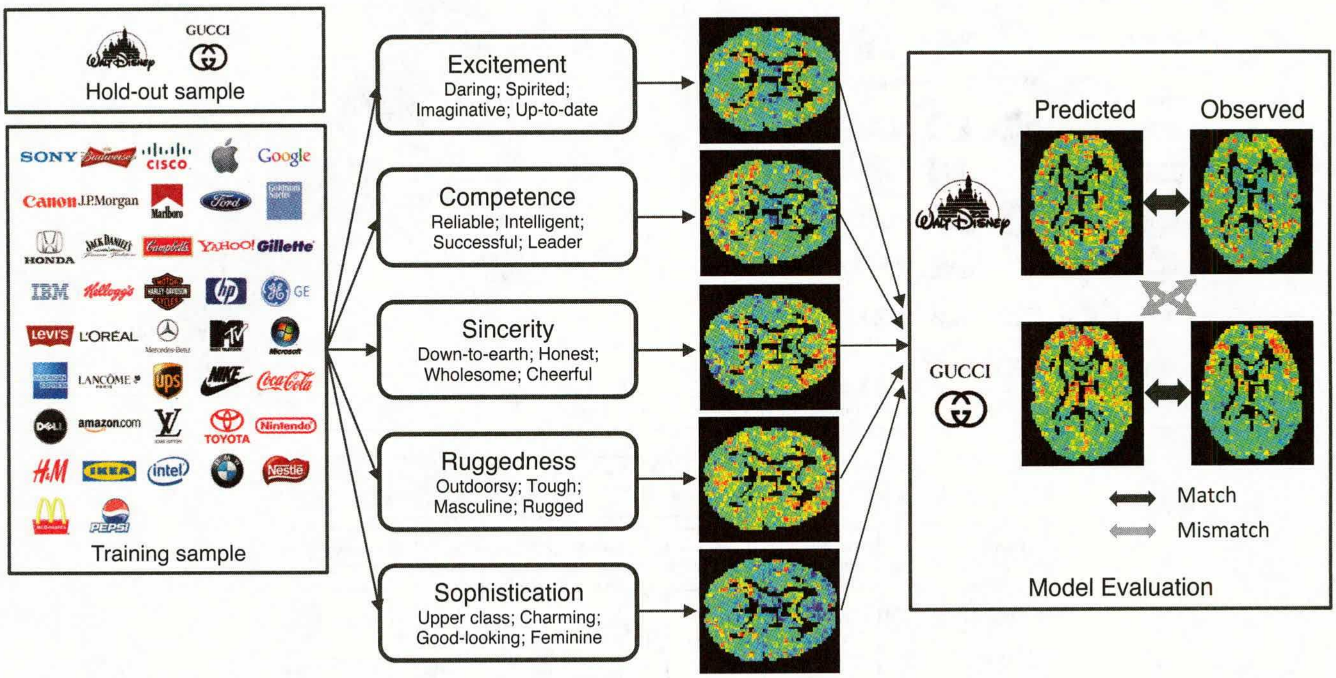

[a]For each iteration, two brands were held out of the training set (e.g., Disney and Gucci), and model calibration was done using the remaining 42 brands in the training set.

[b]Neural signatures of brand association were estimated using brands' personality features derived from participants' ratings.

[c]Coefficients for the five personality features are depicted in single-axial slice, with color representing image intensity.

[d]Cross-validation is completed by using trained neural signatures to predict observed neural responses to holdout brands. The predicted image for the holdout brand is calculated as a linear combination of the personality features of the holdout brands, weighted by the estimated coefficients associated with each feature.

lap in time. We first used a general linear model in SPM8 to estimate a single fMRI image for each of the 176 brand presentations using method LS-S in Mumford et al., whereby each event was modeled as an impulse function convolved with a double gamma hemodynamic function. The beta values estimated for the first regressor of the brand of interest were used as the brain activation patterns associated with a brand at a particular repetition time (for robustness checks using alternative methods of estimating representative fMRI images, see the Web Appendix).

Using brain images for each brand at each repetition time, we standardized the activation levels for each voxel by z-scoring over the 176 files. Then, for each brand, we averaged the four brain images of the four repetition times to obtain the averaged fMRI image associated with thinking about the brand. Finally, we applied the individual gray matter mask to include voxels within the gray matter.

In-sample model training. To infer the engagement of specific mental representations from pattern of neural responses, we took a model-based approach in which the decoding of brain activation patterns is guided by quantitative models capturing psychological features underlying specific mental representations (Mitchell et al. 2008; Naselaris et al. 2011; Poldrack 2011). The underlying hypothesis of our approach is that neural representation of consumer brands is related to the strength of association of an individual brand to its personality features. That is, we assume that neural response in voxel v to brand j is given by

$$y_j^v = c_1^v f_{1,j} + c_2^v f_{2,j} + \ldots + c_n^v f_{n,j}, \tag{1}$$

where $f_{n,j}$ is the value of the nth personality feature for brand j, and c_n^v is a scalar parameter that specifies the degree to which the nth feature activates voxel v. More specifically, c_n^v defines the relationship between the brain activation level and the brand personality features.

We performed model-based decoding using a cross-validation approach in which the model was repeatedly trained using 42 of the 44 available stimulus brands and then tested using the two holdout stimulus brands. We denote the neural response y_j^v in voxel v to brand j as $y_j^v = c_1^v f_{1,j} + c_2^v f_{2,j} + \ldots + c_n^v f_{n,j}$ (Equation 1). We trained the model on each iteration using the set of observed fMRI images associated with 42 known brands to obtain c_n^v values through maximum likelihood. More specifically, we reconstructed the relationship between the brain activation level (as dependent variables) and the brand personality features (as independent variables) with the multiple regression approach, using only 42 of the 44 available stimulus brands. We then tested the model performance on the two holdout brands, which are not in the training set.

