Journal of
Theoretical and Applied Electronic Commerce Research

MDPI

Article

Mapping Morality in Marketing: An Exploratory Study of Moral and Emotional Language in Online Advertising

Mauren S. Cardenas-Fontecha [1], Leonardo H. Talero-Sarmiento [2,*] and Diego A. Vasquez-Caballero [3]

[1] Marketing Technology, Universidad Autonoma de Bucaramanga UNABTec, Campus El Jardín, Avenida 42 No. 48-11, Barrio El Jardín, Bucaramanga 680003, Santander, Colombia; mcardenas894@unab.edu.co
[2] Industrial Engineering Program, School of Industrial and Business Studies, Universidad Industrial de Santander, Carrera 27 No. 9, Ciudad Universitaria, Bucaramanga 680002, Santander, Colombia
[3] Psychology Program, Faculty of Psychology, Universidad Cooperativa de Colombia, Campus Bucaramanga, Calle 30A No. 33-51, Bucaramanga 680002, Santander, Colombia; diego.vasquezca@campusucc.edu.co
* Correspondence: lehetasa@correo.uis.edu.co

Abstract

Understanding how moral and emotional language operates in paid social advertising is essential for evaluating persuasion and its ethical contours. We provide a descriptive map of Moral Foundations Theory (MFT) language in Meta ad copy (Facebook/Instagram) drawn from seven global beverage brands across eight English-speaking markets. Using the moralstrength toolkit, we implement a two-channel pipeline that combines an unsupervised semantic estimator (SIMON) with supervised classifiers, enforces a strict cross-channel consensus rule, and adds a non-overriding purity diagnostic to reduce attribute-based false positives. The corpus comprises 758 text units, of which only 25 ads (3.3%) exhibit strong consensus, indicating that much of the copy is either non-moral or linguistically ambiguous. Within this high-consensus subset, the distribution of moral cues varies systematically by brand and category, with loyalty, fairness, and purity emerging as the most prominent frames. A valence pass (VADER) indicates that moralized copy tends toward negative valence, yet it may still yield a constructive overall tone when advertisers follow a crisis–resolution structure in which high-intensity moral cues set the stakes while surrounding copy positions the brand as the solution. We caution that text-only models undercapture multimodal signaling and that platform policies and algorithmic recombination shape which moral cues appear in copy. Overall, the study demonstrates both the promise and the limits of current text-based MFT estimators for advertising: they support transparent, reproducible mapping of moral rhetoric, but future progress requires multimodal, domain-sensitive pipelines, policy-aware sampling, and (where available) impression/spend weighting to contextualize descriptive labels.

Keywords: moral foundations theory; digital advertising; Meta Ads Library; advertising copy; platform governance; sentiment analysis (VADER); SIMON; multimodal persuasion; alcohol marketing; computational text analysis

Check for updates

Academic Editors: Jolanta Kowal and Paweł Weichbroth

Received: 27 November 2025
Revised: 27 December 2025
Accepted: 12 January 2026
Published: 14 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

This paper examines moral language in Meta ad copy—the standardized textual fields (primary text, headline, description) that accompany images and videos in paid placements on Facebook and Instagram. Marketing has long leveraged affect to shape judgment and capture attention [1–3]; neuroimaging shows reward- and emotion-related activation during consumer choice, consistent with hot–cold empathy gaps [4] and with the vulnerability

that bounded rationality creates for strategically crafted appeals [5]. Beyond generic affect, moral language is distinctive: people routinely ascribe moral value to marketplace choices, and intuitionist accounts—rapid, affect-laden judgments followed by post hoc reasoning—explain why moral cues can be consequential [6]. At the same time, prior work indicates that moral language is relatively uncommon in everyday communication, appearing in only a small fraction of conversations and social-media posts [7]. This scarcity underscores both the analytical challenge and the potential informativeness of detecting moral cues when they do appear in commercial contexts.

Moral Foundations Theory (MFT) provides a tractable framework for categorizing moral content into care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, and purity/degradation [8]. In an intuitionist view, moral judgments are not primarily the product of deliberation but arise from fast, affective appraisals that guide judgments of right and wrong [6,9]. In marketing contexts, evidence links morally charged communication to shifts in consumer responses, particularly when appeals align with audience moral identities [10–13]. We adopt MFT in an explicitly exploratory manner, using it as an organizing framework to describe patterns in advertising copy rather than to infer causal effects or persuasive effectiveness. We use MFT pragmatically for three reasons. First, it provides operationally distinct and lexicon-mappable dimensions that have been widely used in consumer research [6,9]. Second, it generates audience-congruence predictions that have been empirically validated in marketing and prosocial behavior contexts [10,11]. Third, it aligns well with existing computational tools designed for large-scale text analysis [14].

