Title
Essays on Consumer Cognition in a Dynamic Marketplace

ISBN
9798189277474

Author
Chen, Vincent Pei-Ming

Publication Date
2026-05-01

Peer reviewed|Thesis/dissertation

Essays on Consumer Cognition in a Dynamic Marketplace

By

Vincent (Pei-Ming) Chen

A dissertation submitted in partial satisfaction of the

requirement for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor Ming Hsu, Chair

Professor Ellen Evers

Professor Heather Haveman

Spring 2026

Essays on Consumer Cognition in a Dynamic Marketplace (Berkeley Edition)

By

Vincent (Pei-Ming) Chen

Abstract

Essays on Consumer Cognition in a Dynamic Marketplace

by

Vincent (Pei-Ming) Chen

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Ming Hsu, Chair

Marketplaces are becoming increasingly dynamic and fast-changing. Cognition plays a central role in how consumers make judgments and decisions in such environments, particularly through processes such as knowledge and memory. However, existing theories of consumer cognition, while conceptually rich, are often insufficiently formalized to generate novel and falsifiable predictions.In this dissertation, I examine three interconnected questions: how brand meanings evolve over time, how similarity and categorization influence consumer spending, and how false memories affect brand judgments. Although these topics span different areas of consumer behavior, they are unified by a common focus on how consumers' mental representations shape judgment and decision-making. Across these projects, I advance a cognitive perspective in consumer research by formalizing previously loosely specified theories—such as those concerning brand meaning and mental accounting—and deriving precise, testable predictions. In doing so, this dissertation contributes to a more rigorous and unified understanding of consumer cognition in dynamic environments.

Introduction

Consumers rarely make decisions based on complete and objective information (Kahneman, 2003; Simon, 1955). Instead, they rely on internal mental representations that summarize products, brands, and their relationships to guide judgment and choice (Hayes-Roth, 1977; Keller, 1993). For example, consumers draw on brand meanings when deciding which brands to purchase (Keller, 1993), use memory to retrieve relevant product information (Nedungadi, 1990), and form mental accounts to reduce cognitive load when evaluating similar products (R. Thaler, 1985).

Importantly, these representations are constructed from past experiences and are often imperfect: they may be distorted (Loftus, 2005), context-dependent (Tversky & Kahneman, 1981), and subject to change over time (Alba & Hutchinson, 1987a; Bettman et al., 1998). Consumers may confuse memoryt of advertisements with their own experiences (Braun, 1999a), or treat objectively equivalent sources of money as belonging to distinct, non-exchangeable mental accounts (R. H. Thaler, 1999a).

While each of these components has long been acknowledged in marketing research (D. A. Aaker, 1991; Keller, 1993; R. Thaler, 1985), they are often conceptualized in ways that lack precise structure, making it difficult to generate clear and testable predictions (Alba & Hutchinson, 1987b). For instance, brand meanings are frequently described in loose, metaphorical terms (J. L. Aaker, 1997b; Fournier, 1998), for example, statements such as "Nike has been defining what it means to be cool." Although it is widely recognized that understanding and shaping brand meaning is important for marketers (D. A. Aaker, 1991; Keller, 1993), relatively few studies have measured this construct quantitatively or examined its behavioral consequences in a systematic way (T. Y. Lee & Bradlow, 2011; Netzer et al., 2011).

Similarly, mental accounting has been widely used to explain a variety of consumer behaviors, such as how consumers allocate spending across different sources of money (Shefrin & Thaler, 1988; R. Thaler, 1985; R. H. Thaler, 1999a). However, existing theories provide limited guidance on how these mental categories are formed or what constrains their boundaries, making them difficult to test rigorously (Cheema & Soman, 2006; Milkman & Beshears, 2009). In the absence of such constraints, interpretations can become flexible: outcomes are often explained post hoc by assigning items to plausible mental accounts. For example, if consumers spend a zoo trip rebate on a skiing trip, it may be attributed to an "entertainment" account; if they instead donate it to the Wildlife Conservation Society, it may be attributed to an "animals" account.

In comparison, consumer memory has been more extensively studied than the other two constructs (E.g., Lynch & Srull, 1982). Nevertheless, important aspects of consumer memory remain underexplored. In particular, prior research has focused primarily on memory for the presence of events (Dick et al., 1990; Janiszewski et al., 2003; Nedungadi, 1990), for example, recalling whether consumers have seen an advertisement of a brand (Burke & Srull, 1988), while relatively less attention has been paid to memory for the absence of events, such as recalling which brands were not involved in a scandal. Understanding such asymmetries is important, as consumers often rely on both types of memory when making decisions.

To address these limitations, a more precise representation of consumer cognition is needed. In response, this dissertation develops computational frameworks for studying how consumers' mental structures can affect their decision-making processes. Rather than modeling consumer decisions purely as functions of stimuli, I model them as arising from operations on consumers' mental representations of those stimuli and prior knowledge.

A key advantage of this framework is that it enables precise, testable predictions about consumer behavior. For example, while branding literature has established a strong theoretical foundation, emphasizing the importance of maintaining consistent brand associations and meanings for consumers (D. A. Aaker, 1991; Keller, 1993), it lacks precise models that generate testable predictions about consumer judgment and decision making. Similarly, mental accounting literature sometimes makes circular arguments, such as consumers tend to purchase two items together because they belong to the same mental account, but mental accounts are operationalized circularly as items that are purchased together (Cheema & Soman, 2006; R. Thaler, 1985). In contrast, our framework yields more precise and testable predictions about consumer judgment and decision making, such as how brand perceptions evolve over time and which products consumers are more likely to purchase jointly.

In Chapter 1, I develop a time-aware natural language processing approach to quantify how the meanings of perceptual brand attributes, such as *Visionary*, *Social*, and *Daring*, shift over time. I show that some apparent changes in brand–attribute associations are driven not by changes in consumer perceptions of the brand, but by shifts in the meanings of the attributes themselves. For example, a longitudinal brand-association survey suggests that Microsoft has become more associated with the attribute *Visionary* over time. However, this increase reflects a shift in what *Visionary* means—from *Rugged* to *Innovative*—rather than a change in how consumers perceive Microsoft. These findings highlight the importance for brand managers of accounting for changes in attribute meaning when interpreting the evolution of brand positioning.

To address this challenge, I introduce a computational approach based on time-aware word embeddings from computational linguistics. This method enables counterfactual predictions of how consumers in the past would have perceived brands today under the assumption that attribute meanings had remained stable—an assumption that cannot be evaluated using survey data or human judgment alone. For instance, if current consumers strongly associate Microsoft with *Visionary*, but the model predicts that consumers in 2001 (using 2001 meanings of *Visionary*) would have made a much weaker association, this provides evidence that the meaning of *Visionary* has shifted. Importantly, the model also generates testable implications that allow its underlying assumptions to be evaluated, addressing a key challenge in validating counterfactual analyses.

In Chapter 2, I build on the classic concept of mental accounting by formally modeling how mental accounts are constructed. This formalization introduces structure and constraints that enable direct comparison with competing explanations of consumer spending behavior. Existing theories of mental accounting provide limited guidance on how mental categories are formed, making them difficult to test rigorously and leaving room for flexible, post hoc interpretations. For example, if consumers spend a zoo trip rebate on a skiing trip, this behavior may be attributed to an "entertainment" account; if they instead donate the money to the Wildlife Conservation Society, it may be attributed to an "animals" account.

To provide a more precise account, I propose a similarity-based model of mental accounting in which consumers are more likely to spend money on products that are similar to the source of the money, regardless of whether the money is a windfall or hard-earned. This approach yields clear, falsifiable predictions and allows for direct empirical comparisons with alternative explanations, such as the widely held view that consumers are more willing to spend windfall income simply because of its unexpected nature.

In Chapter 3, I examine how consumer memory processes shape brand perceptions through systematic distortions. First, I investigate how new attributes acquired by a brand can spill over to competing brands through the formation of false memories, using the Deese–Roediger–McDermott (DRM) paradigm from cognitive psychology. This approach demonstrates how associative structures in memory can lead consumers to misattribute brand characteristics. In a related line of work, I study how consumers fail to adopt optimal retrieval strategies when the task involves recalling the absence, rather than the presence, of an event, highlighting an asymmetry that has received relatively little attention in prior consumer research.

Collectively, these projects advance a unified research program that seeks to both broaden the scope of consumer research and increase the precision of its theoretical constructs. By developing structured, quantitative representations of consumer cognition and integrating insights from computational linguistics and cognitive psychology, this dissertation provides a more rigorous account of how consumers construct, update, and apply mental representations in decision-making. In doing so, it offers new tools and perspectives for understanding how consumers navigate an increasingly complex and dynamic marketplace.

Chapter 1

Evolution of Brand Perception Attribute Meanings over Time

This chapter focuses on one key component of my broader framework: how the meanings of brand attributes evolve over time. By developing a quantitative approach to track shifts in attribute semantics, it aims to disentangle real changes in brand perception from changes in the underlying meaning of the attributes themselves. In doing so, the chapter provides a foundation for more precise measurement of brand positioning dynamics and offers new insights into how consumers interpret and update brand-related information.

1.1 INTRODUCTION

Branding is at heart a dynamic exercise. Given the rapid pace of changes in the landscape of consumer mindsets, competitive forces, and market trends, brands more than ever need to continually adapt and evolve their positioning to sustain relevance and competitiveness. Aaker, in his seminal book *Managing Brand Equity*, emphasized a need for brands to be agile because "a positioning strategy can become obsolete over time, the target market ages, or the association becomes less appealing (or even a source of ridicule) as tastes and fashions change."

Despite its importance, however, the ability of researchers and managers to track the evolution of brand perception remains limited owing to important conceptual and methodological challenges . First, a set of challenges exists involving practical difficulties in collecting and maintaining longitudinal data. Consumer surveys, where participants rate how strongly a brand is associated with attributes such as *Upper Class* or *Authentic*, are widely used in academic and industry research (Gupta & Hagtvedt, 2021; Mizik & Jacobson, 2008). However, while these surveys are common as one-time snapshots, they are less frequently employed longitudinally, given challenges in ensure consistency in sampling, coverage, and other methodological details which are necessary for valid cross-time comparisons.

In recent years, these challenges have become increasingly surmountable given the tremendous advances seen in big data and analytics, which promises to greatly expand the abilities of managers to track how consumer perceptions of products and brands evolve over time and in relation to their competitors. It is increasingly feasible, for example, for a brand to track the degree to which consumers consider a brand to be *Inclusive* or *Social* over time using online user reviews or social media discourse, either in conjunction with or in lieu of traditional consumer surveys (T. Y. Lee & Bradlow, 2011; Rust et al., 2021).

Second, important challenges exist on ensuring the validity of cross-time comparisons, due to the fact that the meaning of words used to describe brands and brand positions may themselves change over time. For example, inferring that a brand has become more *Inclusive* over time, whether based on social media or consumer survey results, depends implicitly on the assumption that meaning of the attribute has not changed significantly over time. Otherwise, any observed change could reflect shifts in the brand perception, the meaning of the attribute, or a combination of the two. Although comparatively less appreciated, consumer researchers have long warned about presence of validity threats in such comparisons. Indeed, as far back as 1975, the pioneering work of Gensch and Golob showed that meaning of attributes could change substantially over time, and urged researchers and practitioners to test for the consistency of attribute meanings before making cross-time comparisons.

Here, a useful analogy can be made with notions of "real" and "nominal" values in macroeconomics and the role of inflation in connecting the two (Diewert, 1988; O'Neill et al., 2017). Just as the nominal wage in a particular year is defined as the observed wage measured using prices from that year, we can define the "nominal brand perception" in a particular year as the observed brand perception, measured with the meaning of brands and attributes of that year. Just as to determine the real wage one needs to adjust for the changing value of money over time (i.e., inflation) by using the currency at a "reference time" to measure prices in different years, similarly we may also need to adjust for changes in attribute meaning to determine "real brand perceptions" using definition of attributes at a specific reference time.

If so, as with wages, direct comparisons using nominal data on brand perception can be highly misleading and may result in either over- or under-estimation of the actual changes, depending on how much the meanings of the measured constructs have changed. The interpretation, for example, if Budweiser is rated by consumers as being less *Social* over time, the inference that Budweiser has *really* become less *Social* in the minds of the consumer is only valid if the meaning of *Social* is largely constant over the period in question. As we will see later, the full picture is considerably more complicated given the changes in the meaning of *Social* during the past 20 years, and the same concern exists for many other, though not all, attributes.

In the current paper, we seek to shed light on these questions by developing a conceptual and analytical framework that aims to improve the ability of researchers and managers to track how brand perception and position evolve over time. In doing so, we contribute to the existing literature in three ways. First, we introduce the conceptual distinction between nominal and real changes in brand perception by drawing upon ideas developed over the past two centuries on the measurement of inflation and the value of money, which, to our knowledge, has not been considered previously. We further connect these ideas to seminal theories of branding and brand perception about how consumer brand perception can vary across time and place. In turn, we develop a conceptual framework to improve existing understanding of fundamental questions in brand perception, including (i) what exactly is being measured in commonly used brand perception surveys, (ii) how comparisons of data from these surveys can be complicated by the introduction of time, and (iii) what adjustments are necessary in order to make meaningful intertemporal comparisons.

Second, combining this conceptual framework with longitudinal brand perception data, large-scale time-aware text analysis, and controlled laboratory experiments, we develop an analytical approach to gain quantitative insights about the evolution of brand perception in ways that would not have been possible using each method alone. Specifically, we show how researchers and managers can compute real changes in brand perception to adjust observed nominal values by how the underlying attribute meaning has changed. We show evidence that, in many but not all cases, attribute meaning can change substantially, such that interpretations based on nominal changes in brand perception can be highly misleading. Finally, we provide examples of how these findings enrich our understanding of changes in brand positioning and can aid branding decisions.

1.2 RELATED LITERATURE

Below we provide a summary of two relevant streams of literature, focusing on topics within brand perception and diachronic linguistics. While each has been subject to intense study, their relevance for the other remains poorly understood, given the scarcity of prior contact between them. Therefore, we have chosen to aim for a high-level overview of the core concepts and methods within each field, as well as their relevance for the questions we address in the current paper, instead of attempting to provide a comprehensive review for each of them in isolation.

1.2.1 Brand Perception and Brand Meaning

Effective brand positioning requires managers to identify and communicate the unique value and distinct attributes of a brand that make it stand out in the marketplace. Such attributes may be physical, tangible, or concrete aspects of the brands, such as the functionality and technical specification of the products (D. A. Aaker, 1991, 2011; De Chernatony, 2010; Keller, 2003, 2016). They may also be intangible, i.e., non-physical characteristics that contribute to the brand's overall identity and perception in the minds of consumers. These intangible attributes have been increasingly recognized as a critical component of brand positioning and management, and constitute an active area of research, including notions of brand personality (J. L. Aaker, 1997a; Mathur et al., 2012), brand anthropomorphism (Aggarwal & McGill, 2012), and brand emotions (Coleman & Williams, 2013; Pham et al., 2013), among others (Brakus et al., 2009; Keller, 2003; Pham et al., 2013; Round & Roper, 2017).

Several methods have been developed to provide managers with timely understanding of how their brands are perceived by consumers. These range from open-ended and unstructured techniques such as free association and free response (Boivin, 1986; Krishnan, 1996; Olson & Muderrisoglu, 1979), to more structured techniques such as psychometric scales (Agres & Dubitsky, 1996), and pairwise similarity judgments (Hauser & Koppelman, 1979). In recent years, as consumer activities increasingly move online, there has been growing effort in leveraging big data methods and measures to infer brand perception at individual or aggregate levels (T. Y. Lee & Bradlow, 2011; L. Liu et al., 2020; Netzer et al., 2011; Wang et al., 2022). An important driver behind the enthusiasm for these methods is that they offer potential of more timely insights, in recognition that consumer perceptions of brands may quickly change over time (Culotta & Cutler, 2016).

Although a much smaller literature, there has also been recognition from the early days of marketing that interpretation of these changes across time can become problematic if meaning of the attributes themselves also change across time. Gensch and Golob (1975), for example, pointed out that whenever researchers compare the attribute scores of brands across different samples of participants, e.g., participants in different segments or time periods, they are implicitly assuming that the meaning of the focal attribute is consistent across those contexts. They further point out that there are many scenarios where this assumption might not hold, such as when knowledge and experience gained between two elicitations of brand perception change how consumers interpret the attributes, or mixing of different cultures over time in which attributes mean different things (J. L. Aaker et al., 2001). More recently, Rutz and Sonnier (2011) provide evidence that household purchases in the malt beverage category show evidence of changes in the latent brand attributes, which can be traced to advertising expenditures over time.

Perhaps the strongest evidence that such changes are a fundamental aspect of consumer behavior comes from the consumer culture theory (CCT) literature, where meaning-making is a core concept that emphasizes how consumers actively create and negotiate meanings through their consumption practices. In the CCT framework, meaning is not fixed or predetermined but emerges through ongoing processes of interpretation, negotiation, and appropriation within consumer cultures. Consumers, brands, and the cultural context actively co-create and reinterpret concepts and attributes such as *Luxury* (Kauppinen-Räisänen et al., 2019), *Wellness* (Thompson & Troester, 2002), or *Authentic* (Leigh et al., 2006), among others. This notion has received empirical support from other areas of research in marketing. For example, Mittal, Kumar and Tsiros (1999) showed attributes that are considered important in determining satisfaction can evolve over time, which in turn has consequences for consumers' behavioral intentions.

Managerially, the fact that attribute meaning can change over time carries important implications for brands and branding. As Rutz and Sonnier (2011) pointed out, stability of attribute meaning is a critical assumption underlying most existing models. The assumption that attribute meaning is fixed across time may therefore yield misleading recommendations, such as leading firms to fixate on old meanings that are no longer relevant, or overlooking the possibility of changes in attribute meaning, for example through advertising.

However, despite the recognition of its potential importance, studies exploring how attribute meaning changes and their managerial implications remain limited. Most practically, for example, when faced with declining ratings of *Social* in consumer surveys, it remains difficult for a brand manager to determine how much, if any, of this decline can be attributed to the changes in the meaning of *Social*. As a result, most studies assume, either explicitly or implicitly, that attribute meanings are stable during the period studied. While this may be a reasonable assumption if the time period involved is limited in duration, it becomes potentially problematic when longer durations or new product categories are involved.

1.2.2 Diachronic Linguistics

Although meaning change occupies a small sliver of the marketing literature, it is a core question within the field of diachronic linguistics which studies languages as they change over time (Deo, 2015; Weinreich et al., 1968). In particular, we draw on the area of semantic change, which studies the different ways in which the meanings of words change over time, the different processes underlying these changes, and most importantly for our purposes, how one can go about quantifying changes.

Specifically, in the semantic change literature, words are typically conceptualized as concrete forms that represent abstract concepts (Deo, 2015). While the forms (words) themselves are changeable, the underlying concepts they represent (their functions) can remain relatively stable, though they may also evolve over time (Traugott & Dasher, 2001). In turn, this allows semantic changes to be identified by investigating the context in which words are used. For instance, Geeraerts et al. (2012) argued that, in Middle English texts, the word "anger" gradually emerged as a dominant word to express anger of ordinary people by the turn of the 15th century based on the fact that it often occurred in the context where the experiencers were not in a position of authority, as opposed to the word "wrath," which were the dominant word in 1400s but primarily occurred in the context of kings and the God.

Linguists have devoted much theoretical work to categorize different types of meaning changes (e.g., Urban, 2015; Deo, 2015; Geeraerts et al., 2012). Some well-established examples include when the new meaning of a word is inclusive of its original meaning (*broadening*); when the meaning of a word becomes more specific over time (*narrowing*); when a word gains a new meaning that is similar to the original meaning (*metaphorical change*); when a word gains a new meaning that often co-occurs, either spatially or temporally, with the original meaning (*metonymic change*).

More directly relevant for our purposes is the question of how the forms and the functions of the words change over time. This is often studied from two perspectives. The first involves shifts in the meaning of words or expressions over time, known as *semasiology,* i.e., given a particular form of a word, how its function changes (Deo, 2015; Weinreich et al., 1968). For example, linguists have noted the ways in which the word "literally" has come to mean something opposite of its previous meaning, e.g., "I literally died laughing" (Deo, 2015). The second involves how the same meaning or concept comes to be expressed by different words, known as *onomasiology,* i.e., how people change the forms they use to denote the same function (Deo, 2015; Weinreich et al., 1968). For instance, the meaning of *joyful* used to be expressed by the word "merry," but is now much more commonly expressed by words like "happy." See Appendix B for a visual explanation.

Both perspectives are relevant for brand managers to understand the evolving meaning of their brands in terms of attributes. For example, brand managers who observed that the *Social* rating of Budweiser has dropped may take a semasiological approach and ask whether this change reflects the meaning change of the brand or the attribute, or both. And suppose that they find that it is *Social* that has changed in meaning, say, from face-to-face to online interaction, but they still want to measure Budweiser along the older meaning, they may take an onomasiological approach and ask which words can express the original meaning of *Social* and use those new words in future surveys.

Recently, building upon this conceptual foundation of diachronic linguistics and capitalizing on the increasing availability of large text corpora, computational methods for capturing the evolving meaning of words have been proposed. In the current work, we focus on the class of methods referred to as Diachronic or Dynamic Word Embedding (DWE; Deo, 2015; Kutuzov et al., 2018; Yao et al., 2018), chosen as they not only detect the existence of the changes, but also provide insight into what words mean in different time periods. Methods in this class represent different meanings of words at different time periods using high-dimensional vectors whose spatial relationship represents the semantic similarity between words (Bamler & Mandt, 2017). In this way, meaning change across time can be captured by directly comparing the movement of these vectors across different time periods.

Regardless of the specific Natural Language Processing (NLP) approach taken, it is critical that researchers take steps to evaluate the quality of inferred meaning changes in intertemporal comparisons in terms of how well they align with the ground truth (Deo, 2015; Kulkarni et al., 2015; Kutuzov et al., 2018). This can be challenging, as the gold standard of expert-annotated lists of semantically shifted words can be difficult to obtain given the cost and effort involved (Kutuzov et al., 2018). In contrast to using human experts, an alternative strategy attempts to validate the inference using a synthetic dataset where a semantic change is artificially introduced. For example, Kulkarni et al. (2015) sampled pairs of words from a real-world corpus, denoted the *donor* and *receptor* respectively, where the occurrences of the donor word in the corpus were replaced with the word receptor with a certain probability, which can be seen as the degree of semantic change. They then tested how closely the predicted significance of semantic change matched the degree of semantic change. Rosenfeld and Erk (2018) created synthetic words whose meaning varies between two senses over time following a predetermined path. They then tested how well the semantic change trajectory extracted by their method matches the predetermined path. While being a valuable tool for methodological studies to test the accuracy and sensitivity of semantic shift detection, synthetic methods are less useful for validating real-world semantic changes as the target of inference is pre-existing, rather than injected, changes.

Finally, there exists a broad class of validation strategies that uses predictions of semantic changes on various downstream tasks in the form of practical, application-focused activities that can demonstrate the utility of the upstream methods. For instance, one evaluation task, called word epoch disambiguation, tests the ability of a method to detect the time period in which a word within a particular context occurred (Mihalcea & Nastase, 2012; Popescu & Strapparava, 2015). Another study used detected semantic shifts to predict real-world events such as armed conflicts (Kutuzov et al., 2017) by asking the extent to which change in the contexts of country names in news corpus can be used to predict whether the country is at war or not. A benefit of this class of method is that it does not rely on having access to a gold standard and can be tailored to specific applications in question such as those in marketing.

1.3 METHODS

1.3.1 Longitudinal Survey Data of Consumer Brand Perception

Longitudinal consumer brand tracking data between 2001 to 2020 was obtained from The BAV Group (BrandAsset® Valuator, www.bavgroup.com), containing a rotating panel representative of the U.S. population which were asked to rate each brand on 46 brand perception attributes (see Appendix A for complete list; Gupta and Hagtvedt 2021; Mizik and Jacobson 2008), such that brand attribute scores represent the percentage of respondents who endorsed the brand (selected *Yes* in a binary question). We removed all brands related to cities, countries, hemp/cannabis, individual people, movies, nonprofit organizations, sports franchises, or TV shows. Our final dataset contained brand ratings of 151 brands on 46 attributes covering years 2001 to 2020.

1.3.2 Statistical Procedure of Longitudinal Survey Data

To assess stability of individual attributes, e.g., *Social*, we computed the pairwise Pearson correlation of the ratings for all attributes between 2001 and 2020. If attribute meaning did not change across time, there should not be systematic changes in correlation between individual attributes over this time period (Kline, 2015). To quantify the change in global attribute structure, we took the correlation matrices of the 46 attributes, and vectorized each of the upper triangular parts (as the matrices are symmetric) of the 46 by 46 matrices. We then picked a fixed reference year and computed the Pearson correlation between correlation matrices of each year and that of the reference year. The resulting time series illustrates the cumulative changes of the global attribute structure from the reference time.

1.3.3 Text Corpus

Corpus data were obtained by all articles published in New York Times between January 1996 to December 2023. We use yearly time slices, dividing the corpus into $T = 28$ partitions. After removing rare words (fewer than 100 occurrences in all articles across all years), our vocabulary consists of $V = 159{,}945$ unique words. We then preprocessed the corpus by transforming all letters to lowercase and removing punctuations. To be able to analyze brand names consisting of multiple words, we used the AutoPhrase algorithm (Shang et al., 2017) to detect phrases, defined as a multi-word expression that is used as an independent unit, in a corpus. The resulting words of the detected phrases are then connected using underscore and treated as a single word in the model training process.

1.3.4 Dynamic Word Embedding Model Training

We used the DWE approach, a diachronic method developed in Yao et al. (2018) to quantify changes in attribute meaning. Specifically, we trained a 50-dimensional DWE model on all New York Times articles from 1996 to 2023. Mathematically, DWE optimizes the following objective function:

$$\min_{V(t),t\in\{1,\ldots,T\}} \frac{1}{2}\sum_{t=1}^{T}\|Y(t) - V(t)V(t)^T\|_F^2 + \frac{\lambda}{2}\sum_{t=1}^{T}\|V(t)\|_F^2 + \frac{\tau}{2}\sum_{t=2}^{T}\|V(t-1) - V(t)\|_F^2$$

where $V(t)$ is the word embedding vectors of all the words in time t , T the total number of time points in the corpus, $Y(t)$ the positive pointwise mutual information (PPMI) matrix of words in time t, which is a function of the co-occurrence of every pair of words in time t . $Y(t)$ is computed with a window size $L = 10$.

The first term in the objective function ensures that the inner products of word embeddings align with their PPMI, reflecting semantic relationships at the same time point. The second term is a standard regularization term in NLP. The third term states that word meanings do not change abruptly within a single time point, allowing for alignment of word vectors across adjacent years while maintaining within-year similarity.

Two key parameters govern the training of the vectors. The semantic continuity parameter τ, set to 50, controls how fast we allow the embeddings to change. A value $\tau = 0$ enforces no alignment, such that vectors in each year t are trained only using text from year t, and as $\tau \rightarrow \infty$ converges to a static embedding with $V(1) = V(2) = \cdots = V(T)$, such that no meaning change is allowed. The regularization parameter λ, set to 10, reduces overfitting by penalizing larger weights and encourages the model to prioritize simpler hypotheses.

Word vectors were initialized using New York Times in 1996. The choice of 1996 was chosen to be sufficiently prior to the first year in our BAV dataset, 2001, so as to avoid any potential information from "leaking" if we initialized using years between 2001-2020. That is, the model should not have access to information from the future in a way that can artificially benefit prediction.

Note that as we only have the BAV data between 2001 and 2020, we only use the DWE vectors in that period in all subsequent analyses.

1.3.5 Brand Meaning Word Vector

Having trained the semantic projection model, we extract the brand meaning vector for a brand using the model vector associated with the brand name for each year. For instance, to obtain the brand meaning vector for Nike in 2008, we obtained the model vector corresponding to the token "nike" in 2008.

1.3.6 Attribute Meaning Vector

Attribute meaning vectors for each year were estimated using ridge regression (Hastie et al., 2017). Specifically, each attribute j in year t is a 50-dimensional vector that solves the regularized least squares problem

$$\min_{\beta_j^t} \left\{ \left\| y_j^t - X_j^t \beta_j^t \right\|^2 + \lambda_j^t \left\| \beta_j^t \right\|^2 \right\},$$

where $\|\cdot\|$ is the Euclidean norm, y_j^t is the $N \times 1$ vector of the BAV ratings for all brands on attribute j in year t, X_j^t the $N \times 50$ matrix consisting of all 50-dimensional word vectors pertaining to brand j at time t, and β_j^t the 50×1 vector capturing semantic meaning of attribute j in time t. Thus, the attribute meaning vector is a 50-dimensional semantic vector that best explains contemporaneous BAV ratings of the brands.