Model prediction using holdout sample. After the model was trained, we tested it by presenting the fMRI images (i_1 and i_2) associated with two holdout brands (b_1 and b_2). This process consisted of comparing i_1 and i_2 with the two predicted fMRI images (p_1 and p_2) associated with two holdout brands, where p_1 and p_2 were computed using weights c_n^v and the set of personality features $\{f_{1,k}, ..., f_{n,k}\}$ for the two holdout brands. For example, in an iteration in which Disney and Gucci were excluded from the training, we reconstructed the relationship between the brain activation level and the brand personality features using other 42 brands with Equation 1. Then, using Disney's personality factor scores, we can calculate the predicted activation level for each voxel using Equation 1 and the learned c_n^v values; with these levels, we can create the predicted brain image for Disney. We call the model-predicted brain images p_1 and p_2, and the observed brain images i_1 and i_2, for the two holdout brands.

To evaluate the performance of the model, the model is required to correctly match i_1 and i_2 to b_1 and b_2 using p_1 and p_2, as assessed by which match had a higher correlation value. More specifically, let sel(i) be the vector of values of the selected subset of voxels for image i. We calculated the similarity score between a predicted image, p, and observed image, i, as the Pearson correlation coefficient of the vectors sel(p) and sel(i). The trained model then decided which was a better match—$p_1 = i_1$ and $p_2 = i_2$ or $p_1 = i_2$ and $p_2 = i_1$—by choosing the image pairing with the larger sum of similarity scores. The expected accuracy in matching the two holdout brands to their holdout fMRI images is .50 if the matching is performed at chance levels.

As we described previously, we calculated the similarity between two images using only a subset of the image voxels, following methods proposed in Mitchell et al. (2008). Voxels were selected automatically during training, using only the 42 training brands on each of the leave-two-out cross-validation folds. To select voxels, all voxels were first assigned a stability score using the data from the four presentations of each of the 42 training stimuli. Given these $4 \times 42 = 168$ presentations (168 fMRI images), each voxel was assigned a 4×42 matrix, in which the entry at row i, column j, is the value of this voxel during the ith presentation of the jth brand. We then computed the stability score for this voxel as the average pairwise correlation over all pairs of rows in this matrix. In essence, this calculation assigns highest scores to voxels that exhibit a consistent (across different presentations) variation in activity across the 42 training stimuli (for details, see the Web Appendix).

Significance testing. To calculate statistical significance, we used a permutation procedure to empirically estimate the null distribution (Mitchell et al. 2008). Specifically, we estimated a null model on each iteration by shuffling the fMRI image and brand personality pairing. For example, on a particular iteration, as opposed to using the true brand personality score, we might use Google's personality features to describe Gucci, or IBM to describe Campbell's. Under the null hypothesis that the brand personality framework provides no information about the underlying neural representation, these shuffled brain–brand pairings should yield prediction rates similar to the actual pairings. The null distribution is then calculated using the pooled 600 permuted models from each of the 17 participants, for 10,200 models in total.

BEHAVIORAL RESULTS

Brand Personality Factor Structure

First, we wanted to characterize the set of personality features $f_{n,j}$ associated with our brands using participant ratings of brands on the set of traits outlined in the Aaker (1997) framework (Figure 1, Panel C; see also Table S1 in the Web Appendix). Specifically, we used a factor-analytic approach to summarize variation in trait ratings and reduce collinearity issues. Consistent with previous results, we found that a substantial proportion (86%) of the variance was captured by five factors (Table S2 in the Web Appendix). Further inspection of the factor loadings showed that our results largely replicated those of previous studies (Figure S2 in the Web Appendix) (Aaker 1997). For example, the first factor loaded highly on the traits "trendy," "unique," and "cool"—commonly referred to as the "excitement" factor. The third factor, referred to as "sincerity," loaded highly on traits such as "friendly," "family-oriented," and "down-to-earth." Using this factor-analytic framework, it is possible to characterize each brand (e.g., Apple) as a vector of personality features consisting of these five factors that summarizes the set of characteristics participants associate with these brands (Figure 1, Panel D; Table S3 in the Web Appendix).

Importantly, this association architecture enables us to account for some of the salient similarities and differences between brands apart from their product categories. For example, although Apple and Microsoft reside in the same industry, they elicit highly distinctive associations and are distinguishable in this association architecture (Figure 1, Panel D). In contrast, Disney and IKEA are similar in this framework despite differences in objective features (Figure 1, Panel D). Although this framework by no means captures all characteristics consumers associate with brands, it has been invaluable to researchers by capturing and organizing knowledge in a parsimonious and tractable manner (Aaker 1997).

Robustness of Association Architecture

Furthermore, to investigate the robustness of our framework, as well as the degree to which these trait associations could be generalized to samples from different populations, we surveyed an additional sample of 25 undergraduate students on the same set of traits and brands. We found that the average responses of the trait scores were highly correlated among our neuroimaging participants and the follow-up undergraduate participants (Pearson $r = .86$, $p < 10^{-10}$; Figure 1, Panel E), such that there was considerable agreement between the two samples regarding these brands despite participants' different demographic and socioeconomic characteristics. These results show that this brand personality architecture is considerably robust across samples from different populations, suggesting its utility in organizing the underlying psychological associations.

NEUROIMAGING RESULTS

Brand Personality Traits Can Be Recovered from Brain Activity

Using results from the Aaker (1997) model, we next aim to relate personality factor scores with observed fMRI data associated with viewing brands using a cross-validation approach and test the ability of our framework to discriminate between the previously unseen brands. For each iteration, we held out a different pairing of two brands (e.g., Disney and Gucci) from the training set, and the model was trained using the remaining 42 brands (Figure 2, Panel A). Specifically, training involved regressing the activation level of each voxel on the set of personality features of the training brands obtained from the factor analysis (Figure 2, Panel B). We used the derived maximum likelihood estimates as terms, which we then combined with the personality factor scores of each holdout brand to form a predicted fMRI image. We iterated this leave-two-out train-test procedure 946 times to hold out each of the possible brand pairs (Figure 2, Panel C).

Following training, we evaluated the computational model by comparing these predicted fMRI images with the observed fMRI data of the two holdout brands, evaluated over the 500 image voxels with the most stable responses across training presentations (Figure 2, Panel D). Specifically, given the two holdout brands b_1 and b_2, we calculated their respective predicted images p_1 and p_2 using the set of personality feature $f_{n,j}$ associated with the holdout brands and the set of weights obtained from the training set. Next, using the actual fMRI images i_1 and i_2 associated with the two holdout brands, we asked whether the model was able to correctly match i_1 to p_1 and i_2 to p_2 by choosing the image pairing (i_1 and p_1 vs. i_2 and p_2) that is more highly correlated (Figure 2; for details, see the Web Appendix).