We choose to study ad copy because it constitutes a standardized, comparable, and explicitly optimized textual component of advertising. Ad copy is typically brief and carefully crafted to convey core persuasive cues under tight platform constraints, making it directly comparable across brands and product categories. Prior research shows that verbal information in advertising plays a crucial role in clarifying product attributes and guiding consumer inferences [15], even when other elements are present. Moreover, textual elements, once attended to, tend to receive deeper processing and support more specific interpretations [16]. From a measurement standpoint, copy fields are easily extractable, platform-defined, and systematically optimized by advertisers, making them a tractable and analytically consistent unit for large-scale text analysis.

Studying morality is important in this context because platform governance further shapes moral expression. Meta's review and policy systems restrict violent, discriminatory, or politically sensitive content [17]. Simultaneously, the ad-delivery stack encourages multiple variants of primary text, headlines, and descriptions that are algorithmically recombined with different creatives to optimize outcomes [18], and text quality predicts click-through and related metrics [19]. In practice, marketers tune language to "fit" algorithms and policies, embedding implicit moral cues that preserve resonance without triggering filters [20]. Under these constraints, explicit moral language may be attenuated, while subtler or implicit moral framing can persist—making its detection analytically demanding but potentially informative.

Against this backdrop, we provide a descriptive map of moral and emotional language in Meta ad copy using MFT as the organizing framework. We assemble a cross-market corpus from the Meta Ad Library, focusing on leading global firms selected from established rankings with transparent, annually updated methodologies, and examine categories with heightened public-health and sustainability salience—alcohol and tobacco—given their regulatory scrutiny, responsibility messaging, and alignment with SDG 3 and SDG 12 [21,22]. Consequently, the study is designed as a pilot and hypothesis-generating exercise rather than a definitive assessment of brand-level strategies. Our analysis addresses three primary quantitative research questions: (RQ1) How are moral foundations distributed in Meta ad

copy across brands? (RQ2) How does foundation usage differ within high-salience categories (alcoholic vs. non-alcoholic)? (RQ3) To what extent do moral and sentiment-laden terms co-occur?

To make the complex labeling mechanics transparent, we supplement this corpus-level analysis with a fourth, qualitative objective: to conduct an expert-led "deep dive" into a representative subset of ad exemplars, illustrating the interplay of the unsupervised and supervised channels and the application of the consensus rules.

We conduct a descriptive (non-causal) analysis of textual ad fields; moral signaling conveyed visually or auditorily is undercaptured. Platform policies likely suppress explicit moral language, biasing distributions toward implicit cues. Algorithmic recombination of copy–creative variants complicates precise alignment between text and served impressions. Text embedded in images, sarcasm/irony, and cross-language nuances may be missed or misclassified. Sampling leading global firms and focusing on alcohol/tobacco limits representativeness but targets economically salient actors under strong public-interest scrutiny; the covered time window further constrains generalizability. To preserve cohesion with the rest of the manuscript, Materials and Methods details the Meta Ads Library data source (Graph API v23.0), brand and market sampling, text normalization, and the moral-language estimation pipeline (unsupervised SIMON and supervised classifiers, consensus rules, confidence index, and purity diagnostics). Results reports foundation-level distributions, co-occurrence patterns of moral and emotional terms, and sensitivity checks—without inferring causal effectiveness. Discussion interprets these descriptive patterns in light of multimodal creative, policy filtering, and algorithmic optimization, articulates domain-specific limitations (including ongoing debates around MFT dimensionality and discriminant validity [9]) and outlines implications for future multimodal and multilingual work. Conclusions summarizes the main takeaways and delineates avenues for subsequent research.