1.3.7 Nominal and Real Brand Perceptions

As introduced earlier, we apply the semantic projection approach to infer nominal and real brand perceptions using the brand and attribute meaning vectors. The same approach applies to both cases, except that, for real perceptions, the time index of the attribute meaning vector is in general different from that of the brand meaning vector. To compute brand perception in any year t_1 based on attribute meaning of some reference time t_2, we simply need to project the brand vectors in t_1 onto attribute meaning vectors from t_2. In this way, real perception of a brand i at time t_1 using the meaning of attribute j at reference time t_2 is expressed as:

$$Brand_i\left(attribute_j, t = t_1, t_{ref} = t_2\right) = < brand_i{}^{t_1}, attribute_j{}^{t_2} >,$$

where $brand_i^{t_1}$ is the brand meaning vector of brand i at time t_1 and $attribute_j^{t_2}$ at time t_2. Conversely, nominal perception of the brand corresponds to cases where $t_1 = t_2$:

$$Brand_i(attribute_j, t = t_1, t_{ref} = t_1) =< brand_i^{t_1}, attribute_j^{t_1} >$$

1.3.8 Performance of Human vs. DWE on Predicting Past Consumer Brand Perception

A total of 290 participants were recruited from Prolific, of which 242 completed the entire survey and were included in the subsequent analyses (122 females, $M_{age} = 42$). To construct the test set for human retrospective judgments, we selected 600 binary-choice questions out of the total possible 12,524,272 from 2002 to 2019 which were used in Study 2A. Of the 600, 300 are unstable questions, defined as the questions where the correct answer changed between the year being tested and 2020, and the other 300 are stable questions, where the correct answer did not change. Each participant answered 30 randomly selected questions of each type, for a total of 60 questions.

1.3.9 Contemporaneous Model

The Contemporaneous semantic projection model computes perception ratings by calculating $< brand_i^t, attribute_i^t >$, where $brand_i^t$ denotes the DWE vector of $brand_i^t$ in time t, and $attribute_i^t$ denotes the attribute vector of $attribute_i^t$ in same time t as the brand vector. In other words, in this model the year of attribute and brand vectors are aligned as mentioned in the previous section. As this model corresponds to nominal ratings, no reference years were involved such that both attribute meaning and brand perception are allowed to vary.

1.3.10 Attribute-Constrained Model

The Attribute-Constrained semantic projection model computes perception ratings by calculating $< brand_i^t, attribute_i^{t_{ref}} >$, where t_{ref} is the reference time. In other words, this model uses the outdated attribute meaning but correctly updated brand meaning over time. Three reference years were used—2001, 2010, 2020.

1.3.11 Brand-Constrained Model

The Brand-Constrained semantic projection model computes perception ratings by calculating $< brand_i^{t_{ref}}, attribute_i^t >$. In other words, this model uses the outdated brand meaning but correctly updated attribute meaning over time. Three reference years were used—2001, 2010, 2020.

1.3.12 DWE-Aided Brand Retrospection

A total of N=201 (90 females, M_{age}=41) participants were recruited from Prolific. All participants completed the study. Stimuli were generated by selecting 40 binary choice questions of the year 2001. All questions are about pairs of brands whose relative ranking has flipped during the past twenty years, i.e., unstable pairs. Each participant was randomly assigned into either the Past Meaning or Present Meaning condition.

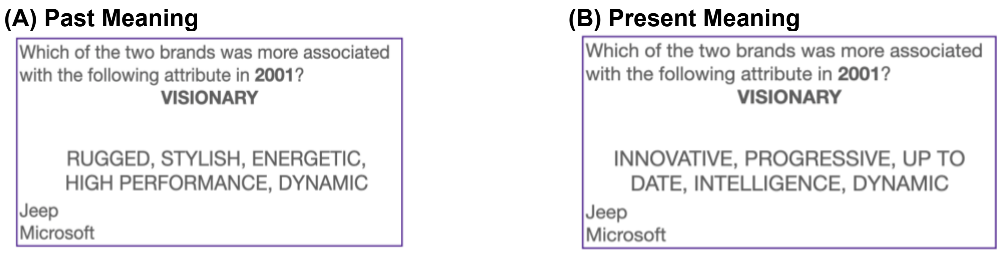

Figure 1.3.12.1 *Example question prompts in (A) Past Meaning and (B) Present Meaning conditions.*

In the Past Meaning condition, we generated the focal attribute in the past by translating the focal attribute in the language of the present (Figure 1.3.12.1A). Specifically, for all the attribute vectors 2020, we calculated their cosine similarity with the focal attribute vector in 2001, and then selecting the top five attributes in terms of the cosine similarity. In the Present Meaning condition, we generated the synonyms of the focal attribute by selecting the top five attributes that are most highly correlated with the focal attribute using their BAV ratings in 2020 (Figure 1.3.12.1B).

1.4 STUDY 1 OVERVIEW: SEMANTIC SHIFT IN BRAND ATTRIBUTE MEANING

While it is well-established that language is constantly changing, and that linguistic shifts can occur rapidly within the span of a decade or even less, the frequency and magnitude with which such changes happen within a branding context remains poorly understood. Even more importantly, it is unclear whether existing ways of measuring brand perception have sufficient sensitivity and temporal resolution to be able to capture such changes. Therefore, we sought in Study 1 to provide initial qualitative and quantitative evidence that attribute meaning changes can occur over the types of timespans that are relevant for managers and consumer researchers. We do so by using a longitudinal brand perception dataset from the BAV Group, a brand tracking firm that maintains one of the largest continuous brand surveys.

This longitudinal panel allows us to approach measuring changes in attribute meaning in a similar way as psychometric evaluation of the construct stability over time using test-retest correlation (Gensch and Golob 1975; Kline 2015). Specifically, we examine the extent to which the correlation structure between attributes changes over time. Under the null hypothesis that no changes have taken place, the correlation between attributes should remain the same. For example, if the meaning of the attribute Social is unrelated to that of Innovative in t_1, it should also be minimally correlated with Innovative in t_2. In contrast, time series of correlation changes provide hints of both the presence of change as well as the temporal trajectory of the pace of change. Changes occurring in a short period of time, for example, should result in a sudden but enduring shift in the correlation structure, whereas those which unfold over a longer period should be reflected in small but systematic drifts.

1.5 STUDY 1 RESULTS

1.5.1 Semantic Shifts in Brand Attribute Meaning

We start with a bird's-eye view of how the relationship between all attributes has changed over time. We do so by visualizing the correlation between pairs of attributes in 2001 and 2020 using multidimensional scaling (MDS, Figure 1.5.1.1). Thus, the distance between a pair of attributes in MDS space provide an indication of the extent to which the attributes share a similar meaning.

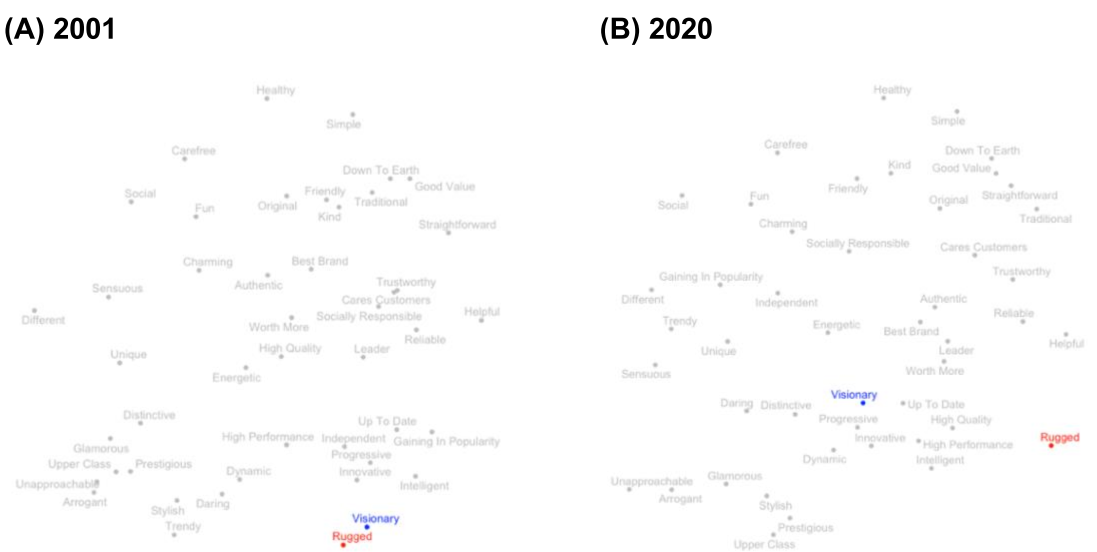

Figure 1.5.1.1 *Changes in attribute correlations between (A) 2001 and (B) 2020 visualized by MDS. Attributes that are closer on the plots are more similar in meanings.*

Comparing these two time points, some pairs of attributes can be seen to remain consistently close together, such as *Original* and *Kind*, whereas others are consistently far apart, such as *Prestigious* and *Down to Earth*. However, not all attribute pair relationships are so stable. *Visionary* and *Rugged* change from being highly correlated in 2001 to highly different in 2020. In contrast, *Visionary* and *Progressive* exhibited the opposite trend, becoming more similar over this time period.

Attribute	Top brands in 2001	Top brands in 2020
Visionary	Jeep	Disney
	Ford	Apple
	Dodge	Tesla
	Nike	Microsoft
	Reebok	Netflix

Rugged	Jeep Dodge Ford Nike Reebok	Jeep Ford Subaru Dodge GoPro

Table 1.5.1.1 *Brands most associated with Visionary and Rugged in 2001 vs. 2020.*

In addition to changes in the correlation structure of attributes, we can also inspect how potential changes in attribute meaning impacts the type of brands that are rated highly on those attributes (Table 1.5.1.1). For example, while brands most associated with *Visionary* in 2001 were those related to automobiles and sports, in 2020 they became dominated by technology brands. In contrast, for the attribute *Rugged*, which did not change substantially, similar types of brands appeared in both 2001 and 2020.

Importantly, these changes varied greatly in terms of their temporal trajectory. To obtain a more granular view of the evolution of global attribute structure over time, we computed how far the overall attribute correlation structure has shifted over time. Figure 1.5.1.2A shows the movement of the correlation matrices using three different reference times, 2001, 2010, and 2020. Certain years, such as 2006, exhibited greater mean change. This is particularly notable when looking at the temporal trajectory of attribute pairs (Figure 1.5.1.2B). Changes between *Rugged* and *Visionary*, for example, was largely driven by a one-time drop between 2006 and 2007, whereas that of *Innovate* and *Social* occurred over the time period of 2006 to 2010 before plateauing. On the other hand, the correlation between *Glamorous* and *Prestigious* remained fairly constant during this time.

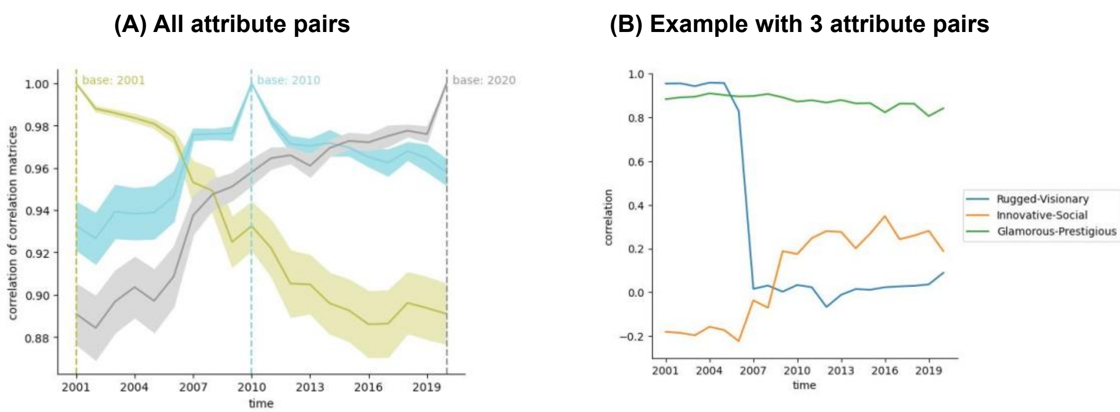

Figure 1.5.1.2 *(A) Correlation between the 2001, 2010, 2020 attribute correlations and all other correlations. Error band indicates standard errors. (B) Change in correlation across time in three pairs of attributes: Rugged-Visionary, Innovative-Social, and Glamorous-Prestigious.*

	Most changed attributes	**2nd most changed attributes**	**3rd most changed attributes**
Any 1 year	Visionary (2006-2007)	Socially Responsible (2013-2014)	Independent (2016-2017)
Any 10 years	Visionary (2006-2016)	Social (2006-2016)	Socially Responsible (2004-2014)

Table 1.5.1.2 *The attributes experiencing the greatest change at different time scales.*

Given the variation in speed of change, the timescale of the comparison becomes critical. Looking at the time intervals of one year, for example, *Independent* and *Rugged* are among the top three most changed attributes (Table 1.5.1.2). However, over longer time intervals such as ten years, the two are replaced by *Social* and *Socially Responsible* in the top three. This suggests attributes can have different rates of change, with some having large short-term changes but not long-term changes (*Independent*), some having large long-term changes but not short-term changes (*Social*), while others having both long- and short-term changes (*Visionary* and *Socially Responsible*).

1.6 STUDY 2 OVERVIEW

It might be tempting to ask if it is sufficient to recover real changes in attribute meaning from changes in the correlation structure between attribute ratings. The answer, unfortunately, is no, as the latter only provides information about changes in the relationship between the attributes and cannot distinguish different possible changes that can take place. Take for example the evolution of the word viral in the past few decades. Simply knowing that viral is becoming more similar to social media is not enough to conclude that viral has changed in meaning, without accounting for the extent to which social media have changed.

To capture meaning change that is necessary for inferring real changes in brand perception, we used a time-aware diachronic NLP technique, DWE (Yao et al. 2018) that considers the relationship between all words in the corpus. On a high level, DWE and related time-aware NLP methods build upon well-established static word embedding methods where words are assigned to a high-dimensional vector, such that words that are more related to each other are assigned to vectors that are closer in the vector space (e.g., Mikolov et al. 2013). DWE generalizes word embedding models by adding a time component, such that the same word across different time periods can be represented as a set of time-indexed vectors placed in a common semantic space, which in turn allows vectors to be compared across time.

As discussed above, the ambiguity in meaning adjustment can be overcome by looking at how surrounding context words, operationalized as nearby word vectors in the semantic space, have changed (Deo 2015; Weinreich et al. 1968). Specifically, by examining how different words move in relation to each other in the semantic space, we can understand meaning changes from the two perspectives introduced earlier. In Study 2A we take a semasiological perspective in how brand or attributes come to be associated with different meanings across time. For example, Apple becoming more associated with Visionary over time. In Study 3B, we adopt an onomasiological perspective where we "translate" the meaning of attributes from the past into the language of the present to help participants predict how consumers perceived brands in the past, e.g., the 2020 meaning of Visionary has become a poor indicator of its 2008 meaning, which can instead be better approximated by the 2020 meaning of High Performance.

1.6.1 Study 2A Overview

Study 2A formalizes the above intuition using DWE. Rather than requiring an a priori specific set of context words as was necessary in the past, approaches such as DWE allow researchers to use a data-driven approach where *all* words in the corpus can be used as potential context words depending on their stability. In turn, this provides us with an approach to decompose all brand perception, real or nominal, in terms of their component parts: (i) the totality of the brand meaning at any particular time t, and (ii) the meaning of the attribute at a reference time t_{ref}.

More formally, the perception of brand i at time t on some attribute j at reference time t_{ref} can be captured by a similarity metric, e.g., inner product, between a semantic vector $brand_i^t$ capturing the meaning of the brand at time t and a semantic vector $attribute_j^{t_{ref}}$ capturing the meaning of the attribute in reference time t_{ref}:

$$Brand_i\left(attribute_j, t, t_{ref}\right) =< brand_i{}^t, attribute_j{}^{t_{ref}} >.$$

Following Grand et al. (2022), we will refer to this as the *semantic projection* of brand i on attribute j, but generalizing to include the time indices necessary in our case. An intuitive case of applying this approach is nominal perception— how a brand is perceived on some attribute can be captured by measuring the similarity between the brand vector and the contemporaneous attribute vector.

For example, the nominal *Visionary* rating of Facebook in 2020 can be captured by projecting the 2020 brand meaning vector of Facebook to the 2020 attribute meaning vector of *Visionary* (Figure 1.6.1.1A), i.e., $< Facebook^{2020}, visionary^{2020} >$. We will refer to this as the *Contemporaneous semantic projection model*, as the attribute meaning vector and the brand meaning vector are from the same time.

Importantly, this approach to quantify meaning can be applied to any pair of semantic vectors without requiring them to be from the same time, which offers the benefit of dissociating the time index of brand meaning and attribute meaning. For example, we can express how Facebook in 2020 is perceived on the *Visionary* attribute in 2008, i.e., $Facebook\left(visionary,\ t = 2020,\ t_{ref} = 2008\right)$, by projecting the

2020 brand meaning vector Facebook to the 2008 attribute meaning vector of (B), i.e., $< Facebook^{2020}, visionary^{2008} >$ (Figure 1.6.1.1B). We refer to this as the *Attribute-Constrained semantic projection model*, as the attribute meaning vector is constrained to the reference time.

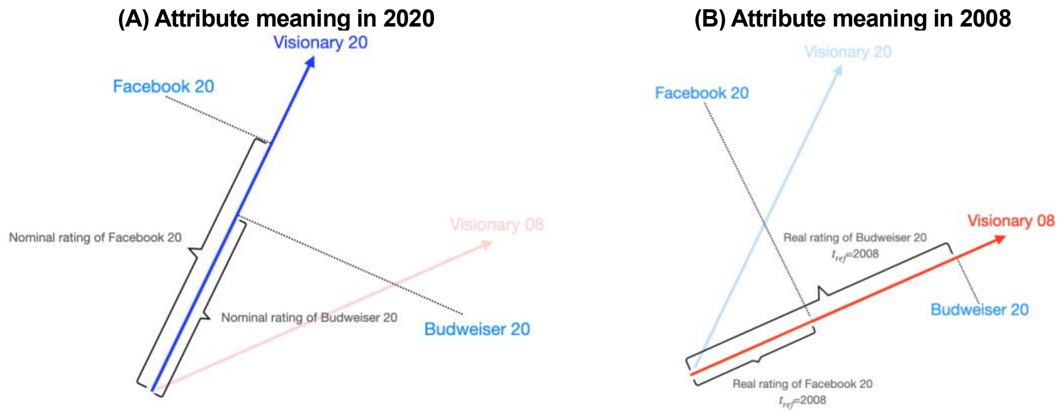

Figure 1.6.1.1 *How Facebook was perceived in 2020 based on (A) the meaning of the attribute Visionary in 2020 and (B) 2008, illustrated using semantic projection of brand vectors on attribute meaning vectors. Brackets indicate strength of association.*

1.6.2 Study 2A Results

DWE Can Capture Evolution of Consumer Perception. Using the approach outlined above, we show how a semantic projection model trained on the New York Times corpus can be used to recover real brand perception by decomposing nominal brand perception into component parts corresponding to (i) a brand vector containing brand meaning at a particular year, and (ii) a separate attribute vector containing meaning of the attribute.

Specifically, we obtained the attribute meaning vectors by finding the semantic vectors that best explain BAV ratings of the brands for each year to visualize how attribute meaning evolved over 20 years using t-SNE projection (Figure 1.6.2.1; Maaten and Hinton (2008)). For example, in early 2000s, *Visionary* was close to *Rugged*, but then quickly moved toward *Innovative*, *Progressive*, *Intelligent*. In contrast, *Upper Class*, *Prestigious*, and *Glamorous* move in unison with one another over time.

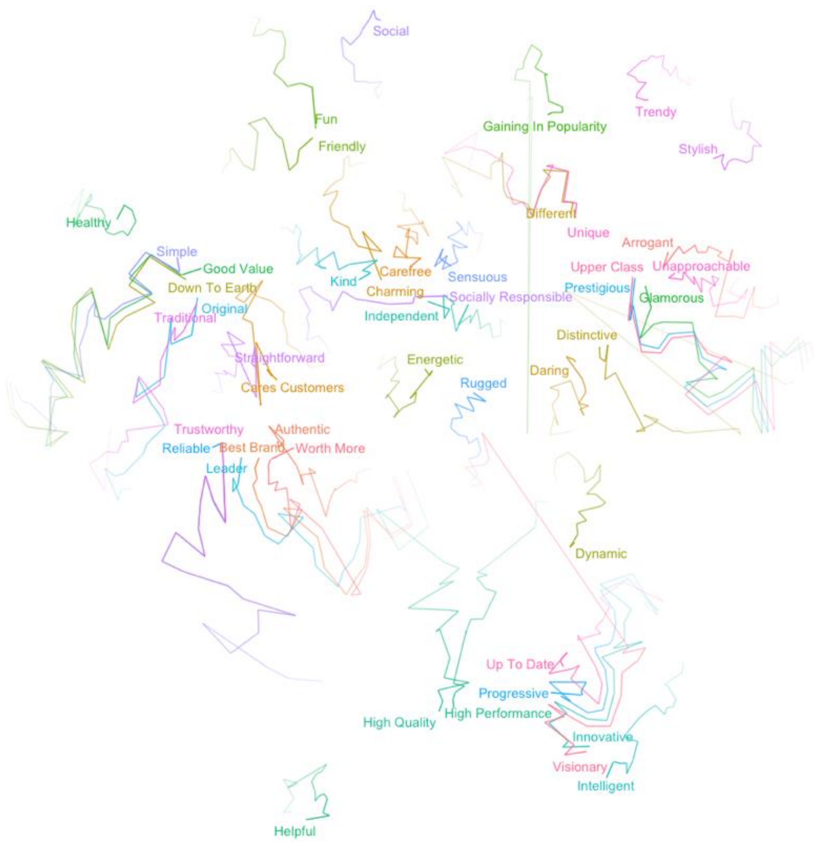

Figure 1.6.2.1 *Visualizing data using t-SNE projection of the attribute vectors. Each line represents the temporal trajectory of an attribute in the semantic space from 2001 to 2020. The lines are shaded such that the trajectories in 2001 are the lightest and those in 2020 are the darkest. The attribute labels are positioned at the attribute vectors in 2020.*

To evaluate the quality of the decomposition, we assessed the extent to which the trained model can provide out-of-sample predictions of BAV ratings that are both accurate and stable in the sense of not exhibiting significant temporal trends over time. This is important as the quality of the trained model requires that (i) the information contained in the underlying corpus reflects the evolution of brand meaning and attribute meaning over time, and (ii) this information can be captured in our semantic projection model and our operationalization of brand and attribute meaning vectors. If accuracy fluctuates greatly or contains a systematic trend, it may reflect shortcoming of either the corpus, for example if the New York Times becomes systematically more or less representative of US consumers during certain periods, or assumptions underpinning our semantic projection model.

Specifically, we tested the quality of our semantic projection model by asking the extent to which the trained model was able to make accurate out-of-sample predictions of brands in terms of their scores on different attributes. We tested, for each year, all possible pairs of brands within that year and asked the extent to which the trained model was able to select which of the two brands scored higher on each attribute. For example, in 2015, whether the model was able to correctly identify Facebook as being rated higher in *Social* compared to Disney. This metric quantifies the ability of each model to correctly identify which one of a pair of brands scored higher on a given attribute at a specific time. To ensure that all predictions are out-of-sample, we used the attribute and brand meaning vectors in time $t-1$ to predict the BAV ratings in time t. At the aggregate level, we found that, across all pairs of brands and all attributes and years available in the BAV dataset, the Contemporaneous semantic projection model was able to perform significantly greater than chance, averaging 82% across all years ($\chi^2(1)$=583,103, $p<10^{-8}$). Moreover, no significant time trend was observed when investigating across all years (β=1, t=1.759, $p>0.05$; Figure 1.6.2.2).

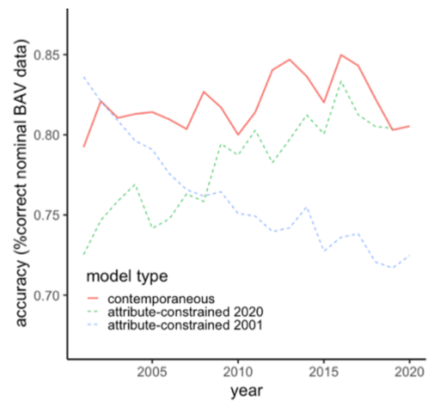

Figure 1.6.2.2 *Performance of the Contemporaneous semantic projection model and the Attribute-Constrained models with different reference times over time. Here we show the accuracy of the models on all possible binary-choice questions.*

Importantly, in the previous analysis, we allowed attribute meaning to change in accordance with the year of BAV data. That is, we asked whether our semantic projection model is able to recapitulate nominal BAV brand ratings. In contrast, if we were to constrain attribute meaning to some particular year, model performance steadily declining in years farther away from the reference time. Specifically, we evaluated model performance constraining attribute meaning to a reference time of 2001 and 2020, respectively, and tested on BAV ratings from 2002 to 2019. Consistent with idea that attribute meaning change is an important determinant in changes in nominal brand ratings, we found that constraining attribute meaning to these years resulted in significant time trends and declining performance as time of interest is farther away from the reference time (t_{ref}=2020: $\beta=5.26\times10^{-3}$, t=8.37, $p<10^{-6}$; t_{ref}=2001: $\beta=-5.44\times10^{-3}$, t=−12.16, $p<10^{-6}$; Figure 1.6.2.2).

Having shown that our approach is capable of decomposing nominal ratings into semantic vectors corresponding to brand and attribute meaning, it now becomes possible to understand the evolution of brand meanings and adjust for attribute meaning changes of any brand within our dataset (Figure

1.6.2.3). To facilitate the interpretation of brand meanings, we visualize the brand trajectories using Principal Component Analysis (PCA, Figure 1.6.2.3).

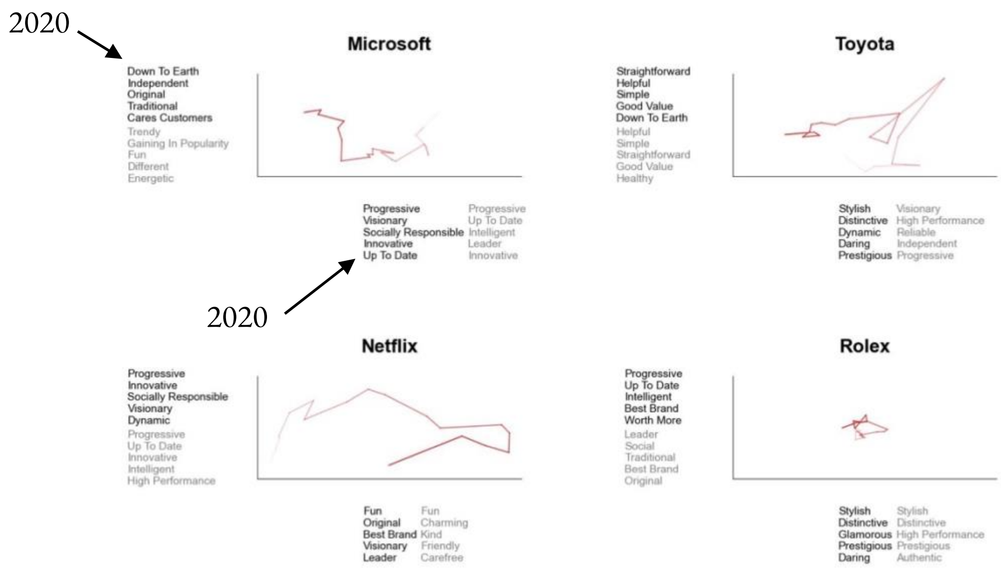

Figure 1.6.2.3 *Temporal trajectories of selected brands along the first (x-axis) and second (y-axis) principal components (PCs) of their DWE vectors since their first appearance in the corpus to 2020. Labels along axes represent meaning of corresponding PCs, with meanings of attributes in their 2020 (dark gray) and 2001 (light gray) senses. See Appendix C for all brands.*

The semantic trajectory of brands serves two functions. First, it illustrates the movement of brand meanings in the semantic space. For instance, Microsoft has become less *Progressive* and *Visionary*, but more *Down To Earth* and *Independent* over time. In contrast, Rolex has roughly the same set of associations over time. It is worth pointing out that the attribute labels merely describe the difference in brand association over time, as opposed to a brand's absolute positioning. For instance, the presence of the attribute *Prestigious* and *High Performance* on the first principal component of Toyota's DWE vectors indicates that Toyota has moved a lot, in this case decreased, along those attributes, but this does not necessarily mean that Toyota is highly associated with them.

Second, more subtly, it showcases the importance of specifying the reference time when interpreting brand meaning changes. Take Toyota as an example. When interpreted using attributes in their 2020 senses, Toyota has become less *Stylish* and *Distinctive* over time. However, when interpreted using attributes in their 2001 senses, Toyota has become less *Visionary* and *High Performance*. Importantly, in both cases Toyota has the same real meaning changes. The fact that Toyota is best described by a difference set of attributes appears to be largely due to change in attribute meanings.