Under the null hypothesis of no association, the predicted fMRI image for a brand will be equally predictive of the matched brand as with the unmatched brand. In contrast, we found that the overall hit rate for iterating over all of the possible combinations of holdout data was 58% and highly significant as assessed using a permutation test obtained by independently training 10,200 single-participant models with randomly shuffled personality features of brands ($p < 10^{-5}$; see the Web Appendix). These results are thus consistent with our hypothesis that brand personality exists in the mind of the consumer a priori (H_1).

Furthermore, we found that the predictive power was strongly modulated by the psychological similarity of brands as measured by correlation of trait ratings. Separating the brand pairs into quartiles on the basis of psychological similarity, we found that performance in classification was substantially better when brands were dissimilar, in which the average hit rate was 63% ($p < 10^{-7}$). In contrast, predictive accuracy was not significantly different from chance when brands were highly similar (Figure 3, Panel A). This modulation of prediction rate by psychological similarity thus argues against the likelihood that our results were driven by some unrelated factors. Moreover, the fact that we were unable to distinguish neural responses to different brands when their personality features were sufficiently similar can be interpreted as a boundary condition in which the brain data no longer contain sufficient resolution to differentiate brand personality representations.

Finally, these results were robust to several variations in specific analytical process, including the method of extracting representative fMRI response to the brands (Figure S7), similarity metric (Figure S8), voxel selection (Figures S9–S10), excluding visual cortex voxels through masking (Fig-

Figure 3

BRAND PERSONALITY TRAITS CAN BE RECOVERED FROM BRAIN ACTIVITY

A: Predictive Accuracy Is Modulated by the Psychological Similarity of Brands[a]

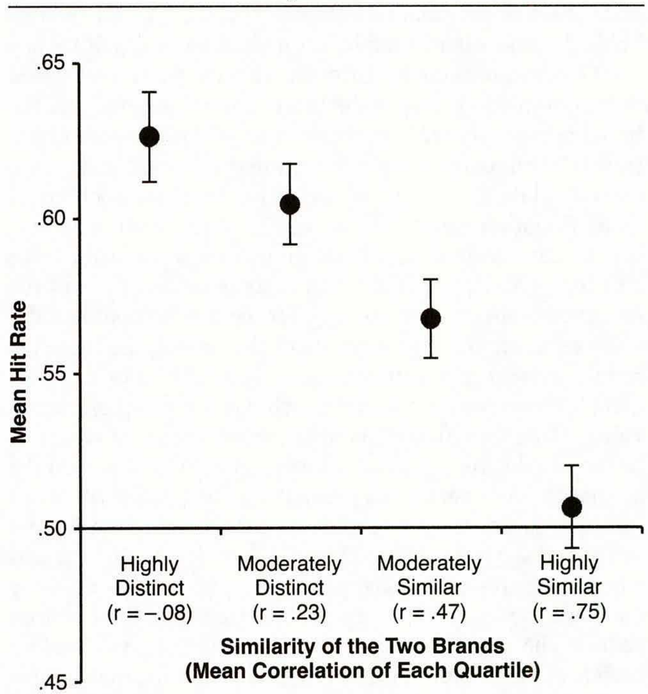

B: Neural Similarity of Brands Is Modulated by Psychological Similarity[b]

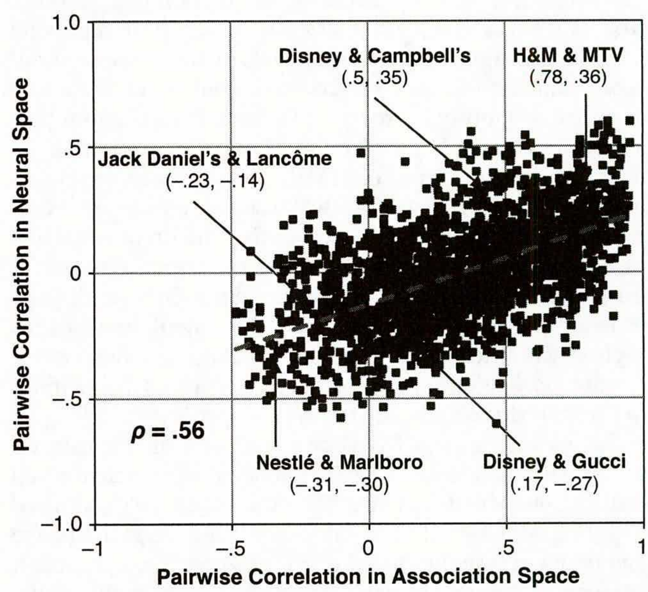

[a]The overall hit rate for holdout classification was 58%. Separating the brands by subjective similarity into quartiles as assessed according to correlation of trait ratings, we find a significant relationship between hit rate and subjective similarity. Error bars indicate standard errors.

[b]Panel B presents, for each brand pair, the correlation between predicted and observed brain images evaluated over the 500 image voxels with the most stable responses across training presentations (y-axis) against similarity in brands' psychological properties as measured using correlation of trait ratings (x-axis).

ure S11), and controlling for physical properties of brand logos (Figure S12; for details, see the Web Appendix).

Neural Similarity of Brands Is Modulated by Psychological Similarity

To examine the relationship between the psychological organization of brands and the discriminability of the associated brain images more systematically, we compared, for each brand pair, the correlation between predicted and observed brain images, evaluated over the 500 image voxels with the most stable responses across training presentations, against psychological similarity in brand meaning as measured by correlation of trait ratings (Figure 3, Panel B). We found that strength of neural correlation is robustly modulated by the similarity of brands' psychological properties (Pearson $r = .56, p < 10^{-7}$), such that brands that were more similar at the psychological level were also more highly correlated at the neural level (Figure 3, Panel B). For example, H&M and MTV are highly similar in their psychological associations as measured using a correlation index (Pearson $r = .78$), whereas those for Disney and Gucci are highly distinct (Pearson $r = .17$) (see Figure S3 and Table S3 in the Web Appendix). Consistent with this pattern, neural signatures associated with H&M are more similar to those associated with MTV than Disney with Gucci (Pearson $r = .36$ vs. $r = -.27$, respectively). We obtained similar results using Euclidean distance as a measure of similarity (see Figure S7 in the Web Appendix). These results underscore the notion that the brand personality framework provides a reasonable first-order approximation of the mental representation, consistent with our Assumption 2.