2. Materials and Methods

2.1. Data Source and Sampling Frame

We assembled a multi-market corpus of brand advertising messages retrieved from Meta's Ads Library via the Graph API (v23.0). The corpus covers seven global beverage brands—Coca-Cola, Pepsi, Red Bull, Heineken, Hennessy, Nescafé, and Corona—across eight English-speaking markets: the United States, Great Britain (GB), Ireland (IE), Canada, Australia, New Zealand, South Africa, and Singapore. The analysis is intentionally restricted to English-language ad copy because the `moralstrength` toolkit and its underlying moral dictionaries are validated for English text only [14]. Extending the pipeline to other languages would require automated translation, which risks distorting culturally embedded moral cues and introducing systematic bias in moral foundation assignment [23]. To prioritize internal validity and interpretability of the moral-language estimates, we therefore limited the corpus to English-speaking markets and treat multilingual extension as a priority for future work.

Brand selection followed the 2025 *Forbes* Global 2000 list (Food, Drink, and Tobacco https://www.forbes.com/global2000/list/, accessed on 15 October 2025) and focused on multinational beverage firms with high consumer visibility and policy relevance. From a marketing perspective, relying on the Forbes Global 2000 provides a transparent and externally validated benchmark of brand prominence, ensuring that the analysis targets firms with substantial brand equity and symbolic influence [24,25]. Prior research shows that highly salient, widely recognized brands play a disproportionate role in shaping category meaning, consumer associations, and normative expectations, making them appropriate focal actors for examining communicative strategies such as moral framing [26].

By concentrating on leading global brands, we reduce noise from idiosyncratic or local positioning strategies and increase the likelihood that observed moral-language patterns reflect systematic tendencies rather than firm-specific anomalies. For RQ2, we contrasted alcoholic (Heineken, Hennessy, Corona) versus non-alcoholic (Coca-Cola, Pepsi, Red Bull, Nescafé) categories given their differing regulatory salience and responsibility messaging.

Because disclosure rules vary by country, harvesting proceeded in two stages. First, in GB and IE—where comprehensive archive search is available—we discovered official brand pages using disjunctive queries (for example, "Coca-Cola" OR "Coke") and retained candidates containing brand tokens, bounded by a minimum delivery date of 1 January 2021. We appended a curated list of verified page IDs to mitigate recall gaps. Second, we projected discovered IDs to the other English-language markets and retrieved ads with a uniform field set: unique identifiers; creation and delivery timestamps; text fields and creative snapshot URL; language and platform metadata; impression and spend bounds (lower/upper/average where available); demographic and geographic delivery aggregates; and bylines.

For cross-country consistency, we respected the API constraint on `ad_type`: `ALL` in GB/IE and `POLITICAL_AND_ISSUE_ADS` elsewhere, with graceful fallback if unsupported. Because this can bias the pool toward issue ads outside GB/IE, we present GB+IE as a commercial benchmark and either stratify or cautiously pool other markets. The primary analysis window is 1 January 2021 through 15 October 2025; legacy creatives earlier than 2021 are included only when required by archive constraints and are flagged in robustness checks.

Ads with multiple text bodies were exploded so that each body constituted a distinct observation linked to its parent ad. The resulting analysis variable underwent Unicode NFKC folding and whitespace compaction. The final dataset contains $N = 758$ text bodies and serves as the basis for all quantitative analyses (RQ1–RQ3). Data are public, aggregated, and non-personal; no sensitive attributes were accessed or inferred.

2.2. Depuration and Normalization

Prior to each run we removed residual outputs from previous executions to prevent contamination. Normalization preserved lexical content while harmonizing diacritics and spacing (for example, "Nescafé" → "Nescafe"). We retained analytic variables (normalized text, identifiers, timestamps, delivery metadata) and deduplicated at both the parent-ad and text-body levels.

2.3. Moral Language Estimation

To ensure robust classification of moral content, we implemented a dual-estimator framework utilizing the `moralstrength` library [14]. This approach combines unsupervised semantic proximity with supervised lexical mapping to triangulate moral signals in short-form advertising text.

2.3.1. Unsupervised Estimator (SIMON)

The unsupervised channel (SIMON) quantifies the semantic similarity between the document vector $\mathbf{v}_t$ and the foundation-specific vector $\mathbf{v}_f$ using cosine similarity:

$$C_f(t) = \frac{\mathbf{v}_t \cdot \mathbf{v}_f}{\|\mathbf{v}_t\| \, \|\mathbf{v}_f\|}, \quad f \in \{\text{care}, \text{fairness}, \text{loyalty}, \text{authority}, \text{purity}\}, \tag{1}$$

where $C_f(t) \in [-1, 1]$. To facilitate linear comparison across models, raw similarities are rescaled to a decile metric:

$$U_f(t) = 5\big(C_f(t) + 1\big) \in [0, 10]. \tag{2}$$

the model identifies the primary foundation (top_SIMON) and measures classification confidence through the winning margin Δ_U,

$$\text{sim_max}(t) = \max_f U_f(t), \qquad \Delta_U(t) = U_{(1)}(t) - U_{(2)}(t). \tag{3}$$

and records the argmax label $\text{top_SIMON}(t) = \arg\max_f U_f(t)$.