Next, we proceed to convert nominal ratings to real ratings, which can help explain otherwise anomalous changes in brand perception. For example, through 2001 and 2005, Dodge and Toyota were viewed as highly *Visionary* in nominal terms, more so than Microsoft and Dodge (Figure 1.6.2.4). This, however, changed dramatically in 2007, when Microsoft experienced a sudden increase in *Visionary* and Dodge, a corresponding drop.

However, the interpretation of these movements changes considerably after adjusting for the meaning change in *Visionary*. First, using 2020 as the reference time, the real brand perception of Microsoft on *Visionary*, rather than improving over time, has instead been declining slowly over time, albeit still much higher than Dodge. Using the reference time of 2001, the real perception of Dodge on *Visionary* can be observed to be declining steadily over time, from well above Microsoft to below it by the mid-2010s. In contrast, Microsoft has largely remained largely constant during this period. Therefore, the sudden changes in nominal ratings for these brands were due in large part to changes in the meaning of *Visionary* occurring around 2006-2008, rather than consumer perception of the brands themselves.

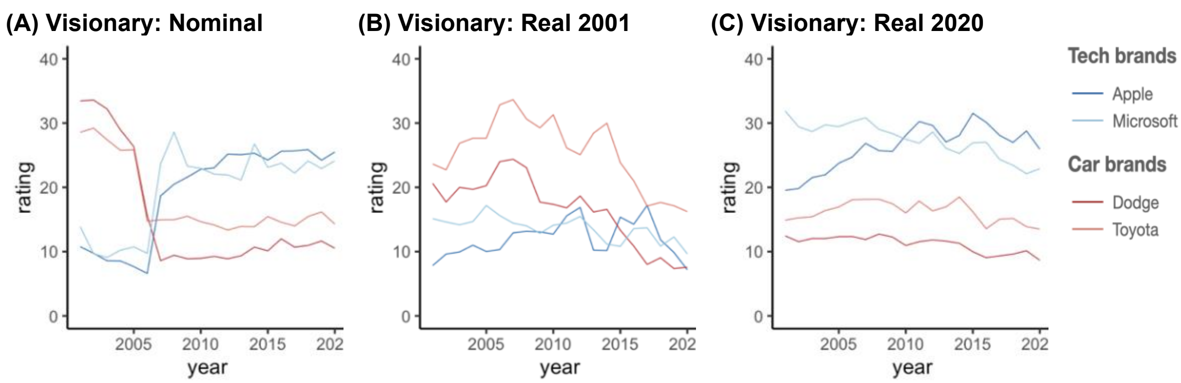

Figure 1.6.2.4 *The (A) nominal rating, (B) inferred real rating using 2001 and (C) 2020 as the reference time of selected brands in the car and technology sector.*

1.6.3 Study 2B Overview

To further illustrate the value of our approach, we compared performance of the Contemporaneous semantic projection model against performance of human participants in retrospective judgments of brand perception. We constructed a set of binary-choice questions drawn from those in Study 2A. In each retrospective judgement question, participants were asked to choose which of two brands was rated higher on some attribute in a particular year, e.g., "Which brand was considered by a panel of representative consumers to be more *Energetic* in 2010, *Disney* or *Yoplait*?"

We further divided the questions into two types—*stable* and *unstable*. Stable questions consisted of those where the correct answer did not change between the year being tested and 2020, which was the most recent year of BAV data that we had access to. For example, according to the BAV data, in both 2010 and 2020, *Disney* was considered more *Energetic* than *Yoplait*, and indeed throughout the entire period of our data. Conversely, unstable questions are defined as those where correct answer changed across years. For example, whereas *Dell* was considered more *Distinctive* than *Facebook* in 2006, this changed after 2011 (See Appendix D).

This is important because whereas we can precisely control which years of information the DWE model had access to in Study 2A, it is not possible to do so with human participants. Thus, a DWE model trained on a 2012 corpus will not be able to rely on information from 2018. In contrast, human participants can use their current knowledge to "fill in" any gaps in their knowledge about the past. In particular, if the brand relationships did not change between the past and present, participants could be highly accurate by simply using their current knowledge. The inclusion of unstable questions rules out this possibility, such that accuracy requires participants to flexibly account for changes in meaning of both attribute and brand.

1.6.4 Study 2B Results

To compare performance, we calculated, for each question, the paired difference in the accuracy between Contemporaneous semantic projection model and that of the human participants, and binning across three time periods: 2002-2008, 2009-2013, and 2014-2019 (Figure 1.6.4.1). We found that in the most recent time period (2014-2019), the performance of human participants was slightly worse than that of the model on the stable questions (D=3.87%, t=3.165, p<0.002), but not significantly different on unstable questions (D=-5.10%, t=-1.01, p=0.314).

Performance on the two question types diverges even more when extended to the previous two time periods. For stable questions, human participants performed worse than the model in the first two time periods (2002-2008: D=10.06%, t=8.449, $p<10^{-6}$; 2009-2013: D=13.67%, t=11.05, $p<10^{-6}$). The difference in performance becomes more pronounced for unstable questions (2002-2008: D=28.13%, t=30.91, $p<10^{-6}$; 2009-2013: D=13.37%, t=8.996, $p<10^{-6}$), and a significant interaction exists between the time and the question type (stable vs. unstable, β=-0.0079, t = -2.607, p=0.009).

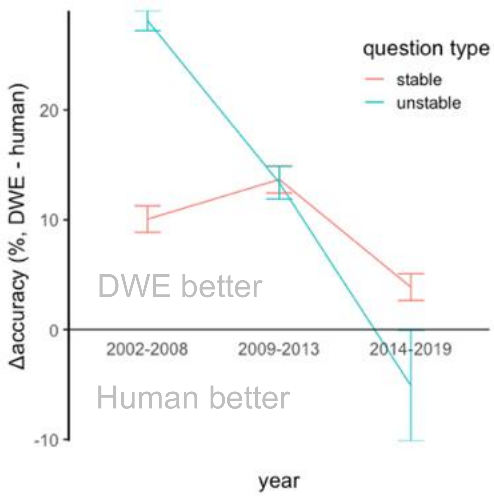

Figure 1.6.4.1 *DWE outperforms human participants in retrospective judgments of brand attribute associations across time. The error bars indicate standard errors.*

Taken together, these results provide a demonstration of the practical value of our approach given that retrospective judgement is perhaps the only alternative available to managers in reconstructing brand meaning across time. If human participants were able to give accurate retrospective judgments of BAV brand perception across time, it would suggest that marketers could instead use human raters to adjust for attribute meaning changes over time. However, consistent with past findings in psycholinguistics and diachronic linguistics, human retrospective performance was both worse overall and systematically declined over times, suggesting that they were unable to correct for changes in brand meaning, attribute meaning, or both.

1.7 STUDY 3 OVERVIEW

Although our approach provides a way of decomposing nominal brand perception into their component parts, it relies on an important assumption regarding the validity of the attribute meaning change inferred from the data. Study 3 therefore seeks to provide evidence supporting the validity of these inferences. Obtaining direct validation for meaning change, however, is a challenging task, as it relies on a counterfactual dataset that does not exist. Consider the case when one wants to evaluate Coca-Cola in 2008 on some attribute in its 2015 meaning. Direct validation requires us to take a consumer from 2008 and transport her to 2015 to learn about the meaning of the attribute in 2015, or alternatively transport a consumer from 2015 to 2008 to learn about the perception of Coca-Cola in 2008. Without the ability to time travel, we turn to another strategy: by designing downstream tasks that rely on the validity of our approach (Kutuzov et al., 2017; Mihalcea & Nastase, 2012; Popescu & Strapparava, 2015).

Specifically, we will build on the observation in Study 2 that human retrospective performance declines systematically over time, suggesting that human retrospection contains systematic errors due to lack of sufficient update of either attribute or brand meaning changes. If so, a test for the validity of attribute meaning generated by the semantic projection model can be formulated based on the extent to which retrospective performance improved when participants are provided with correctly-dated attribute meaning that are generated by the semantic projection model, i.e., the descriptors are chosen based on the *onomasiological* change. That is, if retrospection errors partially reflect the use of incorrectly dated attribute definitions, correcting these errors should improve retrospection performance. On the other hand, if recovered changes are spurious, it should not result in improved retrospection performance. Accordingly, in Study 3A, we use our model to show human retrospection errors contain systematic errors due to the use of outdated definitions of attributes. In Study 3B, we show that accuracy in retrospective judgments improve when participants are provided with the correctly-dated definition of attribute as generated by our model.

1.7.1 Study 3A Overview

We begin by assessing the extent to which the failure to adjust for attribute meaning change is indeed a significant source of error using human retrospection data from Study 2B. To do so, we repurposed the different models from Study 2A and use them to assess different types of systematic errors in human retrospection. First, we used the Contemporaneous semantic projection model to capture a human rater who updates both brand and attribute meaning over time (Figure 1.7.2.1C). Although it was the most accurate model in terms of predicting original BAV ratings, its purpose here is that of a negative control, as human performance was found in Study 2A to be significantly worse than the Contemporaneous model

Second, we used the Attribute-Constrained semantic projection model to capture systematic errors in updating attribute meaning. This model seeks to capture errors arising from raters who apply an outdated meaning of an attribute, but where brand perceptions are updated over time in nominal terms. Specifically, we fixed the attribute meaning to some reference time while allowing brand meaning to be updated over time (Figure 1.7.2.1A). Third and finally, we use the Brand-Constrained semantic projection model to capture errors arising from raters who update attribute meaning over time, but whose brand knowledge is "stuck" in the particular reference time. Specifically, we allow attribute meaning to change over time but fix the brand meaning at a particular reference time (Figure 1.7.2.1B).

1.7.2 Study 3A Results

Comparing model fits across the three classes of models, we found that the Attribute-Constrained model using attribute meaning from 2020 performed the best (ΔLL=215 compared to a null model with no predictors, with 95% C.I. [170, 254]; Figure 1.7.2.1). In contrast, the results do not support the hypothesis that retrospection participants have the accurate attribute meaning in years past, as the Contemporaneous model performed poorly (ΔLL=34, 95% C.I. [15, 50]). It also does not support the hypothesis that participants simply had incorrect attribute meaning, rather than the incorrectly dated attribute meaning. If this were the case, the Attribute-Constrained model from 2020 should not outperform its 2010 or 2000 counterparts, nor other models such as the Brand-Constrained models. Instead, we observe a continuous decline from 2020 to 2001 for Attribute-Constrained model with the reference time of 2020, 2010, and 2001 respectively, consistent with the idea that retrospection participants did not sufficiently account for the changes in attribute meaning over time.

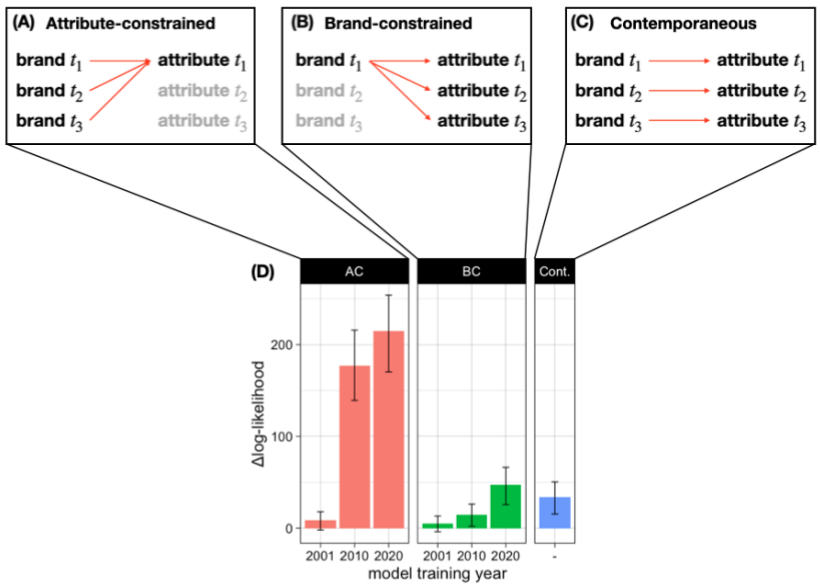

Figure 1.7.2.1 *Human errors in retrospective judgments of brand attribute association reflect failure to adjust for attribute meaning changes. (A) Attribute-constrained model allows brand perception to change over time but attribute meaning is fixed to a specific reference time. (B) Brand-constrained model allows attributes to change but fixes brand perception to a specific reference time. (C) Contemporaneous model allows both brands and attributes to vary across time. (D) Model goodness of fit comparison. The y-axis represents the log-likelihood improvement compared to a random null model. Error bars: bootstrap 95% C.I. with 10,000 resamples.*

1.7.3 Study 3B Overview

Building on findings from Study 3A showing that the use of outdated attribute meaning constitutes an important source of error in human retrospection, we next sought to test if correcting these errors can improve retrospection performance. To this end, we tested whether providing human participants with the correctly-dated attribute meaning, i.e., the attribute meaning corresponding to the reference year, as generated by our model would improve their retrospective judgments.

That is, one can explain to someone in 2020 what the word *gay* meant in 1950 by showing words close to its meaning in 1950, i.e., its 1950s synonyms. Here, there are two options in one's choice of synonyms in their 1950s or 2020 version. For example, if one were to describe *gay* in the language of 1950, the closest synonym will be *merry.* In the language of 2020, however, one would instead use *happy.* Given that our goal is to explain the meaning of older words to individuals in the present, we used synonyms in their modern sense, e.g., use *happy* to describe what *gay* meant in 1950.

Similar to the human retrospection experiment in Study 2A, the participants were asked to judge which brand out of a given pair is more associated with a specific attribute, e.g., *Social*, in a given year according to the BAV data. Unlike Study 2A, participants were provided with a set of related words that corresponds to one of two conditions. In the Past Meaning condition, participants were given the synonyms of the focal attribute from 2001. In the Present Meaning condition, which serves as a control, participants were given the synonyms of the focal attribute from 2020. For example, for *Visionary*, this corresponds to *Rugged* and *Stylish* in the Past condition, and *Innovative* and *Progressive* in the Present condition, reflecting substantial changes in the meaning of *Visionary* over the past two decades. On the other hand, *Upper Class* has almost the same set synonyms in the two time points, reflecting the fact that *Upper Class* has largely retained its meaning over time (See Appendix E for all synonyms used).

1.7.4 Study 3B Results

Access to Past Meaning by DWE Improves Human Retrospection. Comparing performance of the Past Meaning and Present Meaning conditions, we found that those in the Past Meaning condition performed significantly better than those in the Present Meaning condition (β=4.8%, t=3.008, p=0.003; Figure 1.7.4.1A). We further broke down the difference in accuracy by attributes (Figure 1.7.4.1B). Consistent with our hypothesis, participants in the Past Meaning condition performed better than those in the Present Meaning condition for the majority of the attributes. Thus, consistent with the idea that human retrospection errors at least partially reflect the use of outdated attribute definitions, we found that their retrospective performance improved when provided with the correctly-dated attribute meaning as generated by our model. In contrast, under the null hypothesis that our model does not capture attribute meaning as used by human participants to make judgments, we should observe no significant difference in retrospection accuracy between participants in the two conditions.

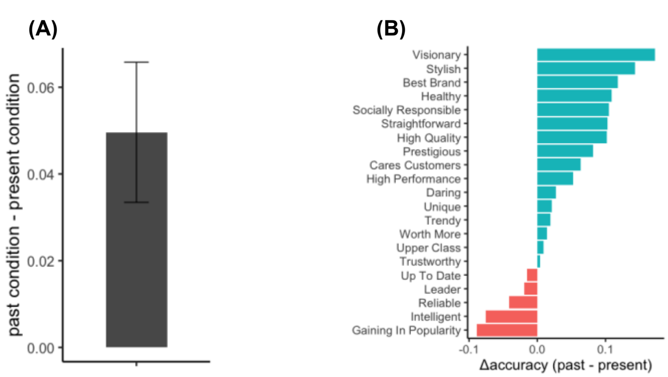

Figure 1.7.4.1 *Human participants' performance improved when provided with correctly-dated attribute*

meanings, assessed on (A) mean accuracy and (B) across attributes.

1.8 DISCUSSION

Whether launching a product, repositioning a brand, or changing a pricing model, the timing of the decision can play a crucial role in determining its success (D. A. Aaker, 1991; Carlson et al., 2019; Round & Roper, 2017; Villas-Boas, 2018). Compared to tangible attributes such as price and functionality, however, much less guidance is available for practitioners regarding potential biases and pitfalls that can threaten the validity of insights regarding the evolution of more intangible aspects of the brand such as brand image (Carlson et al., 2019; Gensch & Golob, 1975). In particular, a foundational challenge exists where comparisons across time may no longer be valid due to the fact that what is being measured in terms of the intangible aspects of the brand changes (Gensch & Golob, 1975; Rutz & Sonnier, 2011).

By introducing ideas and tools from NLP and specifically diachronic linguistics (Deo, 2015; Hamilton et al., 2016b, 2016a; Kutuzov et al., 2018), we address the issue at both conceptual and empirical levels. Conceptually, we introduce the distinction between real and nominal change to provide a conceptual framework for adjusting the measurements to account for the changes due to drifts in the measurement, and quantify the movements in semantic meaning using movements in a semantic space. This in turn, allows us to overcome the theoretical difficulty of meaning changes by using DWE to provide a common coordinate system in which we can represent the meaning of brands, attributes, and other context words, which can serve as semantic landmarks.

We further demonstrate that this distinction is not merely of theoretical interest but can have substantial implications. Specifically, we found that the movement of attribute meaning can vary greatly across attributes and time, with some staying roughly the same across time but others change significantly in meaning, such that nominal and real ratings can differ significantly for certain brands at certain times. Thus, in the same way that real purchasing power may be either over- or under-estimated if using nominal purchasing power (Diewert, 1988; O'Neill et al., 2017), overlooking this conceptual distinction for brand perception can also lead to inaccurate interpretations of the direction, speed, and magnitude of brand meaning changes.

The feasibility of this approach depends on a number of necessary conditions. First, it depends on the availability of a suitable set of natural language text that covers the time period in question. Moreover, the underlying corpus itself must contain information capturing the evolution of brand meaning and attribute meaning over time and allow for intertemporal comparisons. Our choice of the New York Times corpus reflects these concerns as its size, coverage, and readership have been relatively stable over the course of the several decades. It is unclear, however, how robust our proposed approach will be to cases where high quality text corpus is not available, for example those involving a niche industry or in emerging markets where the quality of available text may change suddenly. Overcoming these technical challenges will further allow managers to use user-generated content (UGC) to provide consumer perspectives that are not filtered through intermediaries such as the news media, which would especially be valuable in markets involving novel product categories or those characterized by rapid changing trends.

Second, our approach requires that the meaning of attributes can be identified and captured by corresponding vectors in the word embedding space (Deo, 2015; Hamilton et al., 2016b). In our approach, attribute meaning was recovered in a data-driven approach using BAV ratings for all brands on a particular attribute. An important underlying assumption, therefore, is that coverage and methodology of BAV ratings remains consistent over time in both how brands are selected, as well as associated sampling and data collection procedures. If this is not the case, recovery of attribute meaning will be biased by effects of these methodological changes. An alternative possibility is to directly extract changes in meaning of terms such as *social* without estimating from the data. This would be akin to how linguists capture changes in the meaning of words such as *gay*. A weakness of this approach is that the word *social* can have many different meanings, and that only some of them reflect changes when it comes to branding. Although not feasible at present due to prohibitive computational demands, future advances in diachronic NLP methods that consider word context, such as those in BERT and GPT, may make this a more viable approach.

Third and perhaps most importantly, we require that changes in attribute meaning be captured by time-indexed versions of associated high dimensional vectors. This is the critical condition that allows time-indexed vectors of attribute meaning and brand meaning to be paired across time. That is, how a brand would be perceived on a particular attribute using its meaning, i.e., its real attribute rating. Note that, unlike the first two conditions, which can be validated with nominal perception only, this third condition cannot, as it relies on the comparability of word vectors across time.

To validate our inference of attribute meaning changes over time, therefore, we contribute a novel method by building on existing validation approaches developed in diachronic linguistics (Kutuzov et al., 2017; Mihalcea & Nastase, 2012; Popescu & Strapparava, 2015). Evaluating the validity of inferences from natural language is known to be challenging and becomes even more so when time is involved. Diachronic linguistics researchers have developed various sophisticated methods to address the limitation that we do not have human participants in the past to consult, ranging from predicting the time period of word within a specific context using detected semantic change to validating against synthetic corpus with artificially introduced semantic change. Here, the fact that we have a "gold standard" in the BAV ratings allows us to ask if participants' performance improved when they were provided with DWE-generated meaning of attributes corresponding to the reference date. That is, we had the advantage of being able to assess improvement by comparing against the benchmark of the original BAV ratings.

1.8.1 Managerial Implications

Our findings have a number of implications for managerial actions related to branding and brand building. Due to changes in attribute meaning, the specific time chosen to be the reference can matter greatly. Depending on the specific managerial goals, a brand manager may want to either reference to the original meaning or the contemporary, i.e., present, meaning. For instance, a brand manager at Microsoft may choose to index to the original meaning if the goal is to reposition the brand on the original meaning of *Visionary*. In this case, they should use the real *Visionary* rating of Microsoft with 2007 as the reference time to interpret changes in brand perception.

Alternatively, the brand manager may wish to understand how Microsoft has evolved on the most up-to-date meaning of *Social*. In this case, they should use the real *Social* rating with the most recent year as the reference time. Yet another use is when a firm wants to know how the attributes have changed over time. For example, a social media company may want to know how the meaning of *Social* has evolved over time to identify managerially relevant trends. In this case, managers should use the semantic trajectories of *Social* for insights on any changes that may have taken place.

Distinguishing between real and nominal ratings can also have an important impact on our interpretation of the time course of changes. Managers may select different actions depending on whether the change is a sudden one or a slow, accumulated one (D. A. Aaker, 2011; Carlson et al., 2019; Round & Roper, 2017; Villas-Boas, 2018). For example, in our earlier example, Microsoft experienced a sudden increase in ratings of *Visionary* around 2008 in nominal terms. Interpreted naively, a manager could mistakenly attribute these changes to branding actions that Microsoft conducted prior to that time. Converting these to real changes using either pre-2008 or post-2008 attribute meaning, however, the sudden improvement can almost entirely be explained by changes in the meaning of *Visionary*. Indeed, using its 2020 meaning, Microsoft actually experiences a slow decline. Thus, risks of misinterpretation are present even if Microsoft only compared itself to rival brands, such as Apple. Using only nominal ratings, both brands appear to be close competitors throughout the past two decades. However, adjusting for meaning change, it is apparent that Microsoft once had a noticeable lead on the attribute but steadily lost this edge.

While the discussion above considered only a single attribute, brand perception inherently concerns multiple attributes (D. A. Aaker, 2011; J. L. Aaker, 1997a; De Chernatony, 2010; Keller, 2003). Therefore, an important benefit of an NLP approach to deriving real changes in brand perception is the scale it affords in allowing us to generalize to multiple attributes, including those that may not have existed in the past. To the extent that our method provides more scientifically valid comparisons of where brands sit on intangible attribute space, managers can use this information to obtain a more complete and up-to-date understanding of their brand as possible.

More generally, the conceptual distinction between real and nominal is relevant in any consumer and managerial phenomenon where time plays an important role (Carlson et al., 2019). Product adoption curves inherently have a temporal component (Bass, 1969; Rogers, 1976). Consumer trends and fads evolve over time (Bikhchandani et al., 1992). Some forms of marketing strategies such as nostalgia marketing leverage the meanings in different time periods (Pecot et al., 2019). Consumer loyalty toward brands can also shift over time in response to various factors (Brakus et al., 2009; Keller, 2016; Reichheld, 2001). In these and other applications, our approach opens the door for consumer researchers to use NLP tools to supplement traditional sources of data and analytical methods.

1.8.2 Future Directions

Taken together, our study can be viewed as part of a broader effort to incorporate large-scale text analysis and NLP to inform consumer research, and opens up a number of novel avenues for additional inquiry (Berger et al., 2020, 2021, 2023; Berger & Milkman, 2012; Büschken & Allenby, 2016; Cascio Rizzo et al., 2023; Gabel et al., 2019; Hovy et al., 2021; Humphreys & Wang, 2018; Kim et al., 2024; Kronrod et al., 2023, 2023; J. K. Lee, 2021; J. Liu & Toubia, 2018; Netzer et al., 2019; Packard et al., 2023; Packard & Berger, 2021; Rocklage et al., 2023; Tirunillai & Tellis, 2014; Villarroel Ordenes et al., 2019). For example, an important question concerns how managers should adjust for attribute meanings that may vary across different cultures, in order to obtain valid cross-country comparisons. Prior studies suggest that brand perception may not only differ quantitatively but also qualitatively, such that while Spanish consumers distinguish between a brand's *Passion* in ways that American and Japanese consumers do not, Americans distinguish between a brand's *Ruggedness* in ways that their Spanish and Japanese counterparts do not (J. L. Aaker et al., 2001). Such conclusions, however, depend on inferences based on a limited number of time points. It is possible, therefore, that each culture, at some point in the past, did perceive these attributes. Another, non-mutually exclusive, possibility is that one country influences others in how attributes are perceived, e.g., if US firms influenced how other nations' perception of *Rugged* (Steenkamp, 2019; Zhou et al., 2008).

Beyond documenting the existence of changes, progress on these questions also provide important insights into processes underlying the evolution of attribute meaning changes and changes in brand perception, which in turn would be valuable in predicting when and why changes occur. For example, what caused the meaning of *Visionary* to shift from *Rugged* to *Innovative*? When did that shift start to emerge and then stabilize? It is likely that only some of these forces are directly under the control of individual firms. Nevertheless, an understanding of the forces underlying changes in (real) brand perception will be critical to help managers prepare and respond to these changes.

Indeed, the fact that consumer knowledge can change over time, sometimes dramatically so, can pose an important but underappreciated challenge to efforts to test the replicability of research findings, for both applied and basic research (Easley et al., 2000; Lynch et al., 2015; Urminsky & Dietvorst, 2024). For example, a study could fail to replicate because certain elements of the study are no longer applicable, e.g., stimuli, consumer needs, etc. If the study is about innovation, a failure to adjust for these changes, e.g., by directly using 1990 products, would likely no longer be a reasonable test today.

Moreover, given the speed of advancement in NLP and Large Language Models (LLM) over the past decade, our ability to be able to answer these questions will undoubtedly progress rapidly in the future. These will undoubtedly improve the accuracy, precision, temporal resolution, and scope of our predictions. In particular, future advances will likely allow us to make use of context-aware models such as BERT or GPT (Gillioz et al., 2020), which allow for finer distinction of word meaning by disambiguating between different types of word contexts. Although highly powerful in that, the high computational burden and size of input data necessary to train these models have at present restricted users to rely on pre-trained models that have limited transparency and flexibility. In contrast, we chose a DWE approach due to its relatively light computational burden and ability to perform on relatively small text corpora, which allows us to tailor model training to the particular goals of our analytical approach, such as choosing a specific corpus to use or restricting training to specific time periods.

Finally, our study highlights the continuing relevance of methodologically rigorous consumer surveys for researchers and managers. Indeed, despite access to ever more sophisticated computational tools and growing availability of large-scale text corpora, approaches such as ours critically depend on presence of high-quality, longitudinal consumer surveys such as those by BAV (Berger et al., 2022; Deo, 2015; Garg et al., 2018). In the case of inflation, the issues are of sufficient macroeconomic and political significance that national governments invest considerable financial and physical resources in measuring core underlying concepts (Diewert, 1988; O'Neill et al., 2017). Given the lack of such a large-scale, coordinated effort in consumer research, our work helps to bridge this important gap. Thus, despite the fact that our inferences and validation findings are far from perfect and remain limited in the number of brands and the time intervals studied, given the importance of and challenges involved in the questions we study, even the modest advances we make here constitute a significant and productive step that can be built upon by future work.

Chapter 2

Windfall, Similarity, and Mental Accounting

This chapter examines how mental representations of money and product relationships jointly shape consumer spending decisions. Building on my broader framework, it introduces similarity as a key mechanism that constrains how consumers mentally allocate and spend windfall gains. By formalizing the role of similarity in mental accounting, the chapter moves beyond existing explanations to generate more precise predictions about when windfall effects arise and how consumers decide which purchases to make.

2.1 INTRODUCTION

Imagine that you just won a $50 raffle from a gym you visited before. Imagine also that you have been considering going to a concert for a while. How likely would you spend the $50 to buy the concert ticket? Now suppose that instead of a raffle, you received the extra $50 from working overtime in the gym. How would your purchase intention of the concert ticket differ in this situation? What about if you won the raffle from a theater instead of a gym?

Consumers often receive various forms of windfalls, defined as the unexpected money that consumers receive from different sources in their daily life. They might receive a surprise rebate from online shopping, win a bet at the casino, inherit an unexpected fortune of family relatives, or receive a tax refund that they have long forgotten about.