Brand Personality Contents Are Distributed Widely Across the Brain

Having assessed the predictive validity of our decoding framework, we aimed to characterize the set of brain regions where predicted neural response for holdout brands best correlated with the observed responses. To do so, we calculated the correlation coefficient of the predicted and observed fMRI response at each voxel location and selected the set of regions where brain activity was significantly correlated with model predictions (see the Web Appendix). Consistent with connectionist models of distributed representation (H_2), we found that the set of predictive voxels were distributed throughout the brain (Figure 4 and Table 1; see also Figures S6 and S13–S17 in the Web Appendix). In contrast, these regions are not visible using a standard univariate generalized linear modeling approach that ignores information contained in the spatially distributed set of brain regions (Figure S18 in the Web Appendix).

To understand the cognitive functions in which these regions were most involved, we conducted an exploratory reverse-inference analysis using Neurosynth (Yarkoni et al. 2011), correlating our activation map with the neural activation maps for each term in the Neurosynth database (Figure 4). We found that our activations were distributed across several types of cognitive functions but particularly those implicated in previous studies of semantic knowledge (inferior frontal gyrus), imagery (premotor and visual cortex), and emotional processing (anterior and posterior cingulate gyrus), consistent with the notion that brand knowledge consists of a complex mix of thoughts, images, and feelings that consumers associate with brands.

DISCUSSION

The application of neuroscientific methods to marketing has a history that is brief in existence but long on controversy (Ariely and Berns 2010; Plassmann, Ramsøy, and Milosavljevic 2012). In a particularly high-profile incident, the *New York Times* published an op-ed titled "You Love Your iPhone. Literally" by the brand consultant Martin Lindstrom (2011), which prompted a group of 44 neuroscientists to cosign a response letter condemning the article. Whatever the scientific merits of the claims (and indeed, the data have never appeared in a peer-reviewed format), at the heart of the study lies a set of questions of great interest to marketers, consumer researchers, and the lay public alike. Namely, what is the set of thoughts and feelings that occur when people think or interact with the products that they own or are considering purchasing?

Here, we take an important step toward bridging this gap and begin to provide a neuroscientific framework to address these questions. More specifically, using a decoding approach in conjunction with factor-analytic techniques, we formally test our ability to infer mental representations of brands using a set of intermediate psychological features to model the underlying representational space (Haynes and Rees 2006; Mitchell et al. 2008; Norman et al. 2006). Compared with the "where"-type questions that are the focus of traditional localization approaches, these "what"-type questions have become addressable only in recent years (Haynes and Rees 2006; Mitchell et al. 2008; Norman et al. 2006) and, to our knowledge, have not been addressed in consumer neuroscience.

First, consistent with our hypothesis that brand personality traits exist a priori inside the mind of the consumer (H_1), we found that we were able to predict what brand consumers were thinking about solely on the basis of the relationship between brand personality and brain activity. In particular, because participants in our study were not prompted on traits such as "daring," "reliable," and "wholesome" until after the scanning session, our likelihood of predicting what brands participants are thinking of should be at chance if such associations did not come across the consumers' thoughts. In contrast, previous studies have typically elicited subjective ratings online during scanning (Schaefer et al. 2006; Schaefer and Rotte 2010; Yoon et al. 2006), thereby leaving open the possibility that brand-related processing was at least in part induced by the specific stimuli used during the experiment.

Moreover, although the reported predictive accuracy rates were lower than rates observed in more basic perceptual domains (Haxby et al. 2001; Kay et al. 2008), they are comparable to those observed in previous studies of higher-level cognitive processes, including those involving consumer choice (Knutson et al. 2007; Van der Laan et al. 2012), some of which may be attributable to our decision not to include a fixation screen after every brand logo presentation. We did so because pilot participants stated that they found the number of fixation screens between brands to interfere with their ability to process brand traits; however, we acknowledge that this may have resulted in reduced efficiency in extrac-

Figure 4

BRAND PERSONALITY CONTENTS ARE DISTRIBUTED WIDELY ACROSS THE BRAIN

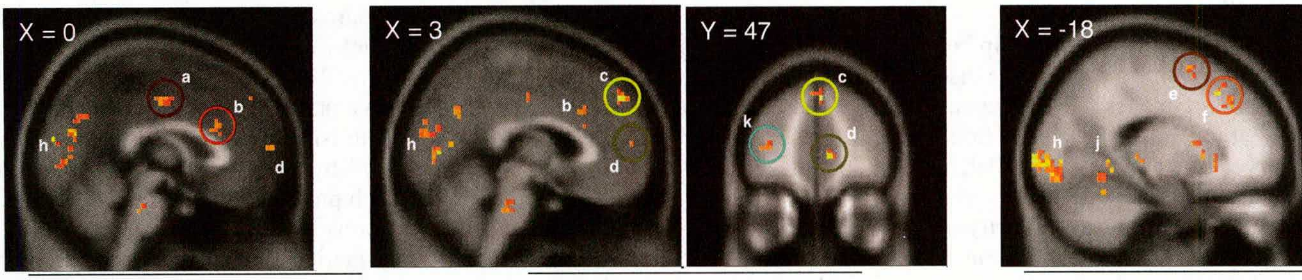

Region/Feature	*Probability*
a *Anterior Cingulate*	
Experiencing	.85
Empathetic	.79
b *Middle cingulate*	
Autobiographical memory	.87
Perspective	.81

Region/Feature	*Probability*
c *Dorsomedial prefrontal*	
Personality traits	.85
Interpersonal	.81
d *Medial prefrontal*	
Autobiographical memory	.87
Person	.79

Region/Feature	*Probability*
e *Premotor*	
Abstract	.81
Imagery	.77
f *Dorsolateral prefrontal*	
Familiarity	.83
Recollection	.82