2.3.2. Supervised Estimator

The supervised channel utilizes a lexicon-based unigram estimator to map foundation-specific evidence to calibrated probabilities $S_f(t) \in [0, 1]$. Using the primary model m^*, we extract peak probabilities and margin scores:

$$\text{sup_max}(t) = \max_f S_f^{(m^*)}(t), \qquad \Delta_S(t) = S_{(1)}^{(m^*)}(t) - S_{(2)}^{(m^*)}(t), \tag{4}$$

recording the argmax label $\text{top_SUP}(t)$. These scores serve as high-precision filters for ranking and diagnostic purposes.

2.3.3. Corpus-Wide Classification and Error Awareness

We applied this dual-modeling approach to the full corpus ($N = 758$) to generate continuous foundation scores and top-label candidates. This comprehensive mapping addresses brand-level distributions (RQ1), category-specific contrasts (RQ2), and moral–sentiment associations (RQ3).

However, short-form advertising introduces lexical noise; promotional jargon (e.g., slogans, calls to action) may mimic moral foundations without expressing moral intent [27]. To mitigate the risk of false positives, we utilize a two-tier analytic strategy: (1) descriptive mapping of the full corpus to maintain scope, and (2) a high-precision depuration process to isolate a smaller set of high-consensus moral signals.

2.4. Consensus Rules and Label Assignment

We define a high-confidence subset based on strict inter-model agreement to maximize the precision of the assigned moral label. Using a priori thresholds ($\tau_S = 0.75$, $\tau_U = 7.5$), we identify observations for which both channels converge on the same moral foundation.

$$I_{\text{cons}}(t) = \begin{cases} 1, & \text{sup_max}(t) > \tau_S \wedge \text{sim_max}(t) > \tau_U \wedge \text{top_SUP}(t) = \text{top_SIMON}(t), \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

We quantify classification ambiguity (mixed_flag) based on narrow winning margins ($\varepsilon = 0.05$):

$$\text{mixed_flag}(t) = \mathbb{1}[\Delta_S(t) < \varepsilon \vee \Delta_U^{\text{norm}}(t) < \varepsilon]. \tag{6}$$

A normalized confidence index $c(t)$ integrates both channels into a single metric:

$$\text{sim_max}^{\text{norm}}(t) = \text{sim_max}(t)/10, \qquad c(t) = \tfrac{1}{2}\big(\text{sup_max}(t) + \text{sim_max}^{\text{norm}}(t)\big). \tag{7}$$

2.5. Purity Diagnostic

To distinguish moral purity appeals from product formulation claims (e.g., "clean," "pure ingredients"), we implemented a diagnostic flag:

$$J_{\text{purity}}(t) = \mathbb{1}\big[\text{top_SUP}(t) = \text{purity} \wedge U_{\text{purity}}(t)/10 < 0.60 \wedge \mathcal{R}(t) = 1\big], \tag{8}$$

where $\mathcal{R}(t)$ identifies marketing-specific keywords. This ensures that physical product descriptions are not conflated with moral signaling.

2.6. Moral–Sentiment Co-Occurrence (RQ3)

Sentiment analysis was performed using VADER [28], selected for its sensitivity to short-form and informal advertising text. For each text t, VADER generates polarity scores and a normalized compound score $s_{\text{compound}}(t) \in [-1, 1]$. These values are evaluated against moral estimates using non-parametric association measures.

2.7. Analytical Strategy

Our descriptive analysis proceeds in three stages: (1) tabulating foundation distributions across brands and markets (RQ1), (2) performing category-level contrasts between alcoholic and non-alcoholic segments (RQ2), and (3) assessing moral–sentiment alignments via Spearman rank correlations between foundation scores and VADER metrics. To maintain statistical rigor, multiple comparisons are corrected using the Benjamini–Hochberg procedure ($\alpha = 0.05$) [29]. Findings are reported for both the full corpus ($N = 758$) and the high-precision consensus subset ($I_{\text{cons}} = 1$).