Research has shown that despite the fungibility of money, consumers tend to spend windfall money differently than hard-earned money. For instance, Arkes et al. (1994) found that the participants who read about winning a radio contest were more likely to spend the money on a color TV than those who read about receiving the same amount of money but from working overtime. O'Curry & Strahilevitz (2001) documented that people are especially likely to choose hedonic products (e.g., a massage) rather than utilitarian ones (e.g., a bike helmet) when spending a lottery prize, compared to when making an ordinary purchase. This phenomenon is known as the *windfall effect* in the literature. While the existence of windfall effects finds a lot of support in the existing literature, little is known about when and why these effects emerge as well as how exactly windfall money is spent.

One set of explanations for the windfall effect is that consumers have a higher marginal propensity to consume for windfall money in general (Arkes et al., 1994). We will refer to this theory as the *general HMPC (Higher Marginal Propensity to Consume) theory*. Under this theory, windfall money is posited to have a relatively lower subjective value than hard-earned money. In other words, the disutility of giving up a dollar of hard-earned cash is higher than the disutility of giving up a dollar received unexpectedly.

As a consequence, consumers more readily spend money received through windfalls than equivalent amounts received through labor.

Several works have replicated this higher spending propensity of windfall money but also documented that the spending appears to increase mostly for hedonic products. For example, Heilman et al. (2002) found that consumers at a supermarket who were given a $1 coupon purchased more items than those who were not and that the additional spending was mostly on treats rather than necessities. We will refer to this as the *hedonic HMPC theory*. These authors propose that receiving a windfall elevates mood and that positive mood subsequently increases purchases of treats. Similar findings, but due to a slightly different proposed process, have been documented by O'Curry & Strahilevitz (2001) and O'Donnel and Evers (2024). Specifically, these sets of authors assume that either the consumption of the good (O'Curry & Strahilevitz, 2001) or the spending of money on hedonic goods (O'Donnell & Evers, 2024) elicit guilt, and show that windfalls alleviate this guilt and therefore lead to increased spending for hedonic goods. While these three explanations differ to some degree in the underlying mechanisms, all predict that windfalls will either increase spending in general or spending on hedonic products, regardless of the source of this windfall.

However, subsequent studies about the windfall effect are not fully consistent with the HMPC view, thus hinting at a more nuanced picture of the windfall effect. In particular, a number of findings suggest that the relationship between the source of the money and target items, in addition to whether the source is a windfall, affects how consumers spend the money. Levav & Mcgraw (2009) found that consumers avoid spending negatively tagged windfall on hedonic products. For example, consumers are less likely to spend gift money from a wealthy brother (vs. a financially strapped brother) on hedonic items like ice-cream sundae. Helion & Gilovich (2014) showed that consumers are more likely to use gift cards to purchase hedonic (vs. utilitarian) items. Perhaps the most suggestive evidence can be found in Reinholtz et al. (2015), where the authors demonstrated that consumers are more likely to spend retailer-specific gift cards on items more typical of the retailer. For example, consumers are more likely to spend a Levi's gift card on jeans as opposed to sweaters compared to an open-loop gift card that can be spent in many retailers. In the cases above, it seems that the more "similar" the source and target items are, the more likely consumers are to spend the windfall.

Indeed, the findings above are consistent with theories of mental accounting, in which consumers mentally assign money into different "accounts," and that money in different accounts is not perfectly fungible (R. H. Thaler, 1999b). We will refer to this explanation as the *MA (Mental Accounting) theory*. Under this perspective, the fact that consumers are more likely to spend retailer-specific gift cards on more typical items of the retailer can be explained by positing the existence of a retailer account in consumers' mind. Both the retailer-specific gift cards and the typical items of the retailer are grouped in this account, whereas open-loop gift cards, which can be used in many different retailers, are not. The mental accounting theories then argue that, as money in the same mental account is seen as more fungible than those in different accounts, consumers are more likely to spend retailer-specific gift cards (vs. open-loop gift cards) on items more typical of the retailer. Similarly, the fact that consumers are less likely to spend negatively-tagged windfall on hedonic items can also be explained by positing the existence of a hedonic account. Since only positively-tagged but not negatively-tagged windfalls are grouped in the hedonic account along with hedonic items, consumers are less likely to spend the negatively-tagged windfalls on hedonic items. In other words, many previous studies about windfall spending can be understood through how these funds are categorized.

Some consumer researchers have explicitly applied mental accounting theories to explain windfall spending. For example, Hodge & Mason (1995) argued that if consumers have a set reference price before purchasing, they have already allocated this amount to a "spend account." They then asserted that spending from windfall savings (vs. hard-earned savings) would result in a more firmly set reference price, which explains the difference in spending. Abeler & Marklein (2017) proposed a formal model of mental accounting to explain why restaurant guests who received a beverage-specific voucher spent more on beverages than those who received a general-purpose voucher of the same value, even when the voucher amount was smaller than the minimum amount they typically spent. In their model, consumers optimize their consumption with a budget constraint. Importantly, consumers are modeled as utility maximizers who vary in terms of how rigid their mental accounts are, ranging from "fully rational" consumers, who allocate budgets regardless of the voucher label and thus exhibit no mental accounting, to "naive" consumers, who ignore the possibility of allocating the windfall coming from beverage-specific vouchers to other items and thus exhibit maximal mental accounting.

However, the mental accounts mentioned in these studies are often ad hoc and, in many cases, appear tautological. For example, the "spend account" explanation proposed in Hodge & Mason (1995) is circular in nature: Windfalls are spent more because they are grouped in the spend account, yet they are grouped in the spend account precisely because they are spent more in the first place. This reflects the issue that the formation process of mental accounts is often vague and underspecified in the literature. Little is known about when and under what conditions money and products will be grouped in the same mental account. Even simple questions such as "is money obtained from winning a lottery grouped in the same mental account as money spent on buying a TV?" are left unanswered by the current explanation. As a result, mental accounting-based theories as described in the literature, while being flexible enough to fit the experimental data well, have limited predictive power and can only make relatively coarse predictions.

To summarize, the current literature on windfall effects demonstrates that consumers appear to spend windfalls differently than hard earned money, however, it is unclear *why* these windfall effects emerge. The general and hedonic HMPC theories, being more descriptive than explanatory, simply states that consumers tend to spend windfall money more readily, and especially on hedonic items, while the MA theory relies on mental accounting as an explanation without specifying what mental accounting entails. In addition, both of these perspectives are consistent with data reported in the extant literature. However, based on the discussion above they would make diverging predictions were we to generalize outside of the contexts studied. According to the first perspective, windfalls, regardless of the source, would always lead to higher spending, especially on hedonic goods. According to the second perspective, the crucial factor instead is the degree to which the source of a windfall matches the target under consideration. Specifically, the category (mental account) in which a consumer gains the money would be predicted to increase spending *in that category only* regardless of whether the goods are hedonic or utilitarian and whether the money is a windfall or hard-earned. As a consequence, if we would like to predict, or even influence, windfall spending, it is crucial to understand the process through which windfall spending differs from regular spending.

We argue that part of the difficulty in dissociating the HMPC and MA theories stems from insufficient stimulus sampling. In the windfall literature, researchers often only selected a small number of items to represent hard-earned/windfall money and the spending items under consideration. For example, in the first study of Arkes et al. (1994) the authors represented hard-earned money by working overtime at a supermarket and windfall money by a sweepstake from a radio station. The spending items under consideration were a color TV. With such a small number of items, it is challenging to differentiate between the two competing theories that explain why consumers are more likely to purchase a color TV with money earned from a radio station sweepstake than with money earned from working overtime at a supermarket: Is the behavioral difference due to the windfall nature of the sweepstake? Or because the sweepstake is grouped in the same mental account as the color TV?

In the current research, we investigate the explanatory power of the three competing theories by testing their diverging predictions on windfall spending across four lab studies and one field experiment. To this end, we make the following conceptual and methodological innovation.

As a first step, we transform mental accounting theory into a generative framework for consumer behavior by formally integrating similarity into the process of mental account formation. As noted earlier, the current form of mental accounting theory cannot generate specific predictions, making it difficult to evaluate against the two HMPC theories. To address this issue, we build upon psychology literature that argued that mental accounting can be viewed as an instance of human categorization process (Henderson & Peterson, 1992), of which similarity is known to play an important role

(Goldstone, 1994; Kruschke, 1992). Under this view, the more similar two items are, the more likely they will be grouped in the same mental account. If such a view is correct, it then follows that mental accounting theory is actually able to make more granular predictions than its current use. More specifically, the theory would predict that the similarity between the source and target will increase the purchase intention of the target item, regardless of whether the money is hard-earned from windfall and whether the products are hedonic or utilitarian. This allows us to compare the explanatory power of the three theories directly.

Second, we adopted extensive stimulus sampling in all of our experiments to address the alternative explanations that previous studies have difficulty ruling out. Specifically, we sampled the products from a wide range of product categories defined externally. In Study 1 and 2, we sampled 496 pairs of products as sources and targets of the windfall. In Study 3A, we sampled 72 pairs of places and products. Even in the field experiment where we were heavily constrained by the logistics of setting up an actual shop, we still sampled 24 pairs of products. By sampling this way, we ensure that there is enough variance along both the dimensions of hedonicity and similarity for us to compare the two theories. Additionally, it helped mitigate potential confounding factors that could arise from sampling a smaller, less diverse set of products.

2.2 METHODS

2.2.1 Participants and Prodecure for Testing Mental Accounting Prediction

This study was pre-registered at https://aspredicted.org/vwnp-zt3r.pdf. We recruited 1,004 participants from Prolific and (following our preregistration) excluded 8 participants who failed the simple English language check at the beginning of the survey, leaving 996 participants for analysis (M_{age}=40 years; 56% female, 42% male, 3% identified as non-binary or preferred not to say). These participants were randomly assigned to either the earmarked condition or the unearmarked condition. In both conditions, participants were asked to imagine that they received a rebate of $50 from shopping. Additionally, for those in the earmarked condition we also explicitly provided the source of this rebate.

Therefore, in both conditions the participants imagined receiving shopping rebates of the same amount, with the only difference being that those in the earmarked condition were aware of the source of this windfall while those in the unearmarked condition were not. Then, all participants indicated their purchase intention for the target item on a 7-point scale, 1 being extremely unlikely and 7 being extremely likely. Source and target items were chosen from a set of 32 consumer products, manually selected from eight Amazon categories: Electronics, Computers, Sports & Outdoors, Home & Kitchen, Pet Supplies, Music & Games, Clothing, and Food. Each product could serve as both the source and the target item, resulting in a total of 496 possible source-target pairs.

Each participant completed 32 trials. Participants in the unearmarked condition answered one question about each of the 32 consumer products. Participants in the earmarked condition answered one question about each of the 32 consumer products as the target item with a randomly selected different source item (with replacement), precluding the source being the same as the target.

For our measure of similarity, we employed semantic similarity. Semantic similarity refers to the degree to which two words convey similar meanings according to their co-occurrence statistics in the text corpus. In consumer research, semantic similarity has been applied to analyze user-generated content, track brand perception shifts, and investigate cultural differences in language use (Berger et al., 2022). Here we use Word2Vec (Mikolov et al., 2013) trained on the New York Times corpus from 1996 to 2023 to compute the semantic similarity. Semantic similarity scores range from 0 (not similar at all) to 1 (extremely similar).

2.2.2 Participants and Prodecure for Testing Marginal Explanation Power of Mental Accounting

This study was pre-registered at https://aspredicted.org/vwnp-zt3r.pdf. We recruited 1,004 participants from Prolific and (following our preregistration) excluded 8 participants who failed the simple English language check at the beginning of the survey, leaving 996 participants for analysis (M_{age}=40 years; 56% female, 42% male, 3% identified as non-binary or preferred not to say). These participants were randomly assigned to either the earmarked condition or the unearmarked condition. In both conditions, participants were asked to imagine that they received a rebate of $50 from shopping. Additionally, for those in the earmarked condition we also explicitly provided the source of this rebate.

Therefore, in both conditions the participants imagined receiving shopping rebates of the same amount, with the only difference being that those in the earmarked condition were aware of the source of this windfall while those in the unearmarked condition were not. Then, all participants indicated their purchase intention for the target item on a 7-point scale, 1 being extremely unlikely and 7 being extremely likely. Source and target items were chosen from a set of 32 consumer products, manually selected from eight Amazon categories: Electronics, Computers, Sports & Outdoors, Home & Kitchen, Pet Supplies, Music & Games, Clothing, and Food. Each product could serve as both the source and the target item, resulting in a total of 496 possible source-target pairs.

Each participant completed 32 trials. Participants in the unearmarked condition answered one question about each of the 32 consumer products. Participants in the earmarked condition answered one question about each of the 32 consumer products as the target item with a randomly selected different source item (with replacement), precluding the source being the same as the target.

For our measure of similarity, we employed semantic similarity. Semantic similarity refers to the degree to which two words convey similar meanings according to their co-occurrence statistics in the text corpus. In consumer research, semantic similarity has been applied to analyze user-generated content, track brand perception shifts, and investigate cultural differences in language use (Berger et al., 2022). Here we use Word2Vec (Mikolov et al., 2013) trained on the New York Times corpus from 1996 to 2023 to compute the semantic similarity. Semantic similarity scores range from 0 (not similar at all) to 1 (extremely similar).

2.2.3 Participants and Prodecure for Reconcilating with Classic Findings 1

This study was pre-registered at https://aspredicted.org/2775-wvfx.pdf. We recruited 1,205 participants from Prolific and randomly assigned them to either the windfall condition or the hard-earned condition. As outlined in the pre-registration, we excluded 6 participants who failed the simple English language check at the beginning of the survey, leaving 1,199 participants for analysis (M_{age}=41 years; 59% female, 38% male, 2% identified as non-binary or preferred not to say).

Participants in the windfall condition imagined that they won a $50 sweepstake from a previous visit to a [place]. Participants in the hard-earned condition imagined that they earned an extra $50 from working overtime at the same place. Participants in both conditions then indicated how likely they were to spend the money to purchase an item. We sampled six locations and twelve items across various domains commonly encountered by consumers in their daily lives. We measured both semantic similarity and human-rated similarity in the same way as before.

2.2.4 Participants and Prodecure for Reconcilating with Classic Findings 2

This study was pre-registered at https://aspredicted.org/sp44-5p7m.pdf. We recruited 603 participants from Prolific and randomly assigned them to one of the three conditions: the control condition, the windfall condition, and the prize condition. As outlined in the pre-registration, we excluded 7 participants who failed the simple English language check at the beginning of the survey, leaving 596 participants for analysis (M_{age}=40 years; 56% female, 42% male, 2% identified as non-binary or preferred not to say).

Participants evaluated ten items sourced from Study 2 of O'Curry & Strahilevitz (2001): athletic shoes, textbook credit, scientific calculator, book bag, bike helmet, Tower Records voucher, dinner for two, tickets to a pop concert, a one-month health club membership, and a massage. All items were priced at $60.

In the control condition, participants rated their purchase intentions for each item on a 7-point scale, considering the $60 price. In the windfall condition, participants imagined receiving $60 from winning a lottery and rated their likelihood of spending this money on each item. In the prize condition, participants imagined winning a lottery that allowed them to choose any one item (worth $60) as a prize and rated their likelihood of selecting each item. After rating purchase intentions, participants in all conditions indicated how similar each item was to a lottery and separately rated each item on hedonic and utilitarian dimensions, all using 7-point scales.

2.2.5 Participants and Prodecure for Field Experiment

This study was pre-registered at https://aspredicted.org/gmnv-r54g.pdf. With the help of a group of undergraduate research assistants, we set up a mini shop on the campus of a large West Coast university. The experiment consisted of two stages. In Stage 1, potential participants were invited to purchase three packs of Oreo cookies and/or a set of three university-themed stickers, each priced at $1. Only those who purchased both items ($2 total) advanced to Stage 2. In Stage 2, participants were informed that they were now taking part in an experiment and consented verbally. They spun a physical wheel that randomly assigned them a $1 coupon labeled either "Oreo cookies" or "stickers," with equal probability. Participants were told they received the coupon because they had purchased the item named on the coupon. Next, participants were shown a display of twelve new items (La Croix, Red Bull, Snickers, Juice, Gum, Mentos, Notebook, Magnets, Pencil, Stress ball, Bookmark, Keychain) and were invited to use their $1 coupon toward the purchase of any item. After selecting their desired item(s), participants paid the remaining balance after the coupon was applied. We ultimately recorded 278 purchase records. We measured the similarity between each item with Oreo and stickers by surveying a separate group of 50 participants on Prolific on a 7-point scale.

2.3 RESULTS

2.3.1 Mental Accounting Theory Predicts How Consumers Spend Windfall Money

As we predicted that similarity between source and target is what drives windfall effects (H1), we ran a linear regression predicting the difference in purchase intention between the earmarked and the unearmarked condition using the semantic similarity between the source and target items. The data revealed that higher similarity indeed led to higher purchase intention (b=0.94, t(15,838)=16.00, p<0.001, see Figure 2.3.1.1). These results show that consumers are more likely to spend windfall money on a product when it is more similar to the source of the money, consistent with the logic of mental accounts as spending categories.

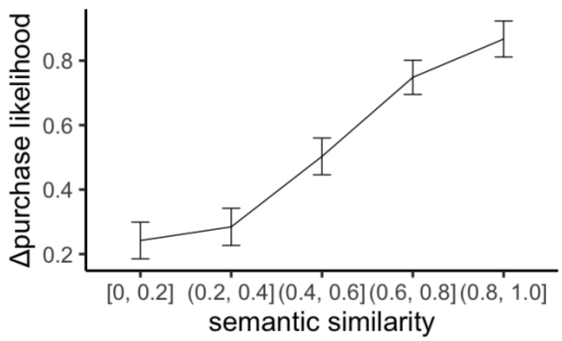

Figure 2.3.1.1 *Purchase likelihoods are higher for products that are more similar to the source of the windfall. Error bars represent the standard error of the mean difference in purchase likelihood for each quintile.*

2.3.2 Mental Accounting Theory Predicts Windfall Spending Beyond Baseline Theories

To test the general HMPC theory, we ran a mixed-effects model predicting purchase intention. The general HMPC model included a dummy variable for windfall as the independent variable, with hedonic and utilitarian perception as covariates, and a random intercept to account for individual differences among participants. This analysis tests whether, when controlling for product characteristics, there is an effect such that windfall money leads to higher purchase intentions. Contrary to predictions of HMPC, we found a significant *negative* effect of windfall, represented by a negative estimate for the dummy (b_{windfall}=-0.52, $t(294.6)$=-4.35, p<0.001).

Next, we turned to test the hedonic HMPC theory. This model can be seen as an extension to the general HMPC model such that we also allow for interaction effects between windfall and the degree to which a product is hedonic and/or utilitarian. The hedonic HMPC model therefore included a dummy variable for windfall, interaction between the windfall dummy and hedonic and utilitarian perception as the independent variable, hedonic and utilitarian perception as covariates, and a random intercept to account for individual differences among participants. Contrary to the prediction of hedonic HMPC, we observed a *negative* interaction between the windfall dummy and hedonic perception ($b_{\text{interaction}}$=-0.12, $t(5856.3)$=-4.50, p<0.001). In other words, if anything, the more hedonic a product is perceived to be, the smaller the effect of windfall on purchase intentions is.

After exploring the two HMPC theories, we next turned to MA. The MA model included similarity between the source and target item as the independent variable, with hedonic and utilitarian perception as covariates, and a random intercept for participants. Consistent with the prediction of the theory, we found a positive effect of similarity ($b_{\text{similarity}}$=0.24, $t(4905.9)$=14.8, p<0.001) such that the higher the similarity was between the source and the target, the more participants indicated they were likely to purchase it.

Finally, we compared the explanatory power of the three models against each other using goodness of fit metrics as measured by the Akaike Information Criterion (AIC). Lower AIC values indicate a better-fitting model. The best-fitting model was the MA model, with an AIC of 22,543. The hedonic HMPC model had an AIC that was 58 points higher, and the general HMPC model had an AIC that was 136 points higher. Following Burnham & Anderson (2004) and our pre-registered analysis plan, we used an AIC difference greater than 10 points as the threshold for determining a meaningful difference between models. Models with an AIC difference greater than 10 were considered substantially worse than the best-fitting model. Thus this result presents strong evidence that the MA model is the best-fitting model of the three.

2.3.3 Comparison with Classic Mental Accounting Findings 1

To reconcile the results with Arkes et al. (1994), we first note one observation: the sources of windfall in the study are more "fun" than those of hard-earned money. Specifically, the windfall came from winning a raffle held by a radio station whereas the hard-earned money came from working at a retail store. This means that the conditions did not only differ in whether the source was a windfall or not, but also across a variety of other associations with that money. As a consequence, if any of those associations is more similar to the target good, it is possible that what appears to be a windfall effect (higher purchase intentions in the windfall condition) may actually be a consequence of higher similarity between source and target in the windfall condition. Looking at Arkes et al. (1994) specifically, there is some evidence that this may have been the case. In Study 2, all participants indicated purchase intentions for a color TV. In the windfall condition, the source of the money was a radio station, which is much more similar to a color tv (0.55 semantic similarity score according to Word2Vec) than the retail store from the control condition (0.37 semantic similarity score according to Word2Vec). Taken together with results in Study 1 and 2, this raises the possibility that the underlying mechanism of the well-known windfall effect could actually be (partially) due to similarity rather than the distinct nature of being windfall money.

As the design of Arkes et al. (1994) does not allow one to tease apart the effects of windfall from the effects of similarity, Study 3A is a conceptual extension of the original study, designed to further dissociate similarity from the nature of the money (i.e., windfall vs. hard-earned). We do so by holding the source of money identical and only varying whether or not that money is a windfall. Additionally, we vary both source as well as target to allow us to estimate the effect of similarity independent of whether the source was a windfall or not.

First, we tested whether the general HMPC, i.e., one of the most widely accepted explanations for windfall effects so far, explains consumers' purchase intention. To do so, we ran a mixed-effects model predicting purchase intention. The general HMPC model included a dummy variable for windfall as the independent variable and a random intercept for individual participants. If general HMPC is correct, then we should see a positive effect of the windfall dummy. However, contrary to this prediction, we found an insignificant, directionally negative effect of the windfall dummy (b_{windfall}=-0.033, t(1197.00)=-0.52, p=0.60). In other words, money being from a windfall, if anything, seemed to *reduce* purchase intentions.

Next, we tested whether our proposed MA theory is better at explaining the effects of windfalls on purchase intentions. The MA model included similarity as the independent variable. If the MA theory is correct, then we should see a positive effect of similarity. Consistent with this prediction, we found a positive effect of similarity ($b_{\text{similarity}}$=0.154, t(1389)=23.3, p<0.001). In other words, the more similar a source of money is to the target item, the more likely participants were to purchase it.

Finally, we tested which of the two models better fit the experimental data. We found that the MA model was the better-fitting model, with an AIC of 58,773 versus an AIC that was 529 points higher for the general HMPC model. This presents strong evidence that the model based on the MA theory fits the data better than the model based on the general HMPC theory.

Importantly, results of Study 3A can also explain why, despite the null effect of windfall, previous studies may have documented higher purchase intention for windfall vs. hard-earned money. The reason lies in insufficient stimulus sampling which then led researchers to mistakenly attribute the observed difference in purchase intention to windfall. This is neatly illustrated in Figure 2.3.3.1A.

(A) (B)

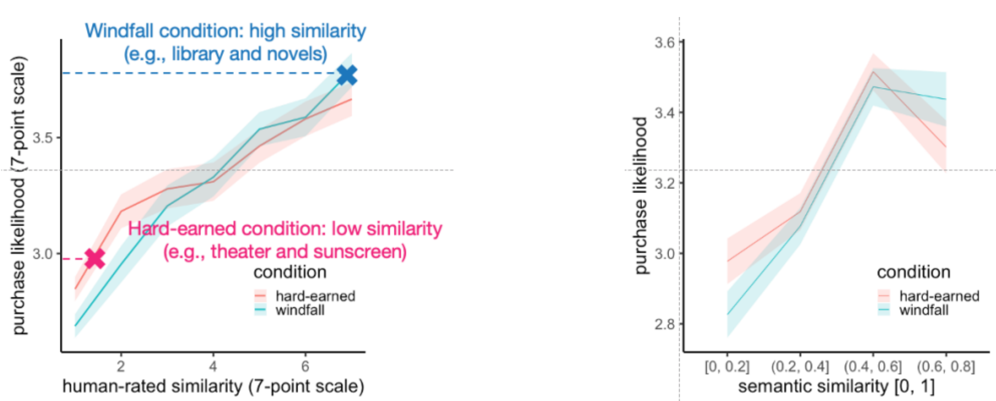

Figure 2.3.3.1 *Purchase intentions are higher for source-target pairs that are more similar as measured using (A) human-rated similarity and (B) semantic similarity. Panel A also shows how it is possible to "replicate" the results of Arkes et al. (1994) even if windfall itself does not increase purchase intention. By selecting more similar source-target pairs for the windfall condition and less similar pairs for the hard-earned condition, one would observe differences in purchase behavior between the two conditions. However, these differences would be driven by similarity, not by the mere effect of receiving a windfall.*

More specifically, if, when testing for windfall effects, researchers select only pairs of items and sources that are high in similarity (e.g., library and novels), and for the hard-earned conditions select only pairs of items that are low in similarity (e.g., theater and sunscreen), then we would still observe that consumers in the windfall condition appear to have higher purchase intention than those in the hard-earned condition. (M_{purchase}=4.08 for library and novels; M_{purchase}=2.86 for theater and sunscreen; Figure 2.3.3.1B).

In fact, if researchers would have sampled high-similarity pairs for the hard-earned conditions and low-similarity pairs for windfall conditions, they would instead find "reverse windfall effects" where consumers appear to spend *less* when having windfall money vs. hard-earned money. This would be akin to selecting gym and running shoes for the hard-earned condition (M_{purchase}=4.06) and pharmacy and concert tickets for the windfall condition (M_{purchase}=3.00).

2.3.4 Comparison with Classic Mental Accounting Findings 2

A similar logic may apply to another classic demonstration of windfall effects featured in O'Curry & Strahilevitz (2001). Specifically, O'Curry and Strahilevitz argue that consumers are more likely to spend windfall money on hedonic goods (HMPC). However, in the previous find of the same chapter, we found no interaction between hedonic perception and windfall conditions. A closer analysis of the two studies reveals a key methodological difference: O'Curry & Strahilevitz (2001) might have inadvertently tested a slightly different hypothesis.

Specifically, the authors did not directly compare windfall *spending* with the spending of hard-earned money. Instead, participants in the windfall condition selected their preferred item to win as a prize in a lottery, whereas those in the hard-earned condition chose items they preferred to purchase with their own money. In other words, O'Curry & Strahilevitz (2001) compared what participants prefer to receive for free, versus what they prefer to spend their money on. Therefore, the difference observed by O'Curry & Strahilevitz (2001) might result from the distinction between choosing and purchasing, rather than reflect a pure windfall effect (O'Donnell & Evers, 2019).

To examine this possibility, we conducted a conceptual replication of O'Curry & Strahilevitz (2001), directly comparing choice versus purchase. We tested two predictions: first, when measuring choice in the windfall condition, we expected to replicate their finding that consumers prefer hedonic items as lottery prizes. Second, when measuring purchase in the windfall condition, we anticipated replicating our own Study 2 results, i.e., consumers do not inherently prefer hedonic items but rather items perceived as similar to the windfall source.

We conducted two distinct sets of analyses to clearly address our research questions. First, we aimed to replicate the original findings from O'Curry & Strahilevitz (2001). For this replication analysis, we used only data from participants in the control condition and the prize condition. We compared two models: the hedonic HMPC model and the MA. The hedonic HMPC model predicted preference (purchase intention in the control condition, and likelihood of choosing in the price condition) purchase intention based on the experimental condition (control vs. prize), hedonic and utilitarian perceptions of items, interactions between condition and these perceptions, and random intercepts for participants. The MA model predicted purchase intention based on the experimental condition, similarity between the source of the windfall and target items, the interaction between condition and similarity, hedonic and utilitarian perceptions of items, and random intercepts for participants.

The hedonic HMPC model revealed a significant positive interaction effect between the prize condition and hedonic perception of items ($b_{interaction}$=0.096, $t(3888)=3.4$, $p<0.001$), indicating that participants preferred more hedonic items as lottery prizes. Interestingly, there was an unexpected significant negative main effect of the prize condition (b_{prize}=-1.1, $t(2894)=-4.7$, p<0.001) indicating that participants were less likely to select the items in general as a prize vs. expected to be willing to spend money on them.

In the MA model, we observed a positive main effect of similarity on purchase intention ($b_{similarity}$=0.2, $t(3419)=8.3$, p<0.001). A goodness-of-fit comparison indicated that the MA model provided a better fit to the data compared to the hedonic HMPC model (difference in AIC=57).

These initial findings replicate the results from O'Curry & Strahilevitz (2001), showing that consumers prefer hedonic items when choosing lottery prizes. However, these findings do not address whether consumers have similar hedonic preferences when spending windfall money. Previous research has shown that consumer preferences can differ substantially when choosing items versus making purchases (O'Donnell & Evers, 2019). Thus, it is possible that the observed effect was due to this distinction between choice and purchase, rather than representing an instance of a windfall effect.