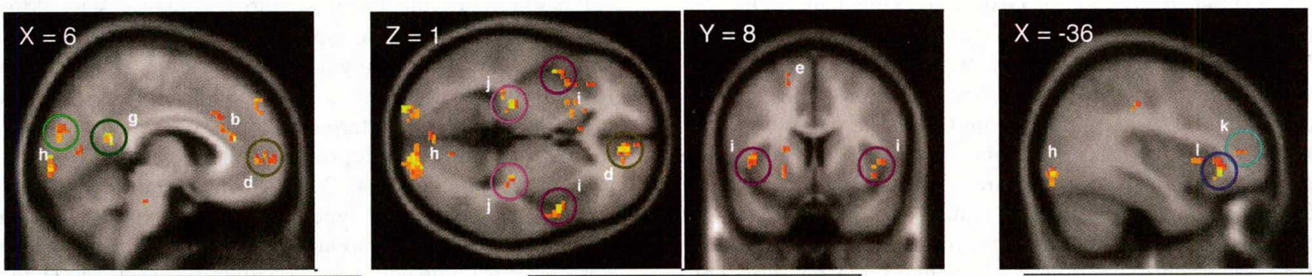

Region/Feature	*Probability*
g *Posterior cingulate*	
Autobiographical	.89
Memories	.79
h *Primary visual*	
Mental imagery	.86
Visual perception	.82

Region/Feature	*Probability*
i *Insula*	
Sensation	.80
Effective	.64
j *Hippocampus*	
Thoughts	.86
Memory	.71

Region/Feature	*Probability*
k *Lateral prefrontal*	
Memory tasks	.84
Retrieval	.73
l *Inferior frontal*	
Semantic	.72
Words	.68

Notes: We show the slice view of the most accurately predicted voxels (i.e., voxels with the highest correlation between out-of-sample prediction rates and actual activations for the average participant). Each panel shows clusters containing at least ten contiguous voxels for which predicted–actual correlation is significantly greater than zero, with $p < .05$ from the permutation test (Table 1). To make inferences about cognitive processes subserved by these regions, we used the meta-analytic tool Neurosynth (Yarkoni et al. 2011) to generate the probability that a specific cognitive process is engaged given activation in a particular brain region. For example, given specific voxel location of the observed activation in the dorsomedial prefrontal cortex (Cluster C), there is a .85 probability that the term "personality traits" was used in a study given the presence of reported activation.

tion of the representative brand fMRI image. Future studies are needed to address the extent to which predictive accuracy can be improved.

Second, we found that neural responses to consumer brands can be decomposed into a basic set of neural activation patterns associated with intangible characteristics of these objects and that these results were robust to several variations in the specific analytical process (see the "Supplementary Results" section and Figures S7–S12 in the Web Appendix). Moreover, our findings are consistent with connectionist models of conceptual knowledge in which brand personality associations emerge from weighted activity across a distributed set of units (H_2) (Binder et al. 2009; Tyler and Moss 2001). That is, with regard to the contentful associations that distinguish one brand from another, the underlying neural representations seem to be akin to previous distributed accounts of conceptual knowledge (Binder et al. 2009; Tyler and Moss 2001) reflecting the complex array of cognitive processes that are engaged.

Notably, within this distributed set of brain regions, we found brand personality contents present in both mPFC and lPFC regions (Figure 4). On the surface, that we found brand personality contents in mPFC regions may seem at odds with previous findings in Yoon et al. (2006) that mPFC activity is lower during brand processing than person processing. Both sets of findings, however, are consistent with the notion that the mPFC exhibits a gradation of activation levels in person judgment tasks. That is, as opposed to "all-or-none" activation, the mPFC has been previously shown to exhibit lower activity in judgment of out-group individuals relative to in-group individuals (Volz, Kessler, and Von Cramon 2009) and in judgment of more dissimilar individuals relative to more similar individuals (Mitchell, Macrae, and Banaji 2006). Under this interpretation, reduced mPFC

Table 1

VOXEL LOCATIONS OF BRAIN REGIONS WHERE PREDICTED NEURAL RESPONSE FOR HELD-OUT BRANDS WERE SIGNIFICANTLY CORRELATED WITH THE OBSERVED NEURAL RESPONSES

Cluster		*Voxel*[c]				
Size[a]	*Correlation Coefficient*[b]	*X*	*Y*	*Z*	*L/R*[d]	*Region*
184	.65	18	−94	−5	R	Lingual gyrus
11	.63	−12	38	55	L	Superior frontal gyrus
15	.60	51	11	−8	R	Superior temporal gyrus
23	.57	6	−52	16	R	Posterior cingulate
145	.55	−12	−97	−8	L	Lingual gyrus
36	.54	6	35	16	R	Anterior cingulate
17	.53	3	47	40	R	Medial frontal gyrus
15	.50	−18	26	43	L	Superior frontal gyrus
10	.49	36	−34	−2	R	Subgyral
14	.48	−21	11	58	L	Middle frontal gyrus
14	.47	−45	2	1	L	Insula
16	.47	−3	−7	43	L	Cingulate gyrus
23	.46	51	2	−2	R	Superior temporal gyrus
14	.46	−36	29	−8	L	Inferior frontal gyrus
12	.46	−9	26	28	L	Cingulate gyrus
11	.45	21	−37	−5	R	Parahippocampal gyrus
26	.44	9	47	1	R	Medial frontal gyrus
25	.43	3	−79	4	R	Lingual gyrus
32	.42	−3	−79	22	L	Cuneus
13	.42	−33	53	13	L	Superior frontal gyrus
14	.40	27	41	31	R	Superior frontal gyrus
28	.39	−12	26	−5	L	Caudate
10	.37	3	−64	28	R	Precuneus

[a]Cluster size (voxels).
[b]Correlation coefficient between the predicted and the observed brain images.
[c]Voxel location (X, Y, Z) in Montreal Neurological Institute coordinate (mm).
[d]Laterality of activation (L = left hemisphere, R = right hemisphere).

activation reflects the notion that brand judgment only weakly draws on anthropomorphic features and processes. An alternative possible explanation is that these two studies engage fundamentally different aspects of mPFC functioning. For example, whereas locally distributed response patterns in the mPFC reflect brand personality, mean response differences in the mPFC may instead reflect some other process that is known to engage mPFC—for example, valuation processes widely observed in neuroeconomic studies (Plassmann et al. 2008; Rangel, Camerer, and Montague 2008). Indeed, this is a general limitation in exploratory reverse inferences, including those using probabilistic meta-analytic techniques such as Neurosynth (Yarkoni et al. 2011). Future studies combining the approach outlined in the current study and that of Yoon et al. (2006) are needed to address these issues.