To investigate this possibility, we conducted a second set of analyses using data only from the control condition and the windfall condition. In this analysis, the hedonic HMPC model revealed no significant interaction between the windfall condition and hedonic perception of items ($b_{interaction}$=-0.02, $t(3861)$=-0.8, p=0.4), suggesting that the preference for hedonic items as a prize does not also appear in a windfall condition. In contrast, the MA model continued to show a significant positive main effect of similarity on purchase intention ($b_{similarity}$=0.24, $t(3695)=9.7$, $p<0.001$).

Taken together, these results clarify the role of choice versus purchase as a confounding factor in the original study. When we addressed this by replacing the prize condition with the windfall condition, the previously observed interaction between windfall and hedonic perception disappeared. This suggests a more nuanced understanding of consumer spending behavior: Consumers prefer hedonic items when selecting lottery prizes but do not necessarily prefer hedonic items when deciding how to spend windfall money. Importantly, similarity consistently influenced consumer spending decisions across both scenarios.

2.3.5 Field Evidence for Similarity-Based Mental Accounting

Next, we moved beyond hypothetical scenarios to test whether similarity influences how consumers spend windfall money in real-world, incentive-compatible purchasing decisions. We conducted a field experiment at a pop-up shop on a university campus where potential participants could purchase food, drinks, and accessories at an affordable price. Upon making an initial screening purchase, participants were randomly given a surprise coupon, labeled as either coming from one of two sources, with equal probability. They were then invited to spend their coupon on any products in the shop. As the coupon could be applied to any item in the shop, it served as a windfall. By observing participants' choices, we tested whether they were more likely to spend the coupon on items similar to its labeled source. This design enabled us to examine the effect of similarity in an incentive-compatible and ecologically valid context.

We ran a linear regression predicting the difference in the number of items purchased that were more similar to Oreo versus those more similar to stickers, using the type of coupon participants received as the predictor. If consumers are more likely to spend on items similar to the source of the money, we should observe a pattern where participants who received the Oreo coupon purchase more Oreo-like items, while those who received the sticker coupon purchase more sticker-like items (and hence fewer Oreo-like items).

Consistent with the prediction, we found a directional trend that consumers in the Oreo coupon condition bought more Oreo-like items (M=0.43), whereas consumers in the sticker coupon condition bought less Oreo-like items (M=0.23). The difference is marginally significant (b=0.20, $t(276)$=1.7, p=0.09).

The findings support the ecological validity of similarity as a factor influencing how consumers spend windfall money. In previous studies, we presented the sources of money and target items with minimal contextual cues. It is possible that, while the effect of source-target similarity on purchase intention does exist and is detectable in controlled lab settings, it has little practical influence on real purchases either because real-world consumers don't pay as much attention to the source of money or because the effect is not enough to move purchases when real money is involved. The current results therefore support the practical relevance of similarity in the real-world marketplace.

2.4 GENERAL DISCUSSIONS

The windfall effect refers to the well-documented tendency for consumers to spend windfall money more readily than hard-earned money (Arkes et al., 1994), despite the economic principle that money is fungible and should be spent similarly regardless of its source. This effect has been attributed to the idea that windfalls have a higher marginal propensity to consume (general HMPC), particularly for hedonic purchases (hedonic HMPC). A third explanation, grounded in MA, has also been proposed but in its current form lacks the precision needed to generate specific, testable predictions. The present research shows that by incorporating similarity into the MA framework, we can better account for the windfall effect. In contrast, the widely cited HMPC explanations fall short of explaining the phenomenon.

Across four controlled lab studies and one field experiment we find that consumers spend money dramatically differently based on the similarity between the source of this money and the target item. When the source of the money is more similar to the target item, consumers are more likely to purchase the target item. Crucially, this holds true regardless of whether the money is windfall or hard-earned. On the other hand, we find no evidence that consumers spend money differently based on whether it is a windfall. If anything, the effect of windfall on spending, after controlling for the effect of similarity, is slightly negative but not consistently so.

More specifically, we find that purchase intention is higher for items that are more similar to the source of the windfall, both in hypothetical scenarios (Study 1, 2, 3a, 3b) and in real consumption contexts (Study 4), supporting the MA theory for the windfall effect. By contrast, merely receiving a windfall does not increase purchase intention overall, nor does it specifically increase purchase intention for hedonic items, providing no support for either version of the HMPC theories. Furthermore, a regression model based on MA offers a better fit to the data than models based on HMPC theories (Study 2).

To reconcile our findings with prior research supporting HMPC theories, we conducted conceptual replications and extensions of two foundational studies (Arkes et al., 1994; O'Curry & Strahilevitz, 2001). We show that the evidence for general HMPC observed in prior work can be attributed to selective stimulus sampling (Study 3A): when a broader range of stimuli varying in similarity is used, the windfall effect disappears. In addition, we find that previous studies cited as evidence for hedonic HMPC did not measure purchase intentions after receiving a windfall, but rather asked participants which items they would prefer as a prize after winning a lottery. When we adjust the experimental design to measure actual purchase intentions following a windfall (while keeping all other elements constant), we find no increased spending on hedonic items (Study 3B).

Our results suggest that windfall itself does not directly increase purchase intention, unless the source of the windfall is similar to the target item under consideration. In other words, the windfall effect should be seen strictly as a phenomenon that researchers have observed under specific conditions, but not as a general theory for explaining experimental findings involving windfalls. These findings are notable in light of the popularity of the two HMPC theories, where consumer researchers often cite as explanations for windfall-related phenomena.

2.4.1 Theoretical Contribution

This research sheds light on understanding how windfalls are spent by consumers compared to hard-earned money. Prior work has argued that windfall increases spending due to its surprising nature. We show that windfall per se does not have a direct impact on purchase behavior. Instead, we propose that windfall effects are governed by the perceived similarity between the source of the money and the target item. Drawing on theories of categorization and associative memory, we argue that greater similarity increases the likelihood that the two are represented within the same mental account, thereby shaping subsequent spending decisions (Henderson & Peterson, 1992; Rosch & Mervis, 1975).

Our work also contributes to the mental accounting literature by proposing a model that is able to make falsifiable predictions. Despite being often cited to explain a wide range of consumer behavior, much is still unclear about how mental accounts are formed in consumers' minds. This has led to two theoretical issues for mental accounting to be an explanation for consumer phenomena such as windfall effect. First, this can make the mental accounting explanation tautological. For instance, researchers have argued that the windfall effect is due to windfall money being put in a "spend-now account," which increases its purchase intention (Hodge & Mason, 1995). However, the fact that people think windfall money is put in the "spend-now account" is precisely because windfall money is observed to be spent more readily in the first place. Thus, this is not a legitimate explanation for the windfall effect.

Second, the vague specification of mental accounting also makes it so flexible as to be compatible with almost all empirical observations, rendering the theory unfalsifiable. Without specifying and constraining how mental accounts can be formed prior to conducting experiments, researchers can always come up with a post hoc category and say "windfalls are spent more readily on X because windfall and X are grouped in a category," whichever item X turns out to be.

We propose a formal model of mental accounting based on the similarity between source and target items. With the model we are able to make concrete, falsifiable predictions about which items will be grouped in the same mental account by consumers.

Lastly, this research contributes to the use of NLP techniques to understand consumer behavior. Prior research in consumer psychology has demonstrated that NLP methods can be used to predict a range of consumer behaviors (e.g, Packard & Berger, 2024). This research shows that NLP-based similarity measures can also predict how consumers form mental accounts.

2.4.2 Marketing Implications

Our findings also provide industries with flexible tools to influence purchasing behavior by highlighting the perceived similarity between the source of funds and the target purchase. To encourage spending on a specific item, companies can frame the money as more closely related to that item. For example, grocery stores offering rebates can emphasize that the rebate originates from grocery purchases, increasing the likelihood that consumers will use it for future grocery shopping. Alternatively, businesses can also frame their products as naturally aligned with the source of consumers' windfalls. For instance, when consumers receive a tax refund, financial institutions can position their financial products as a way to maximize or reinvest that refund.

Similarly, policy makers and financial literacy campaigns can apply this insight to design interventions that promote consumer welfare, for instance, by framing stimulus checks or bonuses in ways that nudge consumers toward more productive spending, such as through categorical cues or suggested uses.

2.4.3 Future Directions

Our theoretical perspective would also point to new venues for future research questions by providing novel predictions that have not been tested in the windfall and mental accounting literature. If windfall spending is indeed affected by how consumers categorize funds, then theories of mental accounting and categorization would raise new questions about the categorical features of mental accounts. Does similarity have differential effects on the spending of more vs. less typical items of a category? If we prompt consumers to classify the source of money using different levels of abstraction (e.g., rebate from a restaurant vs. rebate from a fastfood restaurant), would it broaden or narrow the effect of similarity on spending?

Another potential area of future research concerns the flexibility of mental accounts. Prior studies have demonstrated that consumers can flexibly assign an item to different mental accounts to justify spending (Cheema & Soman, 2006). What are the potential boundary conditions for the malleability of mental accounts? While restaurants can be assigned to both food and entertainment accounts, it seems highly unlikely that they can be assigned to an education account, no matter how motivated consumers are in justifying their expense. It is therefore worth investigating whether the flexibility of mental accounts is constrained by the similarity between the items (e.g., restaurant) and the intended category (entertainment vs. education account).

Chapter 3

Consumer False Memory and Implications on Decision Making

This chapter extends the framework to the domain of consumer memory, examining how inaccuracies in mental representations influence judgment and decision making. In particular, it focuses on false memory as a systematic feature of cognition rather than a random error, and develops a framework to characterize when and how such errors arise. By distinguishing between different types of memory failures and their behavioral consequences, the chapter provides more precise predictions about how consumers encode, retrieve, and act on imperfect information in marketplace settings.

3.1 INTRODUCTION

Consumer memory plays a central role in consumer decision making. A large literature in marketing and psychology has examined the factors that influence the accuracy and accessibility of consumer memory, as well as the consequences of imperfect memory for brand evaluation and choice. For example, brand recall and memory accessibility influence consideration set formation and subsequent purchase decisions (Nedungadi, 1990). Similarly, the effectiveness of advertising depends heavily on how well consumers remember brand information and advertising claims (Shapiro & Krishnan, 2001).

Despite this extensive literature, two important conceptual gaps remain. First, existing studies tend to focus on situations in which consumers need to recall the presence of information. For instance, when shopping, consumers may need to remember whether a particular store carries the product they intend to buy. Much less attention has been devoted to situations in which consumers must recall the absence of information. In many real-world contexts, however, remembering the absence of information is equally important. For example, consumers may try to recall whether a brand has *not* been involved in a recent scandal or whether a particular product attribute was *not* mentioned in advertising. Failures in such recall may meaningfully influence brand evaluations and choices. Moreover, in competitive environments, advertising and brand communications can interfere with consumers' memory for competing brands, further shaping consumer evaluations (Kent & Allen, 1994).

Second, extant research primarily examines the consequences of forgetting—that is, situations in which consumers fail to recall information that was previously encountered. For example, studies have examined how failures in brand recall affect consideration sets and choice (Nedungadi, 1990). These phenomena can be viewed as false negatives in memory. However, the opposite type of error—false positives, in which consumers mistakenly believe that they have encountered information when they have not—can also have important behavioral consequences. For instance, if consumers misremember an advertisement as being associated with a competing brand, the advertisement may not only fail to benefit the original brand but may inadvertently benefit the competitor (Kent & Allen, 1994).

In this chapter, we examine these phenomena under the broader concept of consumer false memory.

3.2 MEMORY OF ABSENCE VS. PRESENCE

Previous research in memory has consistently found that people are generally better at remembering the presence of information than its absence (Johnson et al., 1993). This phenomenon, often referred to as memory asymmetry, has been documented in a variety of contexts. For example, in police lineup decisions, individuals are typically more accurate and better calibrated when identifying the presence of a suspect than when correctly rejecting a suspect who is absent (Seale-Carlisle, 2024).

While this asymmetry has been studied in eyewitness and memory research, its implications for consumer decision making remain less well understood. This is particularly notable because, unlike in police lineup tasks or many laboratory memory paradigms, consumers are typically free to choose how they retrieve information when making decisions. For instance, when deciding whether a brand avoids unethical business practices, a consumer could attempt to recall direct evidence of ethical behavior (presence) or instead try to recall the absence of negative information. In principle, consumers could adaptively select the retrieval strategy that maximizes accuracy. Such adaptive strategy selection could mitigate the effects of memory asymmetry and lead to more informed consumer choices.

However, we find that consumers do not optimally adapt their retrieval strategies. Instead, their judgments are influenced by whether salient outcomes are framed around the presence or absence of an event. Rather than focusing on confirming the presence of desirable information, consumers often attempt to recall the absence of negative information—a strategy that is particularly vulnerable to error given the inherent difficulty of remembering absences.

3.2.1 Materials and Procedure

We employed a within-subject design with 400 participants recruited from Prolific. We conducted three experiments with slightly different setups .

	Study 1 (N=400)	Study 2 (N=400)	Study 3 (N=400)
Design	Two-condition	2 × 2 factorial	Two-condition
Conditions	Presence Framing vs. Absence Framing	Presence Framing vs. Absence Framing × Presence Payoff vs. Absence Payoff	Nudge vs. No Nudge
Procedure	Encoding → Decoding	Encoding → Decoding	Encoding → Nudge → Decoding
Payoff Structure	Higher for presence	Higher for presence in Presence Framing; Higher for absence in Absence Framing	Higher for absence

Table 3.2.1.1 *Study designs in section 3.2.*

Common Encoding Procedure. In the encoding phase, participants were presented with 36 advertisements, each featuring a different well-known consumer brand (e.g., Ford, Coca-Cola) or a generic product (e.g., apple, lobster). Each advertisement consisted of the name of the brand or product and was displayed for one second. The advertisements were organized into six sections, with items from the same category grouped together (cars, sports, restaurants, food, beverages, and fruits).

To discourage participants from writing down the stimuli during the encoding phase, which could bias the measurement of memory accuracy, two measures were implemented. First, participants were instructed to press the "A" or "L" key each time they saw a brand or product name, and were informed that failure to respond correctly could result in disqualification. Second, participants were ostensibly assigned a decoy task in which they were asked to generate a descriptor for the advertisements within each section they viewed.

Decoding Procedure of Study 1. Following the encoding phase in Study 1, participants completed a series of measures assessing brand memory and willingness to pay (WTP) for a gift card associated with the items they just viewed. Memory judgments were elicited using a 100-point scale, while brand attitudes were measured on a 7-point scale. In the Presence-Framing condition, participants evaluated whether a brand had appeared previously by responding to the question, *"How likely do you think that X appeared in a list?"* In the Absence-Framing condition, participants responded to the complementary question, *"How likely do you think that X did **not** appeared in any list?"*

WTP was elicited with the following question, *"Now you are asked to indicate your willingness to pay for an envelope that contains a $20 gift card associated with X. If X appeared in any of the previous lists, it means no one has claimed the gift card, such that the envelope is worth $20. Otherwise someone already claimed the gift card, such that the envelope is worth $0.*

Appeared in a list (gift card unclaimed)	*$20*
Did not appeared (gift card already claimed)	*$0*

Given your estimated likelihood above, how much are you willing to pay for the envelope?"

Decoding Procedure of Study 2. As in Study 1, in Study 2 participants answered a series of measures assessing brand memory of WTP for a gift card associated with the items they just viewed after encoding. The memoery questions were the same as those in Study 1. The WTP questions differed in terms of the value of the envelops across Payoff conditions., In the Presence-Payoff condition, the envelop had a higher value when participants had viewed the item in the encoding phase, which was the same as in Study 1. In the Absence-Payoff condition, the envelop had a higher value when the participants had not viewed the item in the encoding phase.

Nudge Procedure of Study 3. To manipulate whether participants retrieved the presence or absence of events, we introduced a nudge manipulation between the encoding and retrieval phases in Study 3. Memory questions were framed exclusively in terms of presence, resulting in a simple two-condition design (Nudge vs. No Nudge). Note that in Study 3, the presence of item was associated with lower value. More specifically, WTP was elicited using the following question, *"Now you are asked to indicate your willingness to pay for an envelope that contains a $20 gift card associated with X. If X appeared in any of the previous lists, it means someone already claimed the gift card, such that the envelope is worth $0. Otherwise no one has claimed the gift card, such that the envelope is worth $20.*

Appeared in a list (gift card already claimed) $0

Did not appeared (gift card unclaimed) $20

Given your estimated likelihood above, how much are you willing to pay for the envelope?"

In the Nudge condition, participants read the following instruction nudging them to retrieve the presence of memory: *"Next, you will now answer some questions about the Lists you saw at the beginning based on your memory of the items. Research on human memory has shown that people are generally better at recalling events that did happen rather than identifying things that did not happen. This means that when judging whether something did not occur, it is often more reliable to first ask yourself whether you can remember it happening. If you cannot recall it, you can reasonably conclude that it likely did not. Using this strategy may help you make more accurate judgments and better decisions."*

In the No Nudge condition, participants read the following neutral piece of information: *"Research on human memory has shown that people use different strategies to judge whether events occurred or not, and memory accuracy can vary depending on the context. As you complete the task, be aware that recalling past events is not always perfect, and different situations may influence how easily something comes to mind."*

Immediately after the instruction, participants took an attention check about the instruction they just read. Only participants who passed the quiz were included in the anlaysis.

Measurement. Memory accuracy was quantified using signal detection theory. Specifically, we computed d-prime (d′) as the difference between the z-scores of hit rates and false alarm rates. Higher d-prime represents a higher memory efficiency. Responses to the memory questions were binarized using a threshold of 50 (on the 0–100 scale), with values above 50 classified as affirmative responses.

Decision quality was measured as the expected payoff participants would receive if their WTP responses were implemented in a Becker–DeGroot–Marschak (BDM) auction. The value of each gift card depended on whether the item had appeared previously, with higher value assigned to items that had been shown before. Thus, participants who made higher-quality WTP decisions would earn higher expected payoffs.

3.2.2 Study 1: Replicating Memory Asymmetry

Consistent with prior findings in the false memory literature, we hypothesized that consumers would exhibit higher memory accuracy when retrieving the presence of events than when retrieving their absence.

H1: Participants in the Presence condition will exhibit higher memory accuracy (d′) than participants in the Absence condition.

Because high-quality WTP decisions depend on accurate memory, this asymmetry should translate into differences in decision quality. Accordingly, we hypothesized that consumers would make better WTP decisions when retrieving the presence of events than when retrieving their absence.

H2.1: Participants in the Presence condition will make WTP decisions that yield higher expected payoffs than those in the Absence condition.

H2.2: The effect of framing on expected payoffs will be mediated by differences in memory accuracy.

Results 1. Consistent with H1, participants in the Presence condition exhibited higher memory accuracy than those in the Absence condition. Although participants in the Presence condition showed a slightly lower hit rate (Δhit rate = -0.05, $p = 0.02$; $M_p = 0.69$, $M_a = 0.74$), they exhibited a substantially lower false alarm rate (Δfalse alarm rate = -0.22, $p < 0.001$; $M_p = 0.23$, $M_a = 0.45$). This reduction in false alarms led to significantly higher overall sensitivity ($\Delta d' = 0.54$, $p < 0.001$; $t = 6.895$).

These results support H1 and replicate the well-documented asymmetry in memory for the presence versus absence of information.

Results 2. We next examine whether this memory asymmetry translates into differences in consumer decision quality. Consistent with H2.1, participants in the Presence condition made better WTP decisions than those in the Absence condition. Based on expected payoff calculations from the Becker–BDM auction framework, participants in the Presence condition earned higher payoffs (Δpayoff = 6.74, $p = 0.001$; $t = 5.848$, $M_p = 17.5$, $M_a = 10.7$).

These findings suggest that differences in memory performance can have downstream consequences for consumer decision quality.

A causal mediation analysis using nonparametric bootstrap confidence intervals revealed a significant indirect effect of framing on decision quality through memory accuracy (Average Causal Mediation Effect = 0.81, 95% CI [0.53, 1.14], $p < 0.001$). The direct effect remained significant (Average Direct Effect = 5.92, 95% CI [2.86, 8.97], $p < 0.001$), resulting in a significant total effect (Total Effect = 6.74, 95% CI [3.60, 9.85], $p < 0.001$).

These findings support H2.2 and indicate that memory accuracy partially mediates the effect of presence versus absence framing on decision quality. In other words, poorer decision performance in the Absence condition is driven, at least in part, by reduced memory sensitivity.

3.2.3 Study 2: Manipulate Salient Outcome

In Study 1, we found that, consistent with prior literature, consumers exhibited higher memory accuracy when prompted to retrieve the presence of an event. In that study, higher payoff was associated with the presence of the item during encoding.

Could consumers' memory accuracy differ depending on whether the higher payoff is associated with the presence or absence of an event? On the one hand, because retrieving the presence of an event is more efficient, consumers should exhibit higher memory performance when prompted with presence framing, regardless of the payoff structure. On the other hand, when higher payoff is associated with absence, absence-framed memory questions may be more fluent than presence-framed questions, as the absence of the event is more salient in both cases. If so, consumers may exhibit higher memory accuracy when the memory framing is congruent with the payoff structure.

Formally, we test the following hypotheses:

H3.1: Participants will exhibit higher memory accuracy, and thus earn higher expected payoffs, when memory questions are framed in terms of presence (vs. absence), regardless of the payoff structure.

H3.2: Participants will exhibit higher memory accuracy, and thus earn higher expected payoffs, when the framing of the memory question is congruent with the payoff structure.

Results 3. A 2 × 2 ANOVA on memory accuracy (d′) revealed a significant overall effect ($F(3, 383) = 15.60$, $p < 0.001$). Decomposing this effect, there was a significant main effect of memory framing ($b = 0.525$, $SE=0.110$, $t=4.78$, $p<0.001$), such that participants exhibited higher memory accuracy when memory questions were framed in terms of presence (vs. absence) of events. This supports the memory part of H3.1.

Surprisingly, the main effect of payoff structure was not significant (b=−0.026, SE=0.110, t=−0.24, p=0.813). The interaction between memory framing and payoff structure was also not significant (b=0.021, SE=0.157, t=0.14, p=0.893). These do not support the payoff part of H3.1 and H3.2.

Taken together, these results reveal a nuanced picture of the mechanism through which consumer memory of presence vs. absence can differ. First, we find that difference in memory performance is not driven by congruence of memory retrieval strategy and salient outcome. Second, we find that better memory performance does not necessarily translates to higher-quality decisions, even in cases where better memory performance logically lead to higher-quality decisions.

3.2.4 Study 3: Memory Retrieval Nudge

Study 1 and 2 suggest that retrieval strategies can influence the quality of consumer decisions through their impact on memory accuracy. In turn, this reasoning implies that interventions that nudge consumers to adopt a particular retrieval strategy should improve decision quality. However, as we demonstrate below, this seemingly straightforward implication is not guaranteed to hold.

H4: Participants nudged to retrieve the presence of events will make WTP decisions that yield higher expected payoffs than those in the control condition.

To test H4, we employed a two-condition nudging manipulation in between memory encoding and memory retrieval. In the Nudge condition, participants were reminded that retrieving the presence of events could lead to better memory accuracy. In the No Nudge condition, participants read a netural piece about memory. In addition, a gift card was associated with higher value if it did *not* appear in the encoding phase before, opposite to the payoff of Study 1. In other words, the salient outcome was associated with the absence of events.

Results 3. We did not find support for nudging on consumers' memory performance. Specifically, memory accuracy was no significantly difference between Nudge and No Nudge condition ($\Delta d'$=-0.01, p=0.417, M_{n}=1.450, M_{nn}=1.464). Furthermore, there was no significant difference between quality of WTP decisions (Δpayoff=−1.1, p=0.43, M_{n}=9.2, M_{nn}=10.3).

3.2.5 Discussion

In the context of remembering exposures to brands/products, the presence of memory can lead to better decision quality compared to the absence of memory. This effect is driven by better memory retrieval in the Present-Framing condition than that in the Absence-Framing condition. However, interventions that can improve memory retrieval in the Absence condition failed to improve decision quality because consumers do not actually rely on this improved memory to make their decision.

These findings have important implications for understanding consumer behavior. They suggest that memory asymmetries do not merely influence what people can recall but actively shape how they approach decision-making. Moreover, they highlight a critical gap in consumers' adaptive reasoning: even when given the freedom to select their approach, individuals do not naturally compensate for the known weaknesses in memory for absence.

Understanding this behavioral pattern opens new avenues for interventions aimed at improving consumer choices. For example, strategies that prompt consumers to focus on recalling the presence of positive information or that provide external cues to support memory retrieval could help mitigate the negative effects of this asymmetry. Ultimately, our findings suggest that addressing the cognitive biases inherent in memory could be an essential step toward fostering more informed and confident consumer decisions.

3.3 CONSUMER FALSE MEMORY

Advertising plays an important role in shaping consumer judgments and consumption behavior. In the United States alone, firms spent $138.7 billion on advertising in 2025, reflecting the central role of advertising in modern marketplaces (*Search Ad Spend in the U.S. 2025*, Statista). High-profile advertising events such as the Super Bowl and the Academy Awards further illustrate the scale and visibility of firms' efforts to influence consumers.

A key prerequisite for advertising effectiveness is that consumers remember the advertisement. Indeed, many theories of advertising assume that exposure to an ad influences consumer attitudes through memory-based processes (Keller, 2003; Vakratsas & Ambler, 1999a). However, a large body of research in psychology and consumer behavior shows that consumers often have imperfect and unreliable memories. People may forget having seen an advertisement or even falsely remember having seen one that was never presented (Braun, 1999b; Braun et al., 2002; Schacter, 1999).

In this section, we examine how such memory errors influence consumer attitudes toward brands and their competitors. Specifically, we investigate whether the effects of advertising depend on consumers' memory of the advertisement or whether attitudes can also be influenced by false memories of exposure.

Consumer research has largely focused on the consequences of *memory failure*, such as when consumers are unable to recall having seen an advertisement or fail to remember the advertised brand. For example, a consumer may view an advertisement for Toyota but later fail to recall either the exposure or the brand itself. Existing work therefore emphasizes the role of *memory omission* in attenuating advertising effectiveness.

However, memory failure can take a more consequential form. Beyond forgetting, consumers may *misremember* the source of information. For instance, a consumer who viewed an advertisement highlighting Toyota's positive features may later misattribute the advertisement to Honda. In such cases, competing brands may effectively "free ride" on another brand's advertising due to false memory.

This possibility is not fully addressed by existing theories. Standard models of advertising and associative memory implicitly assume that persuasive effects accrue to the correctly encoded brand, while memory research has primarily examined accuracy rather than downstream marketplace consequences. As a result, current frameworks offer limited predictions regarding how false memory may systematically redistribute the effects of marketing communications across brands.

To address this gap, we examine the consequences of consumer false memory and test the following hypotheses:

H1: If consumers fail to recall seeing a brand's advertisement, their attitude toward the brand will remain unchanged relative to a no-advertisement condition.

H2: If consumers falsely remember seeing a brand's advertisement, their attitude toward the brand will increase relative to consumers who correctly remember not seeing the advertisement.

H3: False memories of advertisements will have effects on brand attitudes comparable to those of correctly remembered advertisements.

3.3.1 Materials and Procedure

The main study employed a within-subject experimental design with N=400 participants. The materials consisted of 60 brand names drawn from six common consumer product categories: car, computer, detergent, shampoo, shoes, and soft drink. For each category, ten brands were selected, resulting in a total of sixty brands. The selected brands included both widely known national brands and less prominent brands in order to generate variation in baseline familiarity and attitudes.

In the second stage, participants evaluated the same set of 60 brands. For each brand, participants first indicated their memory judgment by answering the question "How likely do you think this brand appeared earlier?" on a 100-point scale. Participants then reported their attitude toward the brand by answering the question "How much do you like this brand?" on a 7-point scale.

Half of the brands had been presented during the first stage (Ad condition), while the remaining half had not (No-Ad condition). Because each participant evaluated both types of brands, the study employed a two-condition within-subject design.

3.3.2 Results

Descriptive statistics show that consumers exhibit substantial memory errors when recalling advertisements. Among brands that were actually presented during the exposure stage, only 64% of participants correctly recalled having seen the brand. Conversely, among brands that were not presented, 51% of participants falsely believed that they had seen the brand earlier (Figure 3.3.2.1). These results indicate that consumers' memories of advertising exposure are often inaccurate, with both forgetting and false recognition occurring at meaningful rates.

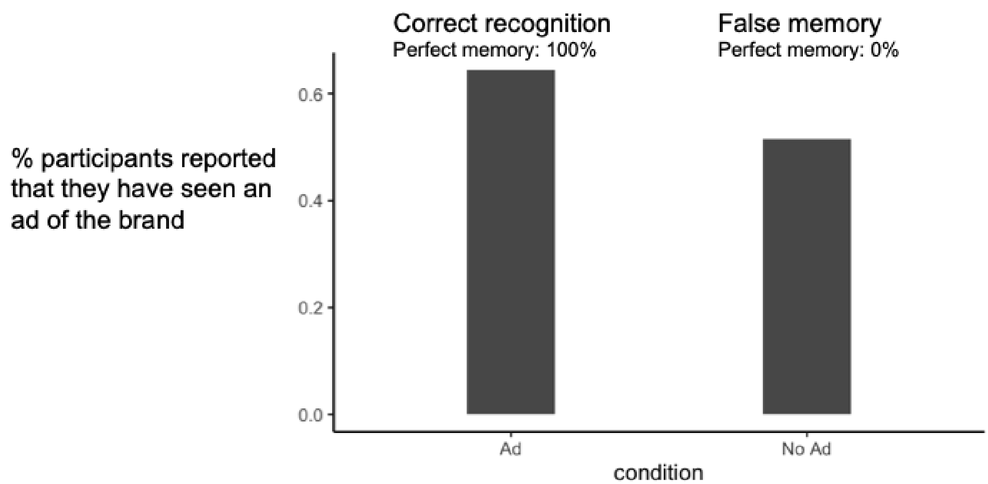

Figure 3.3.2.1 *Participants in the Ad condition are more likely to report that they have seen an ad than those in the No-Ad condition.*

Advertising works as expected when consumers have correct memory of the advertisement. Consistent with prior advertising research, ad exposure increases consumers' attitudes toward the brand when participants correctly remember having seen the advertisement (Figure 3.3.2.2).

However, advertising has little effect when consumers do not remember the advertisement. When consumers fail to recall seeing a brand's ad, their attitudes toward the brand are slightly lower to those in the no-ad condition (M_{ad}=2.3, $M_{no\ ad}$=2.4, t=13.16, p<0.001). This finding suggests that advertising exposure alone is not sufficient to influence consumer attitudes; rather, the effectiveness of advertising depends critically on whether the advertisement is encoded and retrieved from memory.

A more surprising pattern emerges when consumers falsely believe they have seen an advertisement. When participants incorrectly recall having seen a brand's ad, their attitudes toward the brand are almost identical to those observed when the advertisement was actually presented (M_{ad}=3.1, $M_{no\ ad}$=3.2, t=1.59, p=0.112). In other words, false memories of advertising produce effects similar to true exposure. This implies that brands can benefit from advertising "free-riding," where consumers attribute advertising exposure to a brand even though the brand did not actually advertise.

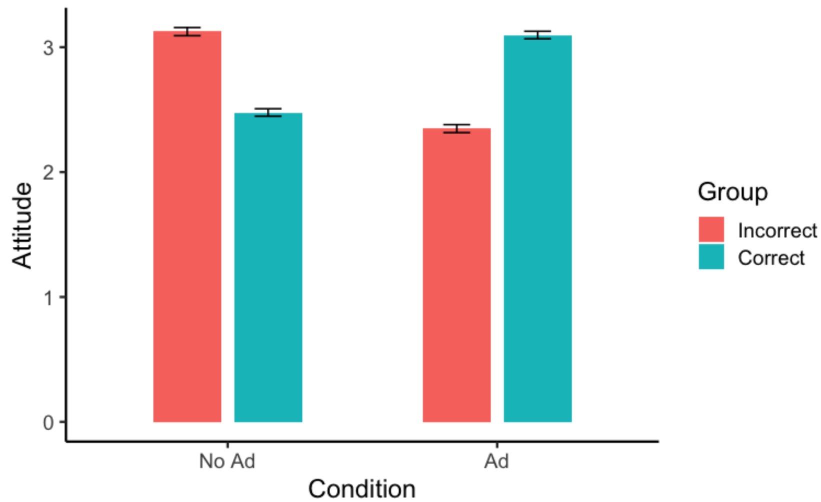

Figure 3.3.2.2 *Consumers who falsely recall seeing an advertisement for a brand exhibit the same increase in brand attitude as consumers who correctly recall seeing the advertisement. Similarly, consumers who falsely recall not seeing an advertisement show the same decrease in brand attitude as consumers who correctly reject having seen the advertisement.*

3.3.3 Discussion

The results in section 3.3 reveal an important strategic implication: false consumer memories can generate unintended spillover benefits for competing brands. In real-world markets where consumers are exposed to a large number of advertisements, some brands may reap the benefits of expensive advertising campaigns without actually paying for them if consumers mistakenly attribute the advertisement to the wrong brand.

The findings fill an important gap in the literature on consumer memory for brand communications. Extant research has primarily focused on the effects of advertising on focal (target) brands, typically evaluating effectiveness based on the extent to which advertising enhances consumers' memory for those brands (Nedungadi, 1990; Shapiro & Krishnan, 2001). In contrast, the present research identifies *advertising free-riding* as an additional and previously overlooked consequence of advertising exposure. Specifically, we show that when consumers misattribute an advertisement to a non-featured brand, their attitudes toward that brand improve as if they had actually seen the advertisement for it.

This perspective suggests that firms should focus not only on designing advertisements that are persuasive but also on ensuring that advertisements are memorable and clearly associated with the focal brand. In addition, firms may wish to design advertisements that minimize the possibility of advertising free-riding, for example by strengthening distinctive brand cues.

Future research could address several important gaps that remain unresolved in the current literature. First, little is known about the *long-term consequences* of advertising free-riding for brand competition and advertising efficiency. Existing work largely evaluates advertising based on its immediate impact on the focal brand (Shapiro & Krishnan, 2001; Vakratsas & Ambler, 1999b), but if advertising benefits can spill over to competing brands through memory errors, this may systematically distort competitive dynamics and lead firms to misestimate the returns to advertising investments. Understanding these longer-term effects is therefore critical for both theory and managerial decision-making.

Second, there is limited understanding of how to *mitigate memory-based free-riding*. While prior research has examined how to enhance memory for advertisements (Ambler & Burne, 1999; Keller, 1987; Unnava & Burnkrant, 1991b, 1991a), it has not addressed how to prevent misattribution of those memories. Identifying effective interventions, such as optimizing repetition strategies (e.g., spaced vs. massed exposures), using interactive formats that strengthen encoding, or targeting consumer segments that are less susceptible to memory errors, would help firms design advertising that more reliably benefits the intended brand.

Conclusion

This dissertation examines how the structure of consumer cognition shapes judgment and decision-making. Across three projects, I investigate how mental representations are constructed, updated, and applied in dynamic environments. In the first project, I show that perceptual attributes can shift in meaning over time. I develop an algorithm that identifies attributes exhibiting significant semantic change and adjusts for these shifts, enabling more accurate interpretation of longitudinal brand–attribute associations. In the second project, I demonstrate that similarity between the source of money and the product under consideration influences consumers' purchase intentions. Specifically, greater perceived similarity increases the likelihood of spending, providing a structured and testable account of mental accounting behavior. In the third project, I show that consumers do not adopt the optimal memory strategy when retrieving the absence of an event. Rather than consistently relying on the more efficient strategy of retrieving the presence of related events, consumers may attempt to directly retrieve absence when it is outcome-relevant, leading to systematic decreases in memory accuracy.

Together, these findings advance the cognitive perspective in consumer research by formalizing previously loosely specified theories—such as those concerning brand meaning and mental accounting—and deriving precise, falsifiable predictions. More broadly, this dissertation provides a unified framework for understanding how consumers construct, update, and apply mental representations in decision-making.

*This dissertation contains co-authored materials with Professor Ming Hsu, Professor Ellen Evers, and Professor Zhihao Zhang.

References

Aaker, D. A. (1991). *Managing brand equity: Capitalizing on the value of a brand name*. Free Press ; Maxwell Macmillan Canada ; Maxwell Macmillan International.

Aaker, D. A. (2011). *Building strong brands*. Simon and schuster.

Aaker, J. L. (1997a). Dimensions of brand personality. *Journal of Marketing Research*, *34*(3), 347–356.

Aaker, J. L. (1997b). Dimensions of Brand Personality. *JOURNAL OF MARKETING RESEARCH*.

Aaker, J. L., Benet-Martinez, V., & Garolera, J. (2001). Consumption symbols as carriers of culture: A study of Japanese and Spanish brand personality constucts. *Journal of Personality and Social Psychology*, *81*(3), 492.

Abeler, J., & Marklein, F. (2017). Fungibility, Labels, and Consumption. *Journal of the European Economic Association*, *15*(1), 99–127. https://doi.org/10.1093/jeea/jvw007

Aggarwal, P., & McGill, A. L. (2012). When brands seem human, do humans act like brands? Automatic behavioral priming effects of brand anthropomorphism. *Journal of Consumer Research*, *39*(2), 307–323.

Agres, S. J., & Dubitsky, T. M. (1996). Changing needs for brands. *Journal of Advertising Research*, *36*(1), 21–30.

Alba, J. W., & Hutchinson, J. W. (1987a). Dimensions of Consumer Expertise. *Journal of Consumer Research*, *13*(4), 411–454. https://doi.org/10.1086/209080

Alba, J. W., & Hutchinson, J. W. (1987b). Dimensions of Consumer Expertise. *Journal of Consumer Research*, *13*(4), 411–454.

Ambler, T., & Burne, T. (1999). The Impact of Affect on Memory of Advertising. *Journal of Advertising Research*, *39*(2), 25–34. https://doi.org/10.1080/00218499.1999.12466598

Arkes, H. R., Joyner, C. A., Pezzo, M. V., Nash, J. G., Siegel-Jacobs, K., & Stone, E. (1994). The Psychology of Windfall Gains. *Organizational Behavior and Human Decision Processes*, *59*(3), 331–347. https://doi.org/10.1006/obhd.1994.1063

Bamler, R., & Mandt, S. (2017). Dynamic Word Embeddings. *Proceedings of the 34th International Conference on Machine Learning*, 380–389. https://proceedings.mlr.press/v70/bamler17a.html

Bass, F. M. (1969). A New Product Growth for Model Consumer Durables. *Management Science*, *15*(5), 215–227. https://doi.org/10.1287/mnsc.15.5.215

Berger, J., Humphreys, A., Ludwig, S., Moe, W. W., Netzer, O., & Schweidel, D. A. (2020). Uniting the Tribes: Using Text for Marketing Insight. *Journal of Marketing*, *84*(1), 1–25. https://doi.org/10.1177/0022242919873106

Berger, J., Kim, Y. D., & Meyer, R. (2021). What Makes Content Engaging? How Emotional Dynamics Shape Success. *Journal of Consumer Research*, *48*(2), 235–250. https://doi.org/10.1093/jcr/ucab010

Berger, J., & Milkman, K. L. (2012). What Makes Online Content Viral? *Journal of Marketing Research*, *49*(2), 192–205. https://doi.org/10.1509/jmr.10.0353

Berger, J., Moe, W. W., & Schweidel, D. A. (2023). What Holds Attention? Linguistic Drivers of Engagement. *Journal of Marketing*, *87*(5), 793–809. https://doi.org/10.1177/00222429231152880

Berger, J., Packard, G., Boghrati, R., Hsu, M., Humphreys, A., Luangrath, A., Moore, S., Nave, G., Olivola, C., & Rocklage, M. (2022). Marketing insights from text analysis. *Marketing Letters*, 1–13. https://doi.org/10.1007/s11002-022-09635-6

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice processes. *Journal of Consumer Research*, *25*(3), 187–217. https://doi.org/10.1086/209535

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, *100*(5), 992–1026. https://doi.org/10.1086/261849

Boivin, Y. (1986). A free response approach to the measurement of brand perceptions. *International Journal of Research in Marketing*, *3*(1), 11–17.

Brakus, J. J., Schmitt, B. H., & Zarantonello, L. (2009). Brand experience: What is it? How is it measured? Does it affect loyalty? *Journal of Marketing*, *73*(3), 52–68.

Braun, K. A. (1999a). Postexperience Advertising Effects on Consumer Memory. *Journal of Consumer Research*, *25*(4), 319–334. https://doi.org/10.1086/209542

Braun, K. A. (1999b). Postexperience Advertising Effects on Consumer Memory: Journal of Consumer Research. *Journal of Consumer Research*, *25*(4), 319–334. https://doi.org/10.1086/209542

Braun, K. A., Ellis, R., & Loftus, E. F. (2002). Make my memory: How advertising can change our memories of the past. *Psychology & Marketing*, *19*(1), 1–23. https://doi.org/10.1002/mar.1000

Burke, R. R., & Srull, T. K. (1988). Competitive Interference and Consumer Memory for Advertising. *Journal of Consumer Research*, *15*(1), 55–68. https://doi.org/10.1086/209145

Burnham, K. P., & Anderson, D. R. (2004). Multimodel Inference: Understanding AIC and BIC in Model Selection. *Sociological Methods & Research*, *33*(2), 261–304. https://doi.org/10.1177/0049124104268644

Büschken, J., & Allenby, G. M. (2016). Sentence-Based Text Analysis for Customer Reviews. *Marketing Science*, *35*(6), 953–975. https://doi.org/10.1287/mksc.2016.0993

Carlson, J. R., Ross, W. T., Coulter, R. A., & Marquardt, A. J. (2019). About time in marketing: An assessment of the study of time and conceptual framework. *AMS Review*, *9*(3–4), 136–154. https://doi.org/10.1007/s13162-019-00148-6

Cascio Rizzo, G. L., Berger, J., De Angelis, M., & Pozharliev, R. (2023). How Sensory Language Shapes Influencer's Impact. *Journal of Consumer Research*, *50*(4), 810–825. https://doi.org/10.1093/jcr/ucad017

Cheema, A., & Soman, D. (2006). Malleable Mental Accounting: The Effect of Flexibility on the Justification of Attractive Spending and Consumption Decisions. *Journal of Consumer Psychology*, *16*(1), 33–44.

Coleman, N. V., & Williams, P. (2013). Feeling like my self: Emotion profiles and social identity. *Journal of Consumer Research*, *40*(2), 203–222. https://doi.org/10.1086/669483

Culotta, A., & Cutler, J. (2016). Mining Brand Perceptions from Twitter Social Networks. *Marketing Science*, *35*. https://doi.org/10.1287/mksc.2015.0968

De Chernatony, L. (2010). *Creating powerful brands*. Routledge.

Deo, A. (2015). *Diachronic semantics*. *1*(1), 179–197.

Dick, A., Chakravarti, D., & Biehal, G. (1990). Memory-Based Inferences During Consumer Choice. *Journal of Consumer Research*, *17*(1), 82–93. https://doi.org/10.1086/208539

Diewert, W. E. (1988). *The Early History of Price Index Research* (No. W2713; p. w2713). National Bureau of Economic Research. https://doi.org/10.3386/w2713

Easley, R. W., Madden, C. S., & Dunn, M. G. (2000). Conducting Marketing Science. *Journal of Business Research*, *48*(1), 83–92. https://doi.org/10.1016/S0148-2963(98)00079-4

Fournier, S. (1998). Consumers and Their Brands: Developing Relationship Theory in Consumer Research. *Journal of Consumer Research*, *24*(4), 343–373. https://doi.org/10.1086/209515

Gabel, S., Guhl, D., & Klapper, D. (2019). P2V-MAP: Mapping Market Structures for Large Retail Assortments. *Journal of Marketing Research*, *56*(4), 557–580. https://doi.org/10.1177/0022243719833631

Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(16), E3635–E3644. https://doi.org/10.1073/pnas.1720347115

Geeraerts, D., Gevaert, C., & Speelman, D. (2012). How anger rose: Hypothesis testing in diachronic semantics. In *Current methods in historical semantics* (pp. 109–132).

Gensch, D. H., & Golob, T. F. (1975). Testing the Consistency of Attribute Meaning in Empirical Concept Testing. *Journal of Marketing Research*, *12*(3), 348–354. https://doi.org/10.1177/002224377501200312

Gillioz, A., Casas, J., Mugellini, E., & Khaled, O. A. (2020). Overview of the Transformer-based Models for NLP Tasks. *2020 15th Conference on Computer Science and Information Systems (FedCSIS)*, 179–183. https://doi.org/10.15439/2020F20

Goldstone, R. L. (1994). The role of similarity in categorization: Providing a groundwork. *Cognition*, *52*(2), 125–157. https://doi.org/10.1016/0010-0277(94)90065-5

Grand, G., Blank, I. A., Pereira, F., & Fedorenko, E. (2022). Semantic projection recovers rich human knowledge of multiple object features from word embeddings. *Nature Human Behaviour*, *6*(7), 975–987. https://doi.org/10.1038/s41562-022-01316-8

Gupta, T., & Hagtvedt, H. (2021). Safe together, vulnerable apart: How interstitial space in text logos impacts brand attitudes in tight versus loose cultures. *Journal of Consumer Research*, *48*(3), 474–491.

Hamilton, Leskovec, J., & Jurafsky, D. (2016a). Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change. *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, *2016*, 2116–2121. https://doi.org/10.18653/v1/d16-1229

Hamilton, Leskovec, J., & Jurafsky, D. (2016b). Diachronic word embeddings reveal statistical laws of semantic change. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers*, 1489-1501,.

Hastie, T., Tibshirani, R., & Friedman, J. H. (2017). *The elements of statistical learning: Data mining, inference, and prediction* (Second edition, corrected at 12th printing 2017). Springer. https://doi.org/10.1007/b94608

Hauser, J. R., & Koppelman, F. S. (1979). Alternative perceptual mapping techniques: Relative accuracy and usefulness. *Journal of Marketing Research*, *16*(4), 495–506.

Hayes-Roth, B. (1977). Evolution of cognitive structures and processes. *Psychological Review*, *84*(3), 260–278. https://doi.org/10.1037/0033-295X.84.3.260

Heilman, C. M., Nakamoto, K., & Rao, A. G. (2002). Pleasant Surprises: Consumer Response to Unexpected In-Store Coupons. *Journal of Marketing Research*, *39*(2), 242–252. https://doi.org/10.1509/jmkr.39.2.242.19081

Helion, C., & Gilovich, T. (2014). Gift Cards and Mental Accounting: Green-lighting Hedonic Spending. *Journal of Behavioral Decision Making*, *27*(4), 386–393. https://doi.org/10.1002/bdm.1813

Henderson, P. W., & Peterson, R. A. (1992). Mental accounting and categorization. *Organizational Behavior and Human Decision Processes*, *51*(1), 92–117. https://doi.org/10.1016/0749-5978(92)90006-S

Hodge, S. K., & Mason, C. H. (1995). Work versus windfall: An exploration of saving on subsequent purchase. *Marketing Letters*, *6*(2), 91–100. https://doi.org/10.1007/BF00994925

Hovy, D., Melumad, S., & Inman, J. J. (2021). Wordify: A Tool for Discovering and Differentiating Consumer Vocabularies. *Journal of Consumer Research*, *48*(3), 394–414. https://doi.org/10.1093/jcr/ucab018

Humphreys, A., & Wang, R. J.-H. (2018). Automated Text Analysis for Consumer Research. *Journal of Consumer Research*, *44*(6), 1274–1306. https://doi.org/10.1093/jcr/ucx104

Janiszewski, C., Noel, H., & Sawyer, A. G. (2003). A Meta-analysis of the Spacing Effect in Verbal Learning: Implications for Research on Advertising Repetition and Consumer Memory. *Journal of Consumer Research*, *30*(1), 138–149. https://doi.org/10.1086/374692

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, *114*(1), 3–28. https://doi.org/10.1037/0033-2909.114.1.3

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, *58*(9), 697–720. https://doi.org/10.1037/0003-066X.58.9.697

Kauppinen-Räisänen, H., Gummerus, J., Koskull, C., & Cristini, H. (2019). The new wave of luxury: The meaning and value of luxury to the contemporary consumer. *Qualitative Market Research: An International Journal*, *22*(3), 229–249.

Keller, K. L. (1987). Memory Factors in Advertising: The Effect of Advertising Retrieval Cues on Brand Evaluations. *Journal of Consumer Research*, *14*(3), 316. https://doi.org/10.1086/209116

Keller, K. L. (1993). Conceptualizing, Measuring, and Managing Customer-Based Brand Equity. *Journal of Marketing*, *57*(1), 1–22. https://doi.org/10.1177/002224299305700101

Keller, K. L. (2003). Brand Synthesis: The Multidimensionality of Brand Knowledge. *Journal of Consumer Research*, *29*(4), 595–600. https://doi.org/10.1086/346254

Keller, K. L. (2016). Reflections on customer-based brand equity: Perspectives, progress, and priorities. *AMS Review*, *6*(1–2), 1–16. https://doi.org/10.1007/s13162-016-0078-z

Kent, R. J., & Allen, C. T. (1994). Competitive Interference Effects in Consumer Memory for Advertising: The Role of Brand Familiarity. *Journal of Marketing*, *58*(3), 97–105. https://doi.org/10.1177/002224299405800307

Kim, S., Lee, S., & McCulloch, R. (2024). A Topic-Based Segmentation Model for Identifying Segment-Level Drivers of Star Ratings from Unstructured Text Reviews. *Journal of Marketing Research*, 00222437241246752. https://doi.org/10.1177/00222437241246752

Kline, P. (2015). *A handbook of test construction (psychology revivals): Introduction to psychometric design*. Routledge.

Krishnan, H. S. (1996). Characteristics of memory associations: A consumer-based brand equity perspective. *International Journal of Research in Marketing*, *13*(4), 389–405.

Kronrod, A., Gordeliy, I., & Lee, J. K. (2023). Been There, Done That: How Episodic and Semantic Memory Affects the Language of Authentic and Fictitious Reviews. *Journal of Consumer Research*, *50*(2), 405–425. https://doi.org/10.1093/jcr/ucac056

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*(1), 22–44. https://doi.org/10.1037/0033-295X.99.1.22

Kulkarni, V., Al-Rfou, R., Perozzi, B., & Skiena, S. (2015). Statistically significant detection of linguistic change. *Proceedings of the 24th International Conference on World Wide Web*, 625–635.

Kutuzov, A., Øvrelid, L., Szymanski, T., & Velldal, E. (2018). Diachronic word embeddings and semantic shifts: A survey. *Proceedings of the 27th International Conference on Computational Linguistics*, 1384-1397,.

Kutuzov, A., Velldal, E., & Øvrelid, L. (2017). *Tracing armed conflicts with diachronic word embedding models*. Association for Computational Linguistics.

Lee, J. K. (2021). Emotional Expressions and Brand Status. *Journal of Marketing Research*, *58*(6), 1178–1196. https://doi.org/10.1177/00222437211037340

Lee, T. Y., & Bradlow, E. T. (2011). Automated Marketing Research Using Online Customer Reviews. *Journal of Marketing Research*, *48*(5), 881–894. https://doi.org/10.1509/jmkr.48.5.881

Leigh, T. W., Peters, C., & Shelton, J. (2006). The consumer quest for authenticity: The multiplicity of meanings within the MG subculture of consumption. *Journal of the Academy of Marketing Science*, *34*(4), 481–493.

Levav, J., & Mcgraw, A. P. (2009). Emotional Accounting: How Feelings about Money Influence Consumer Choice. *Journal of Marketing Research*, *46*(1), 66–80. https://doi.org/10.1509/jmkr.46.1.66

Liu, J., & Toubia, O. (2018). A Semantic Approach for Estimating Consumer Content Preferences from Online Search Queries. *Marketing Science*, *37*(6), 930–952. https://doi.org/10.1287/mksc.2018.1112

Liu, L., Dzyabura, D., & Mizik, N. (2020). Visual Listening In: Extracting Brand Image Portrayed on Social Media. *Marketing Science*, *39*. https://doi.org/10.1287/mksc.2020.1226

Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory*, *12*(4), 361–366. https://doi.org/10.1101/lm.94705

Lynch, J. G., Bradlow, E. T., Huber, J. C., & Lehmann, D. R. (2015). Reflections on the replication corner: In praise of conceptual replications. *International Journal of Research in Marketing*, *32*(4), 333–342. https://doi.org/10.1016/j.ijresmar.2015.09.006

Lynch, J. G., Jr., & Srull, T. K. (1982). Memory and Attentional Factors in Consumer Choice: Concepts and Research Methods. *Journal of Consumer Research*, *9*(1), 18–37. https://doi.org/10.1086/208893

Maaten, L. van der, & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, *9*(86), 2579–2605.

Mathur, P., Jain, S. P., & Maheswaran, D. (2012). Consumers' implicit theories about personality influence their brand personality judgments. *Journal of Consumer Psychology*, *22*(4), 545–557.

Mihalcea, R., & Nastase, V. (2012). Word epoch disambiguation: Finding how words change over time. *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers*, 259–263.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. *Advances in Neural Information Processing Systems*, *26*. https://papers.nips.cc/paper_files/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html

Milkman, K. L., & Beshears, J. (2009). Mental accounting and small windfalls: Evidence from an online grocer. *Journal of Economic Behavior & Organization*, *71*(2), 384–394. https://doi.org/10.1016/j.jebo.2009.04.007

Mittal, V., Kumar, P., & Tsiros, M. (1999). Attribute-level performance, satisfaction, and behavioral intentions over time: A consumption-system approach. *Journal of Marketing*, *63*(2), 88–101.

Mizik, N., & Jacobson, R. (2008). The financial value impact of perceptual brand attributes. *Journal of Marketing Research*, *45*(1), 15–32.

Nedungadi, P. (1990). Recall and Consumer Consideration Sets: Influencing Choice without Altering Brand Evaluations. *Journal of Consumer Research*, *17*(3), 263–276. https://doi.org/10.1086/208556

Netzer, O., Feldman, R., Goldenberg, J., & Fresko, M. (2011). Mine Your Own Business: Market-Structure Surveillance Through Text Mining. *Marketing Science*, *31*. https://doi.org/10.2139/ssrn.1816494

Netzer, O., Lemaire, A., & Herzenstein, M. (2019). When Words Sweat: Identifying Signals for Loan Default in the Text of Loan Applications. *Journal of Marketing Research*, *56*(6), 960–980. https://doi.org/10.1177/0022243719852959

O'Curry, S., & Strahilevitz, M. (2001). Probability and Mode of Acquisition Effects on Choices Between Hedonic and Utilitarian Options. *Marketing Letters*, *12*(1), 37–49. https://doi.org/10.1023/A:1008115902904

O'Donnell, M., & Evers, E. R. K. (2019). Preference Reversals in Willingness to Pay and Choice. *Journal of Consumer Research*, *45*(6), 1315–1330. https://doi.org/10.1093/jcr/ucy052

Olson, J. C., & Muderrisoglu, A. (1979). The stability of responses obtained by free elicitation: Implications for measuring attribute salience and memory structure. *Advances in Consumer Research*, *6*(1).

O'Neill, R., Ralph, J., & Smith, P. A. (2017). *Inflation*. Springer International Publishing. https://doi.org/10.1007/978-3-319-64125-6

Packard, G., & Berger, J. (2021). How Concrete Language Shapes Customer Satisfaction. *Journal of Consumer Research*, *47*(5), 787–806. https://doi.org/10.1093/jcr/ucaa038

Packard, G., & Berger, J. (2024). The Emergence and Evolution of Consumer Language Research. *Journal of Consumer Research*, *51*(1), 42–51. https://doi.org/10.1093/jcr/ucad013

Packard, G., Li, Y., & Berger, J. (2023). When Language Matters. *Journal of Consumer Research*, ucad080. https://doi.org/10.1093/jcr/ucad080

Pecot, F., Valette-Florence, P., & Barnier, V. D. (2019). Brand heritage as a temporal perception: Conceptualisation, measure and consequences. *Journal of Marketing Management*, *35*(17–18), 1624–1643. https://doi.org/10.1080/0267257x.2019.1667414

Pham, M., Geuens, M., & De Pelsmacker, P. (2013). The Influence of Ad-Evoked Feelings on Brand Evaluations: Empirical Generalizations from Consumer Responses to More Than 1,000 TV Commercials. *International Journal of Research in Marketing*, *Forthcoming*. https://doi.org/10.1016/j.ijresmar.2013.04.004

Popescu, O., & Strapparava, C. (2015). Semeval 2015, task 7: Diachronic text evaluation. *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015*, 870–878.

Reichheld, F. F. (2001). *Loyalty rules!: How today's leaders build lasting relationships*. Harvard Business Press.

Reinholtz, N., Bartels, D. M., & Parker, J. R. (2015). On the Mental Accounting of Restricted-Use Funds: How Gift Cards Change What People Purchase. *Journal of Consumer Research*, *42*(4), 596–614.

Rocklage, M. D., He, S., Rucker, D. D., & Nordgren, L. F. (2023). Beyond Sentiment: The Value and Measurement of Consumer Certainty in Language. *Journal of Marketing Research*, *60*(5), 870–888. https://doi.org/10.1177/00222437221134802

Rogers, E. M. (1976). New Product Adoption and Diffusion. *Journal of Consumer Research*, *2*(4), 290. https://doi.org/10.1086/208642

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*(4), 573–605. https://doi.org/10.1016/0010-0285(75)90024-9

Rosenfeld, A., & Erk, K. (2018). Deep neural models of semantic shift. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, *1*, 474–484.