More generally, the methods we outline herein enable consumer researchers to consider a set of research questions not previously testable that center on the idea that spatially distributed fMRI activity patterns may represent a viable signature of hypothesized psychological constructs (Haynes and Rees 2006; Naselaris et al. 2011). This includes, for example, cases in which self-reported perceptions or preferences may be compromised due to factors such as social desirability bias. Existing efforts to control for such biases have largely consisted of randomized response protocols (De Jong, Pieters, and Fox 2010; Warner 1965). These protocols reduce privacy concerns by using a randomization mechanism to "shroud" the participant's response, and they rely on the credibility of the randomization device and feelings of privacy, which have been challenged in recent years (Chaudhuri and Christofides 2013). In contrast, by eliciting neural responses without any overt behavior, passive viewing experiments such as those used in the current study may be able to overcome some of these challenges.

With respect to branding, capturing the mental map of brand personality opens the door for studies addressing several additional questions of interest to consumer researchers and marketers. In particular, by capturing and validating brand personality representations in the brain, a natural next step is to characterize how marketing actions affect these representations and investigate the different cognitive processes that act on these representations. This parallels the trajectory of findings in more basic psychological processes such as working memory, in which discovering the existence of visual working memory contents in extrastriate regions enabled researchers to ask several questions regarding how these representations were affected under different task demands (Chadwick et al. 2010; Lee, Kravitz, and Baker 2013). For example, Lee, Kravitz, and Baker (2013) find that information about object identity was contained in different brain regions depending on whether participants were asked to attend to visual or nonvisual properties of the object.

One set of questions along these lines involves comparing different dimensions of brand knowledge, such as brand experience and brand relationships, as well as how these representations differ across consumer segments. Intuitively, whereas brand personality captures traits that consumers project onto brands (Aaker 1997), brand experience captures responses that brands evoke on the part of consumers (Brakus, Schmitt, and Zarantonello 2009), and brand relationships capture feelings and episodes that consumers have actually experienced with the brands (Fournier 1998). Moreover, research has shown these associations to differ in important ways across segments such as cultural background (Aaker, Benet-Martínez, and Garolera 2001). Therefore, it may be that these constructs are subserved by different mental processes and differ across segments, which would have implications for brand managers in designing marketing activity that can create or affect these dimensions of brand knowledge.

Finally, future studies extending our approach could begin to quantify the extent to which marketing actions affect consumers' mental representations of brand personality, a question of clear interest to brand managers. In our current study, we explicitly assumed that activation patterns elicited by brands remain constant across different repetitions. Although this assumption is likely safe given that our stimuli contained some of the most iconic brands in the world, it limited our ability to make inferences on how brand associations and values are acquired and how they evolve over time (Johar, Sengupta, and Aaker 2005; Van Osselaer and Janiszewski 2001). Future studies combining our approach with dynamic models of inference updating could therefore begin to trace out the processes by which marketing actions affect multiple dimensions of brand knowledge and preference.

REFERENCES

Aaker, David A. (2009), *Managing Brand Equity*. New York: Simon & Schuster.

—— (2012), *Building Strong Brands*. New York: Simon & Schuster.

Aaker, Jennifer L. (1997), "Dimensions of Brand Personality," *Journal of Marketing Research*, 34 (July), 347–56.

——, Verónica Benet-Martínez, and Jordi Garolera (2001), "Consumption Symbols as Carriers of Culture: A Study of Japanese and Spanish Brand Personality Constructs," *Journal of Personality and Social Psychology*, 81 (3), 492–508.

Alba, Joseph W. and J. Wesley Hutchinson (1987), "Dimensions of Consumer Expertise," *Journal of Consumer Research*, 13 (4), 411–54.

Ariely, Dan and Gregory S. Berns (2010), "Neuromarketing: The Hope and Hype of Neuroimaging in Business," *Nature Reviews Neuroscience*, 11 (4), 284–92.

Batra, Rajeev, Peter Lenk, and Michel Wedel (2010), "Brand Extension Strategy Planning: Empirical Estimation of Brand–Category Personality Fit and Atypicality," *Journal of Marketing Research*, 47 (April), 335–47.

Bettman, James R. (1970), "Information Processing Models of Consumer Behavior," *Journal of Marketing Research*, 7 (August), 370–76.

Binder, Jeffrey R., Rutvik H. Desai, William W. Graves, and Lisa L. Conant (2009), "Where Is the Semantic System? A Critical Review and Meta-Analysis of 120 Functional Neuroimaging Studies," *Cerebral Cortex*, 19 (12), 2767–96.

Brakus, J. Joško, Bernd H. Schmitt, and Lia Zarantonello (2009), "Brand Experience: What Is It? How Is It Measured? Does It Affect Loyalty?" *Journal of Marketing*, 73 (May), 52–68.

Bruce, Amanda S., Rebecca J. Lepping, Jared M. Bruce, J. Bradley C. Cherry, Laura E. Martin, Ann M. Davis, et al. (2013), "Brain Responses to Food Logos in Obese and Healthy Weight Children," *Journal of Pediatrics*, 162 (4), 759–64.

Chadwick, Martin J., Demis Hassabis, Nikolaus Weiskopf, and Eleanor A. Maguire (2010), "Decoding Individual Episodic Memory Traces in the Human Hippocampus," *Current Biology*, 20 (6), 544–47.

Chaudhuri, Arijit and Tasos C. Christofides (2013), "Indirect Techniques as Alternatives to Randomized Response," in *Indirect Questioning in Sample Surveys*. Berlin: Springer, 115–49.

Churchland, Patricia S. and Terrence J. Sejnowski (1988), "Perspectives on Cognitive Neuroscience," *Science*, 242 (4879), 741–45.

Clithero, John A., R. McKell Carter, and Scott A. Huettel (2009), "Local Pattern Classification Differentiates Processes of Economic Valuation," *NeuroImage*, 45 (4), 1329–38.

De Jong, Martijn G., Rik Pieters, and Jean-Paul Fox (2010), "Reducing Social Desirability Bias Through Item Randomized Response: An Application to Measure Underreported Desires," *Journal of Marketing Research*, 47 (February), 14–27.