Round, G., & Roper, S. (2017). When and why does the name of the brand still matter? *European Journal of Marketing*, *51*(11/12), 2118–2137. https://doi.org/10.1108/ejm-06-2016-0360

Rust, R. T., Rand, W., Huang, M.-H., Stephen, A. T., Brooks, G., & Chabuk, T. (2021). Real-Time Brand Reputation Tracking Using Social Media. *Journal of Marketing*, *85*(4), 21–43. https://doi.org/10.1177/0022242921995173

Rutz, O. J., & Sonnier, G. P. (2011). The Evolution of Internal Market Structure. *Marketing Science*, *30*(2), 274–289. https://doi.org/10.1287/mksc.1100.0620

Schacter, D. L. (1999). The seven sins of memory: Insights from psychology and cognitive neuroscience. *American Psychologist*, *54*(3), 182–203. https://doi.org/10.1037/0003-066X.54.3.182

Seale-Carlisle, T. M. (2024). Improving the diagnostic value of lineup rejections. *Cognition*, *252*, 105917. https://doi.org/10.1016/j.cognition.2024.105917

Search ad spend in the U.S. 2025. (n.d.). Statista. Retrieved March 17, 2026, from https://www.statista.com/statistics/190275/us-online-display-and-search-advertising-forecast-2010-to-2015/

Shang, J., Liu, J., Jiang, M., Ren, X., Voss, C., & Han, J. (2017). Automated Phrase Mining from Massive Text Corpora. *IEEE Transactions on Knowledge and Data Engineering*, *PP*. https://doi.org/10.1109/TKDE.2018.2812203

Shapiro, S., & Krishnan, H. S. (2001). Memory-Based Measures for Assessing Advertising Effects: A Comparison of Explicit and Implicit Memory Effects. *Journal of Advertising*, *30*(3), 1–13. https://doi.org/10.1080/00913367.2001.10673641

Shefrin, H. M., & Thaler, R. H. (1988). The Behavioral Life-Cycle Hypothesis. *Economic Inquiry*, *26*(4), 609–643. https://doi.org/10.1111/j.1465-7295.1988.tb01520.x

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, *69*(1), 99–118. https://doi.org/10.2307/1884852

Steenkamp, J.-B. E. M. (2019). Global Versus Local Consumer Culture: Theory, Measurement, and Future Research Directions. *Journal of International Marketing*, *27*(1), 1–19. https://doi.org/10.1177/1069031X18811289

Thaler, R. (1985). Mental Accounting and Consumer Choice. *Marketing Science*, *4*(3), 199–214.

Thaler, R. H. (1999a). Mental accounting matters. *Journal of Behavioral Decision Making*, *12*(3), 183–206. https://doi.org/10.1002/(SICI)1099-0771(199909)12:3%3C183::AID-BDM318%3E3.0.CO;2-F

Thaler, R. H. (1999b). Mental Accounting Matters. *Journal of Behavioral Decision Making*, *12*.

Thompson, C. J., & Troester, M. (2002). Consumer value systems in the age of postmodern fragmentation: The case of the natural health microculture. *Journal of Consumer Research*, *28*(4), 550–571.

Tirunillai, S., & Tellis, G. J. (2014). Mining Marketing Meaning from Online Chatter: Strategic Brand Analysis of Big Data Using Latent Dirichlet Allocation. *Journal of Marketing Research*, *51*(4), 463–479. https://doi.org/10.1509/jmr.12.0106

Traugott, E. C., & Dasher, R. B. (2001). *Regularity in Semantic Change*. Cambridge University Press. https://doi.org/10.1017/CBO9780511486500

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458. https://doi.org/10.1126/science.7455683

Unnava, H. R., & Burnkrant, R. E. (1991a). An Imagery-Processing View of the Role of Pictures in Print Advertisements. *Journal of Marketing Research*, *28*(2), 226–231. https://doi.org/10.1177/002224379102800210

Unnava, H. R., & Burnkrant, R. E. (1991b). Effects of Repeating Varied Ad Executions on Brand Name Memory. *Journal of Marketing Research*, *28*(4), 406–416. https://doi.org/10.1177/002224379102800403

Urban, M. (2015). Lexical semantic change and semantic reconstruction. In *The Routledge handbook of historical linguistics* (pp. 374–392). Routledge.

Urminsky, O., & Dietvorst, B. J. (2024). Taking the Full Measure: Integrating Replication into Research Practice to Assess Generalizability. *Journal of Consumer Research*, *51*(1), 157–168. https://doi.org/10.1093/jcr/ucae007

Vakratsas, D., & Ambler, T. (1999a). How Advertising Works: What Do We Really Know? *Journal of Marketing*, *63*(1), 26–43. https://doi.org/10.1177/002224299906300103

Vakratsas, D., & Ambler, T. (1999b). How advertising works: What do we really know? *Journal of Marketing*, *63*(1), 26–43. https://doi.org/10.2307/1251999

Villarroel Ordenes, F., Grewal, D., Ludwig, S., Ruyter, K. D., Mahr, D., & Wetzels, M. (2019). Cutting through Content Clutter: How Speech and Image Acts Drive Consumer Sharing of Social Media Brand Messages. *Journal of Consumer Research*, *45*(5), 988–1012. https://doi.org/10.1093/jcr/ucy032

Villas-Boas, J. M. (2018). A Dynamic Model of Repositioning. *Marketing Science*, *37*(2), 279–293. https://doi.org/10.1287/mksc.2017.1075

Wang, X. (Shane), He, J., Curry, D. J., & Ryoo, J. H. (Joseph). (2022). Attribute Embedding: Learning Hierarchical Representations of Product Attributes from Consumer Reviews. *Journal of Marketing*, *86*(6), 155–175. https://doi.org/10.1177/00222429211047822

Weinreich, U., Labov, W., & Herzog, M. (1968). *Empirical Foundations for a Theory of Language Change*. University of Texas Press. https://books.google.com/books?id=Wr1CAAAAIAAJ

Yao, Z., Sun, Y., Ding, W., Rao, N., & Xiong, H. (2018). Dynamic word embeddings for evolving semantic discovery. *Proceedings of the Eleventh Acm International Conference on Web Search and Data Mining*, 673–681.

Zhou, L., Teng, L., & Poon, P. S. (2008). Susceptibility to global consumer culture: A three-dimensional scale. *Psychology & Marketing*, *25*(4), 336–351. https://doi.org/10.1002/mar.20212

Appendices

4.1 APPENDIX 1A: COMPLETE LIST OF ATTRIBUTES USED IN CHAPTER 1

Arrogant	Intelligent
Authentic	Kind
Best Brand	Leader
Carefree	Original
Cares Customers	Prestigious
Charming	Progressive
Daring	Reliable
Different	Rugged
Distinctive	Sensuous
Down To Earth	Simple
Dynamic	Social
Energetic	Socially Responsible
Friendly	Straightforward
Fun	Stylish
Gaining In Popularity	Traditional
Glamorous	Trendy
Good Value	Trustworthy
Healthy	Unapproachable
Helpful	Unique
High Performance	Up To Date
High Quality	Upper Class
Independent	Visionary
Innovative	Worth More

4.2 APPENDIX 1B: SEMASIOLOGICAL AND ONOMASIOLOGICAL CHANGE OF "VIRAL" AND "BLACKBERRY"

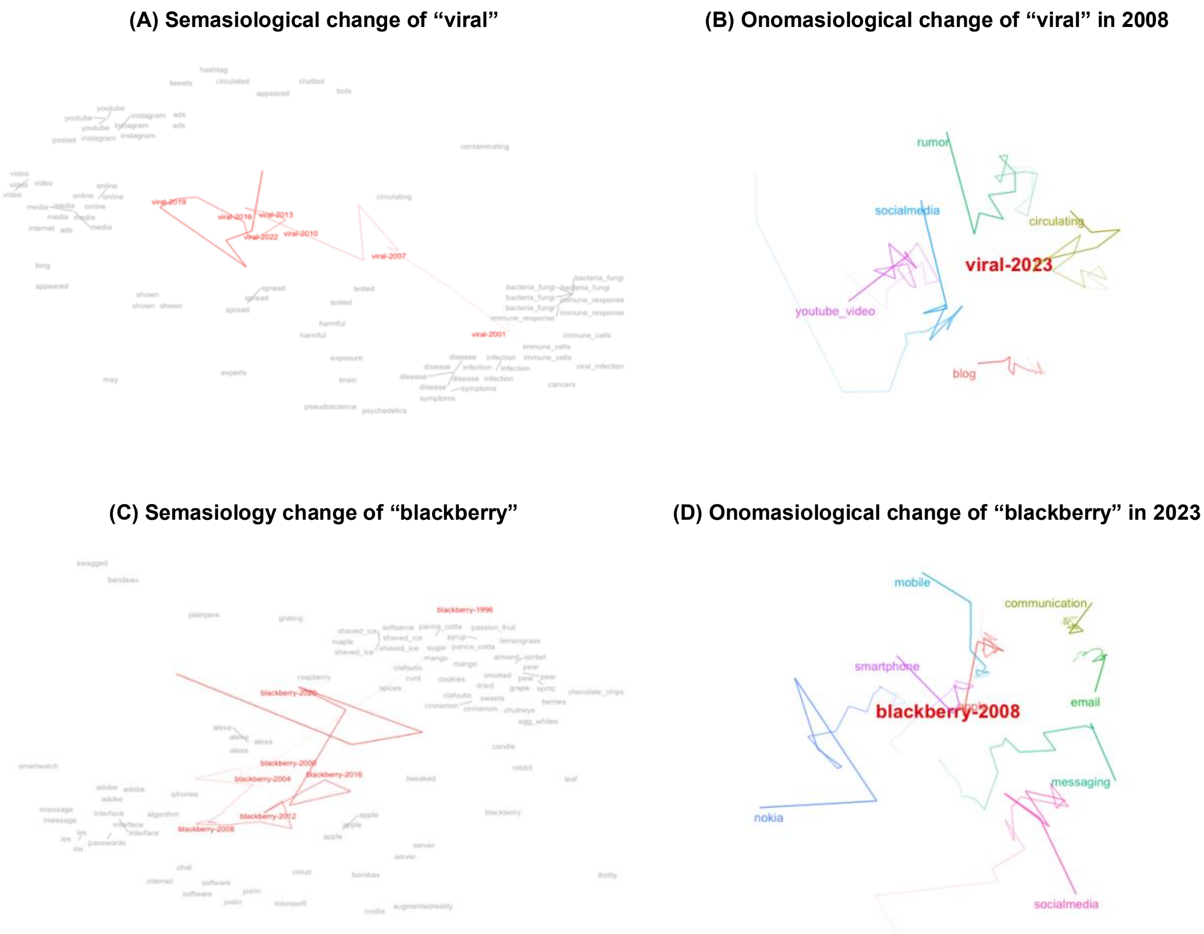

Figure W0.1: (A and C) The semasiological change of the words "blackberry" and "viral". The words in grey are context words, in their 2023 meaning, used to illustrate the meanings of the focal words at different times. (B and D) The onomasiological change of the meanings denoted by the words "blackberry" in 2008 and "viral" in 2023.

4.3 APPENDIX 1C: MEANING CHANGE OF ALL BRANDS


Campbells: 2001-2020
Friendly
Best Brand
Socially Responsible
Kind
Fun
Arrogant
Social
Unapproachable
Best Brand
Daring
Stylish
Glamorous
Upper Class
Prestigious
Unapproachable
Gaining In Popularity
Stylish
Cares Customers
Innovative
Trendy
High Performance Leader
Progressive Trustworthy
Dynamic Traditional
Rugged Cares Customers
Trustworthy Reliable
Canon: 2001-2020
Rugged
Stylish
Energetic
High Quality
Distinctive
Rugged
Visionary
Authentic
Traditional
Daring
Progressive Intelligent
Innovative Progressive
Visionary High Performance
Dynamic Up To Date
Up To Date Innovative
Chevrolet: 2001-2020
Social
Gaining In Popularity
Helpful
Independent
Different
Gaining In Popularity
Progressive
Independent
Up To Date
Innovative
Dynamic High Performance
Stylish Daring
High Performance Stylish
Distinctive Visionary
Daring Dynamic
Chobani: 2010-2020
Authentic Authentic
High Quality Worth More
Worth More Social
Dynamic Leader
High Performance Best Brand
Chromecast: 2014-2020
Helpful
Intelligent
Visionary
Innovative
Progressive
Kind
Charming
Friendly
Fun
Healthy
Simple
Traditional
Straightforward
Kind
Down To Earth
Simple
Kind
Original
Friendly
Straightforward
Innovative Progressive
Progressive Innovative
Up To Date Intelligent
Visionary Up To Date
Intelligent Dynamic
Chrysler: 2001-2020
Upper Class
Prestigious
High Quality
Arrogant
High Performance
Visionary
Arrogant
Unapproachable
High Performance
Rugged
Dynamic Independent
Progressive Visionary
Visionary Progressive
High Performance High Performance
Distinctive Up To Date
Coca Cola: 2001-2020
Straightforward
Helpful
Good Value
Kind
Down To Earth
Helpful
Straightforward
Good Value
Simple
Down To Earth
Helpful Healthy
Simple Helpful
Straightforward Sensuous
Healthy Different
Glamorous Simple
Coke: 2014-2020
Authentic Leader
Worth More Authentic
Best Brand Original
Trendy Best Brand
High Quality Worth More
Compaq: 2001-2020
High Performance
High Quality
Sensuous
Rugged
Stylish
Glamorous
Distinctive
Upper Class
Prestigious
Different
Helpful
Simple
Healthy
Down To Earth
Straightforward
Intelligent
Helpful
Straightforward
Healthy
Simple
Progressive Progressive
Innovative Intelligent
Visionary Up To Date
Intelligent Innovative
Up To Date High Performance
Coors: 2001-2020
Arrogant
High Quality
Prestigious
Distinctive
Worth More
Trendy
Arrogant
Distinctive
Rugged
Social
Leader Leader
Dynamic Gaining In Popularity
Authentic Socially Responsible
High Performance Cares Customers
Sensuous Authentic
Coors Light: 2001-2020
Good Value
Kind
Straightforward
Healthy
Simple
Helpful
Good Value
Simple
Intelligent
Healthy
Best Brand Social
Leader Authentic
Authentic Leader
Up To Date Worth More
Social Best Brand
Crush: 2001-2020
Unapproachable Sensuous
Upper Class Upper Class
Prestigious Glamorous
Glamorous Prestigious
Helpful Unapproachable
Dannon: 2001-2020
Original
Down To Earth
Traditional
Simple
Trustworthy
Original
Carefree
Simple
Fun
Healthy
Authentic
Worth More
Trendy
Progressive
Best Brand
Leader
Gaining In Popularity
Authentic
Worth More
Social
High Quality Gaining In Popularity
Authentic Leader
Trendy Worth More
Worth More Trendy
Dynamic Authentic
Dasani: 2001-2020
Simple
Down To Earth
Straightforward
Healthy
Helpful
Fun
Healthy
Energetic
Simple
Different
Rugged Simple
Different Visionary
Simple Helpful
Social Good Value
Helpful Down To Earth
Del Monte: 2001-2020
Trustworthy
Traditional
Reliable
Straightforward
Cares Customers
Independent
Traditional
Unique
Different
Simple
Authentic Independent
Leader Leader
Worth More Distinctive
High Quality Authentic
Trustworthy Unique
Dell: 2001-2020
Traditional Carefree
Straightforward Healthy
Down To Earth Traditional
Simple Simple
Original Different
Dial: 2001-2020
Trustworthy
Sensuous
Authentic
High Performance
High Quality
Reliable
Trustworthy
Traditional
High Performance
Authentic
Arrogant
Best Brand
Worth More
Distinctive
High Quality
Original
Social
Arrogant
Unapproachable
Sensuous
Innovative Progressive
Up To Date Intelligent
Progressive Dynamic
Gaining In Popularity Helpful
Visionary Innovative
Diet Coke: 2001-2020
Socially Responsible
Cares Customers
Unapproachable
Independent
Kind
Independent
Social
Gaining In Popularity
Friendly
Unapproachable
Helpful Helpful
Simple Intelligent
Rugged Visionary
Intelligent Unapproachable
Healthy Rugged
Disney: 2001-2020
Straightforward
Social
Good Value
Carefree
Friendly
Kind
Down To Earth
Best Brand
Friendly
Cares Customers
Charming Different
Independent Carefree
Unique Innovative
Kind Fun
Sensuous Charming
Dodge: 2001-2020
Simple Healthy
Healthy Simple
Kind Carefree
Social Kind
Carefree Social

Dove Soap: 2001-2020

Dropbox: 2014-2020

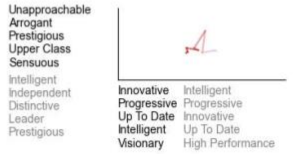

eBay: 2003-2020

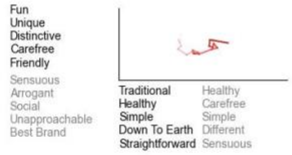

Evian: 2001-2020

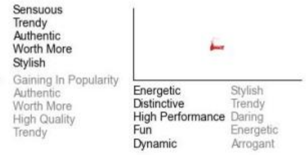

Expedia: 2003-2020

Facebook: 2006-2020

Fanta: 2001-2020

Fiji: 2006-2020

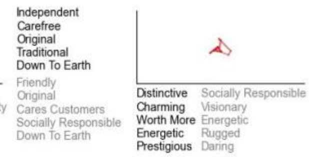

Fila: 2001-2020

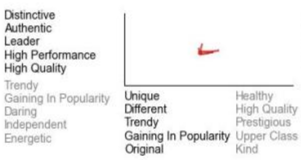

Flickr: 2007-2020

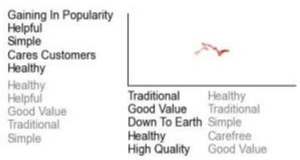

Ford: 2001-2020

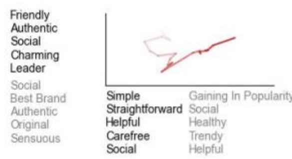

Garmin: 2008-2020

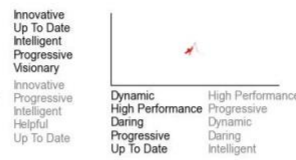

General Motors: 2001-2020

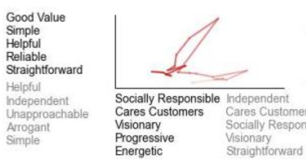

Gmail: 2004-2020

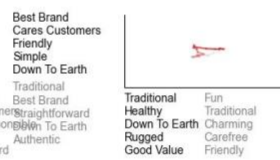

Goldman Sachs: 2001-2020

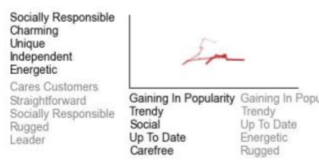

GoPro: 2017-2020

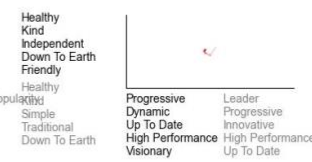

Greyhound: 2001-2020

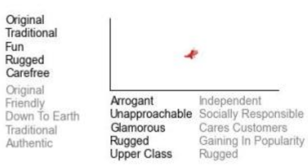

Heineken: 2001-2020

Hershey's: 2001-2020

Hewlett Packard: 2001-2020

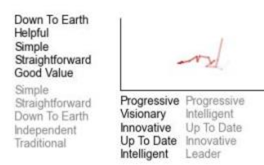

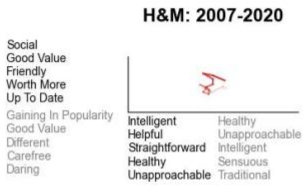
H&M: 2007-2020
Social
Good Value
Friendly
Worth More
Up To Date
Gaining In Popularity
Good Value
Different
Carefree
Daring
Intelligent
Helpful
Straightforward
Healthy
Unapproachable
Healthy
Unapproachable
Intelligent
Sensuous
Traditional

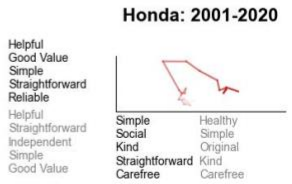
Honda: 2001-2020
Helpful
Good Value
Simple
Straightforward
Reliable
Helpful
Straightforward
Independent
Simple
Good Value
Simple
Social
Kind
Straightforward
Carefree
Healthy
Simple
Original
Kind
Carefree

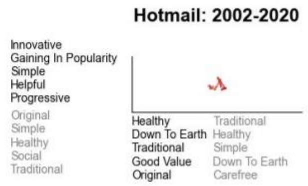
Hotmail: 2002-2020
Innovative
Gaining In Popularity
Simple
Helpful
Progressive
Original
Simple
Healthy
Social
Traditional
Healthy
Down To Earth
Traditional
Good Value
Original
Traditional
Healthy
Simple
Down To Earth
Carefree

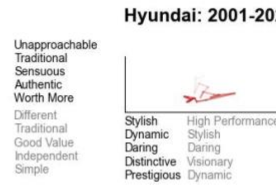
Hyundai: 2001-2020
Unapproachable
Traditional
Sensuous
Authentic
Worth More
Different
Traditional
Good Value
Independent
Simple
Stylish
Dynamic
Daring
Distinctive
Prestigious
High Performance
Stylish
Daring
Visionary
Dynamic

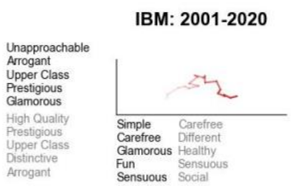
IBM: 2001-2020
Unapproachable
Arrogant
Upper Class
Prestigious
Glamorous
High Quality
Prestigious
Upper Class
Distinctive
Arrogant
Simple
Carefree
Glamorous
Fun
Sensuous
Carefree
Different
Healthy
Sensuous
Social

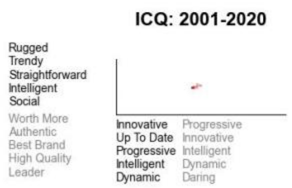
ICQ: 2001-2020
Rugged
Trendy
Straightforward
Intelligent
Social
Worth More
Authentic
Best Brand
High Quality
Leader
Innovative
Up To Date
Progressive
Intelligent
Dynamic
Progressive
Innovative
Intelligent
Dynamic
Daring

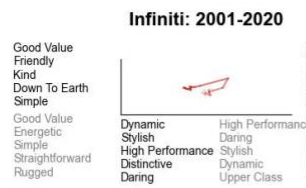
Infiniti: 2001-2020
Good Value
Friendly
Kind
Down To Earth
Simple
Good Value
Energetic
Simple
Straightforward
Rugged
Dynamic
Stylish
High Performance
Distinctive
Daring
High Performance
Daring
Stylish
Dynamic
Upper Class

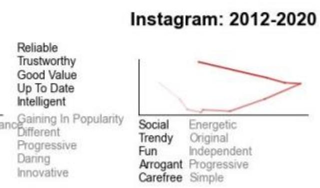
Instagram: 2012-2020
Reliable
Trustworthy
Good Value
Up To Date
Intelligent
Gaining In Popularity
Different
Progressive
Daring
Innovative
Social
Trendy
Fun
Arrogant
Carefree
Energetic
Original
Independent
Progressive
Simple

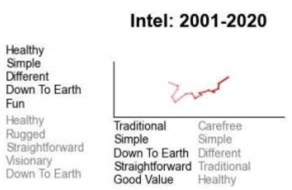
Intel: 2001-2020
Healthy
Simple
Different
Down To Earth
Fun
Healthy
Rugged
Straightforward
Visionary
Down To Earth
Traditional
Simple
Down To Earth
Straightforward
Good Value
Carefree
Simple
Different
Traditional
Healthy

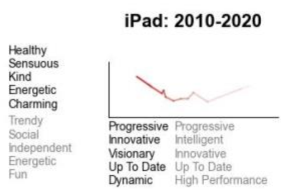
iPad: 2010-2020
Healthy
Sensuous
Kind
Energetic
Charming
Trendy
Social
Independent
Energetic
Fun
Progressive
Innovative
Visionary
Up To Date
Dynamic
Progressive
Intelligent
Innovative
Up To Date
High Performance

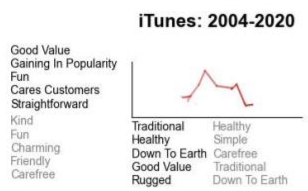
iTunes: 2004-2020
Good Value
Gaining In Popularity
Fun
Cares Customers
Straightforward
Kind
Fun
Charming
Friendly
Carefree
Traditional
Healthy
Down To Earth
Good Value
Rugged
Healthy
Simple
Carefree
Traditional
Down To Earth

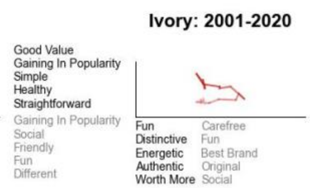
Ivory: 2001-2020
Good Value
Gaining In Popularity
Simple
Healthy
Straightforward
Gaining In Popularity
Social
Friendly
Fun
Different
Fun
Distinctive
Energetic
Authentic
Worth More
Carefree
Fun
Best Brand
Original
Social

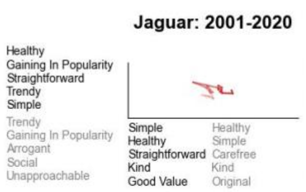
Jaguar: 2001-2020
Healthy
Gaining In Popularity
Straightforward
Trendy
Simple
Trendy
Gaining In Popularity
Arrogant
Social
Unapproachable
Simple
Healthy
Straightforward
Kind
Good Value
Healthy
Simple
Carefree
Kind
Original

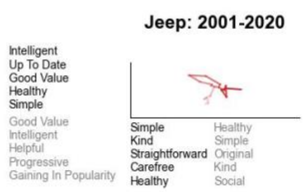
Jeep: 2001-2020
Intelligent
Up To Date
Good Value
Healthy
Simple
Good Value
Intelligent
Helpful
Progressive
Gaining In Popularity
Simple
Kind
Straightforward
Carefree
Healthy
Healthy
Simple
Original
Kind
Social

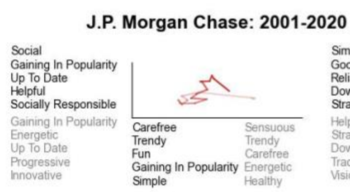
J.P. Morgan Chase: 2001-2020
Social
Gaining In Popularity
Up To Date
Helpful
Socially Responsible
Gaining In Popularity
Energetic
Up To Date
Progressive
Innovative
Carefree
Trendy
Fun
Gaining In Popularity
Simple
Sensuous
Trendy
Carefree
Energetic
Healthy

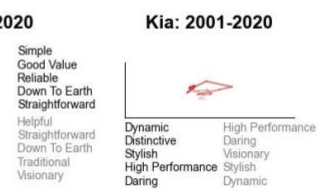
Kia: 2001-2020
Simple
Good Value
Reliable
Down To Earth
Straightforward
Helpful
Straightforward
Down To Earth
Traditional
Visionary
Dynamic
Distinctive
Stylish
High Performance
Daring
High Performance
Daring
Visionary
Stylish
Dynamic


Kindle: 2009-2020
Social
Up To Date
Fun
Friendly
Good Value
Good Value
Reliable
Straightforward
Down To Earth
Trustworthy
Progressive
Innovative
Up To Date
Visionary
Dynamic
Progressive
Up To Date
Intelligent
Innovative
Helpful

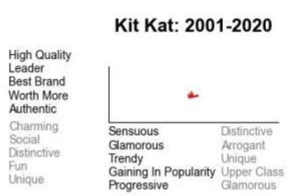
Kit Kat: 2001-2020
High Quality
Leader
Best Brand
Worth More
Authentic
Charming
Social
Distinctive
Fun
Unique
Sensuous
Glamorous
Trendy
Gaining In Popularity
Progressive
Distinctive
Arrogant
Unique
Upper Class
Glamorous

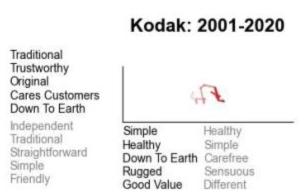
Kodak: 2001-2020
Traditional
Trustworthy
Original
Cares Customers
Down To Earth
Independent
Traditional
Straightforward
Simple
Friendly
Simple
Healthy
Down To Earth
Rugged
Good Value
Healthy
Simple
Carefree
Sensuous
Different

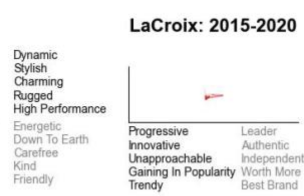
LaCroix: 2015-2020
Dynamic
Stylish
Charming
Rugged
High Performance
Energetic
Down To Earth
Carefree
Kind
Friendly
Progressive
Innovative
Unapproachable
Gaining In Popularity
Trendy
Leader
Authentic
Independent
Worth More
Best Brand


Lenovo: 2006-2020
Energetic
Distinctive
Dynamic
Fun
Trendy
Visionary
Rugged
Trendy
Energetic
Stylish
Progressive
Dynamic
Up To Date
Intelligent
High Performance
Progressive
Intelligent
Innovative
Up To Date
High Performance
Lexus: 2001-2020
Simple
Good Value
Kind
Healthy
Independent
Gaining In Popularity
Good Value
Simple
Helpful
Healthy
Simple
Straightforward
Healthy
Kind
Carefree
Healthy
Simple
Kind
Original
Carefree
LG: 2002-2020
Reliable
Trustworthy
Leader
Worth More
High Quality
Different
Unique
Fun
Independent
Gaining In Popularity
Progressive
Innovative
Up To Date
Dynamic
Visionary
Progressive
Innovative
Intelligent
Up To Date
High Performance
Lincoln: 2012-2020
Different
Unapproachable
Progressive
Innovative
Unique
Trendy
Stylish
Upper Class
Prestigious
Glamorous
Social
Sensuous
Helpful
Straightforward
Carefree
Good Value
Simple
Gaining In Popularity
Different
Down To Earth
Lululemon: 2009-2020
Gaining In Popularity
Visionary
Innovative
Best Brand
Progressive
Innovative
Dynamic
Progressive
Daring
Unique
Stylish
Trendy
Dynamic
Energetic
Sensuous
Trendy
Stylish
Up To Date
High Performance
Leader
Marlboro: 2001-2020
Rugged
High Performance
Reliable
Stylish
Upper Class
Reliable
Cares Customers
Independent
Trustworthy
Down To Earth
Healthy
Different
Gaining In Popularity
Simple
Helpful
Carefree
Different
Simple
Healthy
Fun
Mazda: 2001-2020
Simple
Good Value
Helpful
Down To Earth
Reliable
Helpful
Gaining In Popularity
Good Value
Straightforward
Simple
Dynamic
Stylish
Distinctive
High Performance
Prestigious
High Performance
Daring
Stylish
Visionary
Dynamic
Megabus: 2018-2020
Rugged
Upper Class
Arrogant
Unapproachable
Fun
Different
Unique
Daring
Unapproachable
Rugged
Rugged
Daring
Unique
Distinctive
Stylish
Socially Responsible
Cares Customers
Independent
Visionary
Rugged
Mercedes Benz: 2001-2020
Socially Responsible
Gaining In Popularity
Independent
Up To Date
Visionary
Gaining In Popularity
Progressive
Intelligent
Up To Date
Helpful
Simple
Straightforward
Kind
Healthy
Carefree
Healthy
Simple
Kind
Good Value
Original
Microsoft: 2001-2020
Down To Earth
Independent
Original
Traditional
Cares Customers
Trendy
Gaining In Popularity
Fun
Different
Energetic
Progressive
Visionary
Socially Responsible
Innovative
Up To Date
Progressive
Up To Date
Intelligent
Leader
Innovative
Miller Lite: 2001-2020
Daring
Dynamic
Upper Class
Arrogant
Fun
Daring
Innovative
High Performance
Dynamic
Visionary
Trendy
Authentic
Dynamic
Daring
Stylish
Gaining In Popularity
Social
Leader
Worth More
Trendy
Mini Cooper: 2003-2020
Trendy
Social
Sensuous
Carefree
Energetic
Trendy
Worth More
High Quality
Authentic
Energetic
Simple
Healthy
Straightforward
Kind
Good Value
Healthy
Simple
Kind
Good Value
Helpful
Mountain Dew: 2001-2020
Social
Fun
Original
Carefree
Helpful
Fun
Down To Earth
Friendly
Simple
Original
Trendy
Dynamic
Energetic
Up To Date
Authentic
Worth More
Social
Gaining In Popularity
Authentic
Leader
MSN: 2002-2020
Leader
Traditional
Authentic
Trustworthy
Kind
Social
Carefree
Fun
Friendly
Original
Innovative
Up To Date
Progressive
Visionary
Dynamic
Progressive
Innovative
Up To Date
Intelligent
Dynamic
Myspace: 2005-2020
Intelligent
Innovative
Progressive
Up To Date
Unapproachable
Progressive
Intelligent
Innovative
Up To Date
Dynamic
Progressive
Social
Innovative
Up To Date
Gaining In Popularity
Up To Date
Progressive
Energetic
Innovative
Leader
Napster: 2004-2020
Unique
Innovative
Different
Progressive
Daring
Intelligent
Sensuous
Helpful
Up To Date
Innovative
Up To Date
Innovative
Progressive
Visionary
Intelligent
Progressive
Innovative
Up To Date
Intelligent
High Performance
Nestlé: 2001-2020
Different
Unique
Trendy
Gaining In Popularity
Social
Prestigious
Upper Class
Arrogant
Unapproachable
Healthy
Helpful
Simple
Rugged
Straightforward
Down To Earth
Rugged
Visionary
Healthy
Sensuous
Helpful
Netflix: 2005-2020
Progressive
Innovative
Socially Responsible
Visionary
Dynamic
Progressive
Up To Date
Innovative
Intelligent
High Performance
Fun
Original
Best Brand
Visionary
Leader
Fun
Charming
Kind
Friendly
Carefree
Netscape: 2001-2020
Gaining In Popularity
Up To Date
Cares Customers
Helpful
Socially Responsible
Helpful
Good Value
Gaining In Popularity
Down To Earth
Simple
Innovative
Up To Date
Progressive
Intelligent
Visionary
Progressive
Intelligent
Innovative
Up To Date
Dynamic
Neutrogena: 2001-2020
Helpful
Rugged
Independent
Social
Innovative
Independent
Intelligent
Rugged
Visionary
Progressive
Energetic
Trendy
Authentic
Leader
Social
Worth More
Leader
Authentic
High Quality
Best Brand


Nike: 2001-2020
Sensuous Helpful Unapproachable Kind Straightforward
Different Simple Helpful Sensuous Healthy
Stylish Trendy Charming Sensuous Energetic
Stylish Up To Date Trendy Energetic Visionary
Nintendo: 2003-2020
Arrogant Unapproachable Healthy Sensuous Trendy
Trendy Healthy Good Value Gaining In Popularity Rugged
Simple Healthy Down To Earth Traditional Good Value
Healthy Simple Traditional Carefree Down To Earth
Nissan: 2001-2020
Progressive Visionary Socially Responsible Reliable Unique
Independent Progressive Intelligent High Performance Up To Date
Simple Kind Straightforward Helpful Healthy
Healthy Simple Original Carefree Kind
Nvidia: 2005-2020
Innovative Progressive Visionary Intelligent Up To Date
Intelligent Up To Date Progressive Helpful Reliable
Progressive Intelligent Dynamic Visionary Innovative
Progressive Innovative Intelligent Up To Date Gaining In Popularity
Olay: 2001-2020
Unapproachable Upper Class Innovative Prestigious Glamorous
Innovative Progressive Stylish Up To Date Intelligent
Dynamic Authentic High Performance Trendy Sensuous
Worth More Leader High Quality Gaining In Popularity Authentic
Oldsmobile: 2001-2020
Rugged Arrogant Upper Class Unapproachable Traditional
Rugged Down To Earth Traditional Straightforward Original
Stylish Dynamic High Performance Rugged Prestigious
Stylish High Performance Daring Visionary Dynamic
Pandora: 2009-2020
Innovative Progressive Intelligent Unique Dynamic
Sensuous Glamorous Dynamic High Performance Upper Class
Traditional Healthy Rugged Down To Earth Good Value
Healthy Traditional Simple Sensuous Carefree
PayPal: 2003-2020
Progressive Innovative Up To Date Trendy Gaining In Popularity
Sensuous Worth More High Performance High Quality Progressive
Healthy Simple Good Value Down To Earth Charming
Healthy Simple Sensuous Carefree Down To Earth
Pepsi: 2014-2020
Fun Social Rugged Trendy Gaining In Popularity
Fun Energetic Rugged Carefree Daring
Helpful Straightforward Simple Down To Earth Good Value
Healthy Helpful Simple Straightforward Kind
Perrier: 2001-2020
Helpful Simple Innovative Social Different
Different Healthy Simple Fun Unique
Trendy High Quality Authentic Sensuous Glamorous
Authentic Trendy Distinctive Social Worth More
Pinterest: 2012-2020
Unapproachable Intelligent Visionary Progressive Arrogant
Intelligent Independent Progressive Innovative Up To Date
Progressive Innovative Trendy Up To Date Dynamic
Up To Date Progressive High Performance High Quality Worth More
Polaroid: 2001-2020
High Performance Up To Date Dynamic Visionary Innovative
Trustworthy Best Brand Authentic Innovative Original
Visionary Intelligent Progressive Innovative Dynamic
Progressive Intelligent Leader Up To Date Innovative
Pontiac: 2001-2020
Helpful Traditional Trustworthy Reliable Kind
Straightforward Helpful Simple Good Value Down To Earth
Dynamic Stylish High Performance Daring Distinctive
Stylish High Performance Daring Dynamic Visionary
Porsche: 2001-2020
Reliable Trustworthy Good Value Straightforward Leader
Independent Helpful Intelligent Arrogant Unapproachable
Simple Straightforward Social Kind Healthy
Healthy Simple Original Carefree Kind
Puma: 2001-2020
Straightforward Rugged Fun Distinctive Down To Earth
Visionary Rugged Daring Dynamic Unapproachable
Stylish Glamorous Trendy Daring Prestigious
Stylish Trendy Distinctive High Quality Glamorous
Reebok: 2001-2020
Distinctive Healthy Upper Class Fun Prestigious
Stylish Visionary Rugged Arrogant Unapproachable
Energetic Trendy Stylish Dynamic Authentic
Trendy Energetic Stylish Daring Up To Date
Reese's: 2001-2020
Leader Best Brand Straightforward Authentic Cares Customers
Social Carefree Fun Original Friendly
Trendy High Quality Stylish Worth More Authentic
Authentic Original Distinctive Trendy Worth More
Samsung: 2001-2020
Original Traditional Down To Earth Good Value Simple
Original Traditional Simple Authentic Friendly
Traditional Healthy Simple Down To Earth Rugged
Healthy Carefree Sensuous Simple Down To Earth
Saturn: 2001-2020
Visionary Intelligent Innovative Progressive Rugged
Charming Socially Responsible Intelligent Unique Fun
Stylish Dynamic High Performance Daring Distinctive
High Performance Stylish Daring Progressive Visionary
Sears: 2001-2020
Intelligent Traditional Unapproachable Sensuous Straightforward
Different Simple Traditional Unique Intelligent
Simple Straightforward Healthy Fun Down To Earth
Healthy Carefree Simple Different Rugged

Skype: 2005-2020

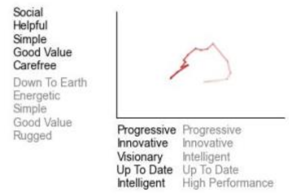

Slack: 2018-2020

Smart: 2012-2020

Snapchat: 2013-2020

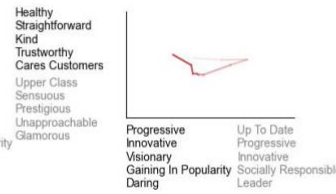

Snickers: 2001-2020

Sony: 2001-2020

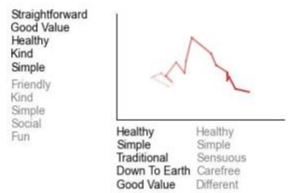

Soylent: 2018-2020

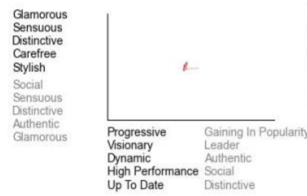

Splenda: 2005-2020

Spotify: 2012-2020

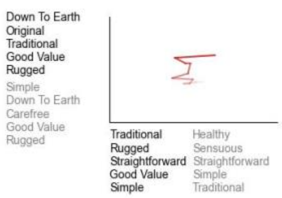

Sprite: 2001-2020

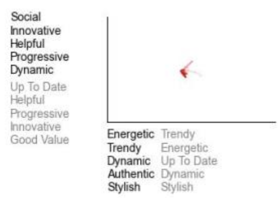

Starbucks: 2001-2020

Subaru: 2001-2020

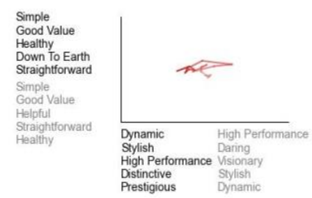

Tencent: 2019-2020

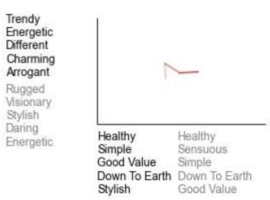

Tesla: 2011-2020

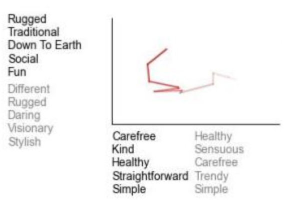

Tivo: 2001-2020

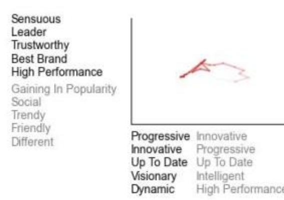

Toyota: 2001-2020

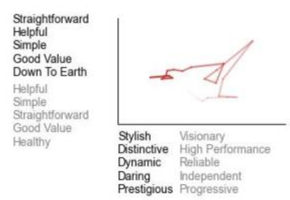

Tumblr: 2011-2020

Twitter: 2009-2020

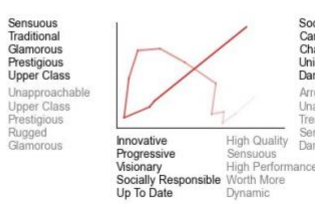

Uber: 2014-2020

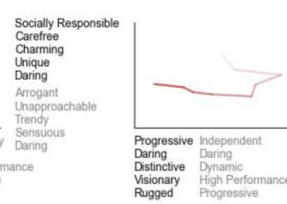

Uniqlo: 2007-2020

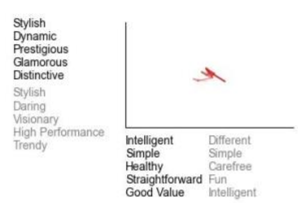

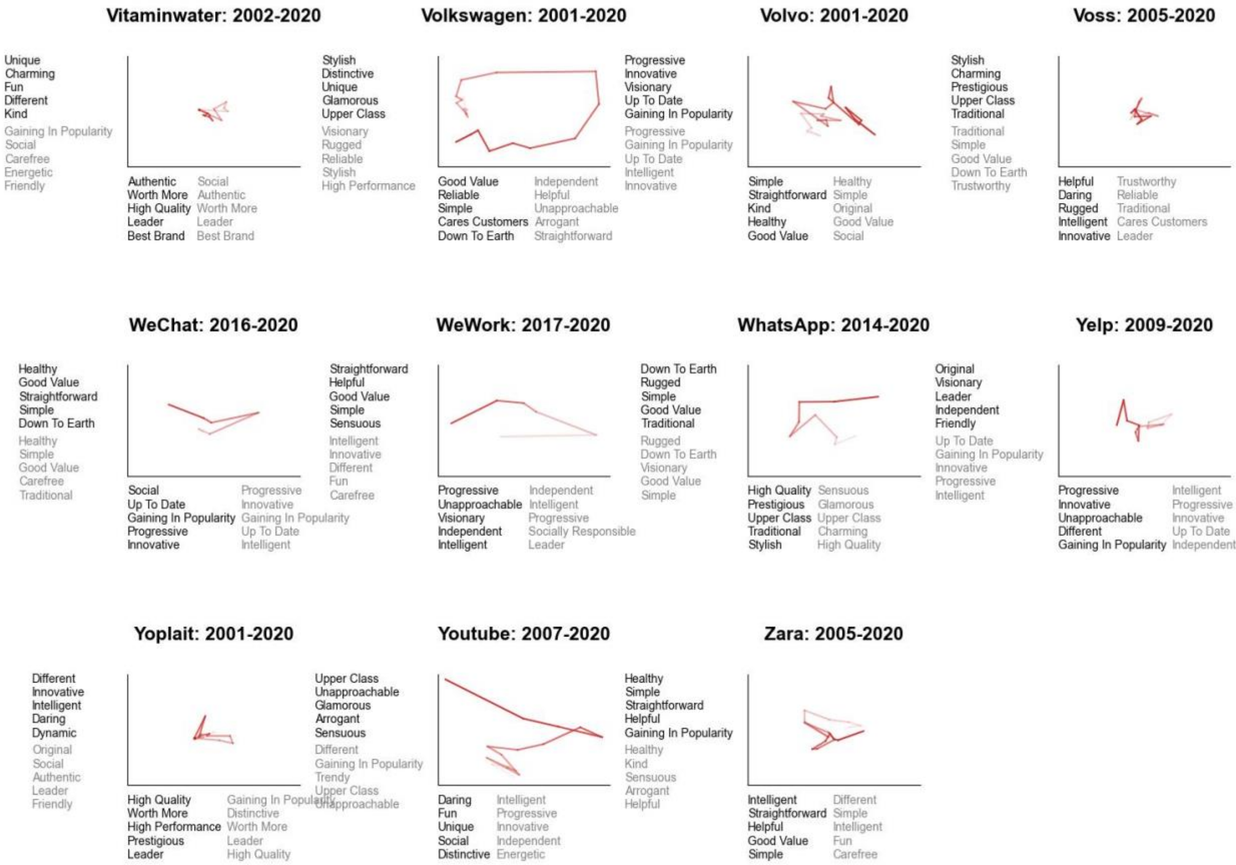

Vitaminwater: 2002-2020
Unique
Charming
Fun
Different
Kind
Gaining In Popularity
Social
Carefree
Energetic
Friendly
Stylish
Distinctive
Unique
Glamorous
Upper Class
Visionary
Rugged
Reliable
Stylish
High Performance
Authentic
Worth More
High Quality
Leader
Best Brand
Social
Authentic
Worth More
Leader
Best Brand
Volkswagen: 2001-2020
Progressive
Innovative
Visionary
Up To Date
Gaining In Popularity
Progressive
Gaining In Popularity
Up To Date
Intelligent
Innovative
Good Value
Reliable
Simple
Cares Customers
Down To Earth
Independent
Helpful
Unapproachable
Arrogant
Straightforward
Volvo: 2001-2020
Stylish
Charming
Prestigious
Upper Class
Traditional
Traditional
Simple
Good Value
Down To Earth
Trustworthy
Simple
Straightforward
Kind
Healthy
Good Value
Healthy
Simple
Original
Good Value
Social
Voss: 2005-2020
Helpful
Daring
Rugged
Intelligent
Innovative
Trustworthy
Reliable
Traditional
Cares Customers
Leader
WeChat: 2016-2020
Healthy
Good Value
Straightforward
Simple
Down To Earth
Healthy
Simple
Good Value
Carefree
Traditional
Straightforward
Helpful
Good Value
Simple
Sensuous
Intelligent
Innovative
Different
Fun
Carefree
Social
Up To Date
Gaining In Popularity
Progressive
Innovative
Progressive
Innovative
Gaining In Popularity
Up To Date
Intelligent
WeWork: 2017-2020
Down To Earth
Rugged
Simple
Good Value
Traditional
Rugged
Down To Earth
Visionary
Good Value
Simple
Progressive
Unapproachable
Visionary
Independent
Intelligent
Independent
Intelligent
Progressive
Socially Responsible
Leader
WhatsApp: 2014-2020
Original
Visionary
Leader
Independent
Friendly
Up To Date
Gaining In Popularity
Innovative
Progressive
Intelligent
High Quality
Prestigious
Upper Class
Traditional
Stylish
Sensuous
Glamorous
Upper Class
Charming
High Quality
Yelp: 2009-2020
Progressive
Innovative
Unapproachable
Different
Gaining In Popularity
Intelligent
Progressive
Innovative
Up To Date
Independent
Yoplait: 2001-2020
Different
Innovative
Intelligent
Daring
Dynamic
Original
Social
Authentic
Leader
Friendly
Upper Class
Unapproachable
Glamorous
Arrogant
Sensuous
Different
Gaining In Popularity
Trendy
Upper Class
Unapproachable
High Quality
Worth More
High Performance
Prestigious
Leader
Gaining In Popularity
Distinctive
Worth More
Leader
High Quality
Youtube: 2007-2020
Healthy
Simple
Straightforward
Helpful
Gaining In Popularity
Healthy
Kind
Sensuous
Arrogant
Helpful
Daring
Fun
Unique
Social
Distinctive
Intelligent
Progressive
Innovative
Independent
Energetic
Zara: 2005-2020
Intelligent
Straightforward
Helpful
Good Value
Simple
Different
Simple
Intelligent
Fun
Carefree

4.4 APPENDIX 1D: EXAMPLES OF STABLE AND UNSTABLE PAIRS USED IN STUDY 2B (CHAPTER 1)

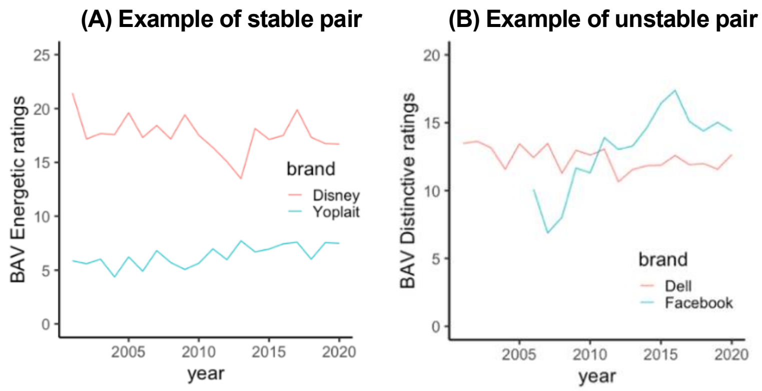

Figure W0.2: (A) A pair of brands, Disney and Yoplait, where participants can use their current knowledge to correctly answer questions about past relationship on the attribute *Energetic*. These pairs are referred to as *Stable*. (B) A pair of brands, Dell and Facebook, where reliance on current knowledge yields incorrect answers on past questions. We refer to these as *Unstable*.

4.5 APPENDIX 1E: ALL SYNONYMS USED IN STUDY 3B (CHAPTER 1)

Focal attribute	**Synonyms in Past condition**	**Synonyms in Present condition**
Best Brand	Authentic	Leader
	Best Brand	Worth More
	Leader	Authentic
	Energetic	High Performance
	Fun	Visionary
Cares Customers	Independent	Trustworthy
	Socially Responsible	Reliable
	Cares Customers	Good Value
	Original	Down To Earth
	Friendly	Kind
Daring	Daring	Dynamic
	Dynamic	Visionary
	Distinctive	Unique
	Stylish	Progressive
	Energetic	Innovative
Gaining In Popularity	Gaining In Popularity	Trendy
	Socially Responsible	Innovative
	Social	Up To Date
	Independent	Different
	Trendy	Progressive
Healthy	Healthy	Kind

	Simple Kind Carefree Straightforward	Simple Straightforward Good Value Cares Customers
High Performance	Daring Dynamic Stylish Progressive High Performance	High Quality Worth More Dynamic Leader Intelligent
High Quality	Daring Trendy Dynamic Stylish Distinctive	Worth More High Performance Prestigious Distinctive Upper Class
Intelligent	Innovative Visionary Intelligent Progressive Dynamic	Innovative Visionary Up To Date Progressive High Performance
Leader	Daring Leader Best Brand Dynamic Authentic	Best Brand Worth More High Performance Authentic Up To Date

Prestigious	Daring Upper Class Glamorous Prestigious Stylish	Upper Class Glamorous Distinctive High Quality Arrogant
Reliable	Dynamic Daring Rugged High Performance Stylish	Trustworthy Cares Customers Leader Worth More High Performance
Socially Responsible	Independent Socially Responsible Cares Customers Friendly Original	Independent Gaining In Popularity Cares Customers Friendly Kind
Straightforward	Traditional Original Down To Earth Trustworthy Good Value	Down To Earth Good Value Simple Trustworthy Cares Customers
Stylish	Stylish Glamorous Trendy Prestigious	Glamorous Prestigious Upper Class Distinctive

	Daring	High Quality
Trendy	Trendy Stylish Glamorous Arrogant Distinctive	Gaining In Popularity Unique Different Social Energetic
Trustworthy	Original Authentic Traditional Daring Best Brand	Reliable Cares Customers Good Value Leader Worth More
Unique	Daring Glamorous Distinctive Sensuous Prestigious	Different Daring Trendy Visionary Dynamic
Up To Date	Visionary Innovative Dynamic Daring Trendy	Innovative Progressive Visionary Intelligent Dynamic
Upper Class	Upper Class Glamorous Prestigious	Prestigious Glamorous Distinctive

	Stylish Daring	High Quality Arrogant
Visionary	Rugged Stylish Energetic High Performance Dynamic	Innovative Progressive Up to Date Intelligence Dynamic
Worth More	Daring Trendy Energetic Authentic Best Brand	High Quality Best Brand High Performance Leader Authentic

4.6 APPENDIX 2A: REGRESSION TABLE FOR STUDY 2 (CHAPTER 2)

Hierarchical Regression Results

	Dependent variable:			
	purchase intention			
	windfall only	windfall + hedonic	similarity only	all
	(1)	(2)	(3)	(4)
condition	-0.516***	0.750***		0.527***
	(0.119)	(0.179)		(0.176)
hedonic	0.276***	0.354***	0.266***	0.354***
	(0.013)	(0.021)	(0.012)	(0.021)
utilitarian	0.301***	0.392***	0.296***	0.392***
	(0.012)	(0.019)	(0.011)	(0.019)
condition × hedonic		-0.119***		-0.140***
		(0.026)		(0.026)
condition × utilitarian		-0.145***		-0.155***
		(0.024)		(0.024)
similarity × condition			0.204***	0.238***
			(0.016)	(0.016)
constant	0.753***	-0.078	0.274***	-0.078
	(0.116)	(0.146)	(0.087)	(0.144)
Observations	5,920	5,920	5,920	5,920
Log Likelihood	-11,333.510	-11,292.410	-11,265.480	-11,188.570
Akaike Inf. Crit.	22,679.030	22,600.810	22,542.960	22,395.140
Bayesian Inf. Crit.	22,719.150	22,654.300	22,583.080	22,455.310
Note:				*p<0.1; **p<0.05; ***p<0.01

4.7 APPENDIX 2B: REGRESSION TABLE FOR STUDY 3A (CHAPTER 2)

Hierarchical Regression Results

	Dependent variable:		
	purchase intention		
	windfall only	similarity only	both
	(1)	(2)	(3)
condition	-0.033		-0.191**
	(0.062)		(0.077)
similarity × condition			0.042***
			(0.013)
similarity		0.154***	0.132***
		(0.007)	(0.009)
constant	3.259***	2.685***	2.782***
	(0.044)	(0.038)	(0.055)
Observations	14,388	14,388	14,388
Log Likelihood	-29,646.950	-29,382.670	-29,382.630
Akaike Inf. Crit.	59,301.900	58,773.330	58,777.270
Bayesian Inf. Crit.	59,332.200	58,803.630	58,822.710
Note:			*p<0.1; **p<0.05; ***p<0.01

4.8 APPENDIX 2C: REGRESSION TABLE FOR STUDY 3B (CHAPTER 2)

Control vs. Windfall condition

	Dependent variable: purchase intention	
	hedonic HMPC	MA
	(1)	(2)
condition	-0.235	-1.167***
	(0.237)	(0.124)
hedonic	0.489***	0.448***
	(0.020)	(0.014)
utilitarian	0.175***	0.152***
	(0.024)	(0.016)
condition × hedonic	-0.023	
	(0.028)	
condition × utilitarian	-0.046	
	(0.033)	
similarity		0.244***
		(0.025)
constant	0.844***	1.164***
	(0.179)	(0.130)
Observations	3,980	3,980
Log Likelihood	-7,675.308	-7,628.711
Akaike Inf. Crit.	15,366.610	15,271.420
Bayesian Inf. Crit.	15,416.930	15,315.440
Note:	*p<0.1; **p<0.05; ***p<0.01	

Control vs. Prize condition

	Dependent variable:	
	purchase intention	
	hedonic HMPC	MA
	(1)	(2)
Condition.	-1.101***	-0.787***
	(0.233)	(0.118)
hedonic	0.490***	0.511***
	(0.021)	(0.015)
utilitarian	0.176***	0.214***
	(0.024)	(0.016)
condition × hedonic	0.096***	
	(0.029)	
condition × utilitarian	0.073**	
	(0.033)	
similarity		0.209***
		(0.025)
condition	0.835***	0.526***
	(0.177)	(0.128)
Observations	3,980	3,980
Log Likelihood	-7,729.490	-7,702.183
Akaike Inf. Crit.	15,474.980	15,418.370
Bayesian Inf. Crit.	15,525.290	15,462.390
Note:	*p<0.1; **p<0.05; ***p<0.01	

4.9 APPENDIX 2D: FIELD EXPERIMENT SETUP PHOTOS

(A) Spinning wheel to randomly assign coupon group.

(B) Cart of goods for sale in the mini shop.

(C) Setup of mini shop.

(D) Happy research assistant and experiment participant.

(E) Participant about to purchase goods.

(F) Research assistant explaining procedure to potential participants.