Deppe, Michael, Wolfram Schwindt, Harald Kugel, Hilke Plassmann, and Peter Kenning (2005), "Nonlinear Responses Within the Medial Prefrontal Cortex Reveal When Specific Implicit Information Influences Economic Decision Making," *Journal of Neuroimaging*, 15 (2), 171–82.

De Quervain, Dominique J.F., Urs Fischbacher, Valerie Treyer, Melanie Schellhammer, Ulrich Schnyder, Alfred Buck, et al. (2004), "The Neural Basis of Altruistic Punishment," *Science*, 305 (5688), 1254–58.

Esch, Franz-Rudolf, Thorsten Möll, Bernd Schmitt, Christian E. Elger, Carolin Neuhaus, and Bernd Weber (2012), "Brands on the Brain: Do Consumers Use Declarative Information or Experienced Emotions to Evaluate Brands?" *Journal of Consumer Psychology*, 22 (1), 75–85.

Formisano, Elia, Federico De Martino, Milene Bonte, and Rainer Goebel (2008), "'Who' Is Saying 'What'? Brain-Based Decoding of Human Voice and Speech," *Science*, 322 (5903), 970–73.

Fournier, Susan (1998), "Consumers and Their Brands: Developing Relationship Theory in Consumer Research," *Journal of Consumer Research*, 24 (4), 343–53.

Gardner, Burleigh B. and Sidney J. Levy (1955), "The Product and the Brand," *Harvard Business Review*, 33 (2), 33–39.

Gazzaniga, Michael S. (2004), *The Cognitive Neurosciences*. Cambridge, MA: MIT Press.

Haire, Mason (1950), "Projective Techniques in Marketing Research," *Journal of Marketing*, 14 (April), 649–56.

Haxby, James V., M. Ida Gobbini, Maura L. Furey, Alumit Ishai, Jennifer L. Schouten, and Pietro Pietrini (2001), "Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex," *Science*, 293 (5539), 2425–30.

Haynes, John-Dylan and Geraint Rees (2005), "Predicting the Orientation of Invisible Stimuli from Activity in Human Primary Visual Cortex," *Nature Neuroscience*, 8, 686–91.

—— and —— (2006), "Decoding Mental States from Brain Activity In Humans," *Nature Reviews Neuroscience*, 7 (7), 523–34.

Hedgcock, William and Akshay R. Rao (2009), "Trade-Off Aversion as an Explanation for the Attraction Effect: A Functional Magnetic Resonance Imaging Study," *Journal of Marketing Research*, 46 (February), 1–13.

Hinton, G.E., James L. McClelland, and David E. Rumelhart (1986), "Distributed Representations," in *Parallel Distributed Processing*, David E. Rumelhart and James L. McClelland, eds. Cambridge, MA: MIT Press, 77–109.

Huettel, Scott A., John W. Payne, Carolyn Yoon, Richard Gonzalez, James R Bettman, William Hedgcock, et al. (2009), "Commentaries and Rejoinder to 'Trade-Off Aversion as an Explanation for the Attraction Effect: A Functional Magnetic Resonance Imaging Study,'" *Journal of Marketing Research*, 46 (February), 14–24.

Johar, Gita Venkataramani, Durairaj Maheswaran, and Laura A. Peracchio (2006), "MAPping the Frontiers: Theoretical Advances in Consumer Research on Memory, Affect, and Persuasion," *Journal of Consumer Research*, 33 (1), 139–49.

———, Jaideep Sengupta, and Jennifer L. Aaker (2005), "Two Roads to Updating Brand Personality Impressions: Trait Versus Evaluative Inferencing," *Journal of Marketing Research*, 42 (November), 458–69.

Kay, Kendrick N., Thomas Naselaris, Ryan J. Prenger, and Jack L. Gallant (2008), "Identifying Natural Images from Human Brain Activity," *Nature*, 452 (7185), 352–55.

Keller, Kevin Lane (1993), "Conceptualizing, Measuring, and Managing Customer-Based Brand Equity," *Journal of Marketing*, 57 (January), 1–22.

——— (2003), "Brand Synthesis: The Multidimensionality of Brand Knowledge," *Journal of Consumer Research*, 29 (4), 595–600.

——— and Donald R. Lehmann (2003), "How Do Brands Create Value?" *Marketing Management*, 12 (3), 26–31.

Klucharev, Vasily, Ale Smidts, and Guillén Fernández (2008), "Brain Mechanisms of Persuasion: How 'Expert Power' Modulates Memory and Attitudes," *Social Cognitive and Affective Neuroscience*, 3 (4), 353–66.

Knutson, Brian, Scott Rick, G. Elliott Wimmer, Drazen Prelec, and George Loewenstein (2007), "Neural Predictors of Purchases," *Neuron*, 53 (1), 147–56.

Kriegeskorte, Nikolaus, Rainer Goebel, and Peter Bandettini (2006), "Information-Based Functional Brain Mapping," *Proceedings of the National Academy of Sciences of the United States of America*, 103 (10), 3863–68.

———, Marieke Mur, Douglas A. Ruff, Roozbeh Kiani, Jerzy Bodurka, Hossein Esteky, et al. (2008), "Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey," *Neuron*, 60 (6), 1126–41.

Lashley, Karl S. (1950), "In Search of the Engram," in *Society of Experimental Biology Symposium No. 4: Physiological Mechanisms in Animal Behavior*. Oxford, UK: Academic Press, 454–82.

Lee, Sue-Hyun, Dwight J. Kravitz, and Chris I. Baker (2013), "Goal-Dependent Dissociation of Visual and Prefrontal Cortices During Working Memory," *Nature Neuroscience*, 16, 997–99.

Levy, Sidney J. (1959), "Symbols for Sale," *Harvard Business Review*, 37, 117–24.

Lindstrom, Martin (2011), "You Love Your iPhone. Literally," *The New York Times*, (September 30), (accessed May 11, 2015), [available at http://www.nytimes.com/2011/10/01/opinion/you-love-your-iphone-literally.html].

McClelland, James L. and Timothy T. Rogers (2003), "The Parallel Distributed Processing Approach to Semantic Cognition," *Nature Reviews Neuroscience*, 4 (4), 310–22.

McClure, Samuel M., Jian Li, Damon Tomlin, Kim S. Cypert, Latané M. Montague, and P. Read Montague (2004), "Neural Correlates of Behavioral Preference for Culturally Familiar Drinks," *Neuron*, 44 (2), 379–87.

Mitchell, Jason P., C. Neil Macrae, and Mahzarin R. Banaji (2006), "Dissociable Medial Prefrontal Contributions to Judgments of Similar and Dissimilar Others," *Neuron*, 50 (4), 655–63.

Mitchell, Tom M., Svetlana V. Shinkareva, Andrew Carlson, Kai-Min Chang, Vicente L. Malave, Robert A. Mason, et al. (2008), "Predicting Human Brain Activity Associated with the Meanings of Nouns," *Science*, 320 (5880), 1191–95.

Mumford, Jeanette A., Benjamin O. Turner, F. Gregory Ashby, and Russell A. Poldrack (2012), "Deconvolving BOLD Activation in Event-Related Designs for Multivoxel Pattern Classification Analyses," *NeuroImage*, 59 (3), 2636–43.

Murawski, Carsten, Philip G. Harris, Stefan Bode, Juan F. Dominguez, and Gary F. Egan (2012), "Led into Temptation? Rewarding Brand Logos Bias the Neural Encoding of Incidental Economic Decisions," *PLOS One*, 7 (3), e34155.

Naselaris, Thomas, Kendrick N. Kay, Shinji Nishimoto, and Jack L. Gallant (2011), "Encoding and Decoding in fMRI," *NeuroImage*, 56 (2), 400–410.

———, Ryan J. Prenger, Kendrick N. Kay, Michael Oliver, and Jack L. Gallant (2009), "Bayesian Reconstruction of Natural Images from Human Brain Activity," *Neuron*, 63 (6), 902–915.

Norman, Kenneth A., Sean M. Polyn, Greg J. Detre, and James V. Haxby (2006), "Beyond Mind-Reading: Multi-Voxel Pattern Analysis of fMRI Data," *Trends in Cognitive Sciences*, 10 (9), 424–30.

Plassmann, Hilke, John O'Doherty, Baba Shiv, and Antonio Rangel (2008), "Marketing Actions Can Modulate Neural Representations of Experienced Pleasantness," *Proceedings of the National Academy of Sciences of the United States of America*, 105 (3), 1050–54.

———, Thomas Zoëga Ramsøy, and Milica Milosavljevic (2012), "Branding the Brain: A Critical Review and Outlook," *Journal of Consumer Psychology*, 22 (1), 18–36.

Poldrack, Russell A. (2011), "Inferring Mental States from Neuroimaging Data: From Reverse Inference to Large-Scale Decoding," *Neuron*, 72 (5), 692–97.

Rangel, Antonio, Colin Camerer, and P. Read Montague (2008), "A Framework for Studying the Neurobiology of Value-Based Decision Making," *Nature Reviews Neuroscience*, 9 (7), 545–56.

Rissman, Jesse and Anthony D. Wagner (2012), "Distributed Representations in Memory: Insights from Functional Brain Imaging," *Annual Review of Psychology*, 63, 101–128.

Schacter, Daniel L. and Anthony D. Wagner (1999), "Medial Temporal Lobe Activations in fMRI and PET Studies of Episodic Encoding and Retrieval," *Hippocampus*, 9 (1), 7–24.

Schaefer, Michael, Harald Berens, Hans-Jochen Heinze, and Michael Rotte (2006), "Neural Correlates of Culturally Familiar Brands of Car Manufacturers," *NeuroImage*, 31 (2), 861–65.

——— and Michael Rotte (2010), "Combining a Semantic Differential with fMRI to Investigate Brands as Cultural Symbols," *Social Cognitive and Affective Neuroscience*, 5 (2/3), 274–81.

Tusche, Anita, Stefan Bode, and John-Dylan Haynes (2010), "Neural Responses to Unattended Products Predict Later Consumer Choices," *Journal of Neuroscience*, 30 (23), 8024–31.

Tyler, Lorraine K. and Helen E. Moss (2001), "Towards a Distributed Account of Conceptual Knowledge," *Trends in Cognitive Sciences*, 5 (6), 244–52.

Van der Laan, Laura N., Denise T.D. de Ridder, Max A. Viergever, and Paul A.M. Smeets (2012), "Appearance Matters: Neural Correlates of Food Choice and Packaging Aesthetics," *PLOS ONE*, 7 (7), e41738.

Van der Lans, Ralf, Bram Van den Bergh, and Evelien Dieleman (2014), "Partner Selection in Brand Alliances: An Empirical Investigation of the Drivers of Brand Fit," *Marketing Science*, 33 (4), 551–66.

Van Osselaer, Stijn M.J. and Chris Janiszewski (2001), "Two Ways of Learning Brand Associations," *Journal of Consumer Research*, 28 (2), 202–223.

Volz, Kirsten G., Thomas Kessler, and D. Yves von Cramon (2009), "In-Group as Part of the Self: In-Group Favoritism Is Mediated by Medial Prefrontal Cortex Activation," *Social Neuroscience*, 4 (3), 244–60.

Warner, Stanley L. (1965), "Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias," *Journal of the American Statistical Association*, 60 (309), 63–66.

Yarkoni, Tal, Russell A. Poldrack, Thomas E. Nichols, David C. Van Essen, and Tor D. Wager (2011), "Large-Scale Automated Synthesis of Human Functional Neuroimaging Data," *Nature Methods*, 8 (8), 665–70.

Yoon, Carolyn, Richard Gonzalez, Antoine Bechara, Gregory S. Berns, Alain A. Dagher, Laurette Dubé, et al. (2012), "Decision Neuroscience and Consumer Decision Making," *Marketing Letters*, 23 (2), 473–85.

——, Angela H. Gutchess, Fred Feinberg, and Thad A. Polk (2006), "A Functional Magnetic Resonance Imaging Study of Neural Dissociations Between Brand and Person Judgments," *Journal of Consumer Research*, 33 (1), 31–40.

Zaltman, Gerald (1997), "Rethinking Market Research: Putting People Back In," *Journal of Marketing Research*, 34 (November), 424–37.

—— and Robin Higie Coulter (1995), "Seeing the Voice of the Customer: Metaphor-Based Advertising Research," *Journal of Advertising Research*, 35 (4), 35–51.