2.8. Consensus-Eligible Subset (N = 12)

From the full corpus, we removed items for which the channels disagreed on the top moral foundation. The remaining cases exhibited inter-model agreement and were partitioned into Strict Consensus (Equation (5)) and Low-Confidence Agreement (LCA). LCA refers to cases in which the supervised and unsupervised channels agree on the top moral foundation but do not meet the strict confidence thresholds specified in Equation (5) and/or are flagged as ambiguous under Equation (6).

During this curation step, we further excluded cases that, despite differences in scope or metadata, corresponded in practice to the same underlying ad copy (e.g., identical text deployed across markets or delivery variants). These removals avoided redundant qualitative interpretation and ensured that each exemplar represented a substantively distinct moral narrative. As a result, manual review identified near-duplicates, and a subject-matter expert curated a final, balanced, consensus-eligible subset for qualitative analysis. Figure 1 illustrates the funnel process from the full corpus used to address the three research questions to the consensus-eligible subset used for qualitative analysis.

2.9. Validation, Governance, and Reproducibility

HTTP requests handled 429/5xx responses with exponential backoff; pagination continued until exhaustion or a configured limit. We persisted intermediate artifacts (discovered page IDs and the final ad table). Each run began by dropping prior unsupervised outputs, model-suffixed columns, and derived flags to prevent leakage across experiments. All thresholds ($\tau_S = 0.75$, $\tau_U = 7.5$, $\varepsilon = 0.05$, $U_{\text{purity}}/10 < 0.60$) were set a priori and varied in sensitivity analyses. Analytic code recorded software versions and random seeds; results are reproducible from the provided scripts and configuration files. Grammar tools such as Grammarly and Generative AI assisted only with language polishing and document structure, and were not used for data collection, modeling, or statistical estimation.

2.10. Ethics and Data Availability

The study uses public, aggregate advertising data and does not involve human subject research as defined by institutional policy; no personal data were processed. Data-collection scripts and derived data sufficient for replication are provided, subject to platform terms, together with instructions for recreating the corpus from the Ads Library.

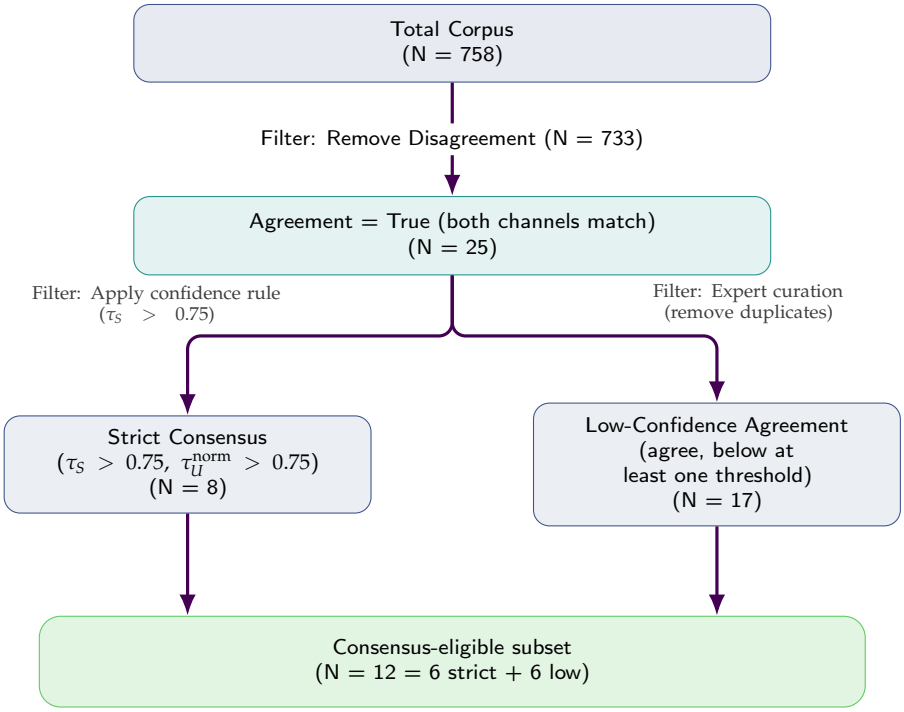


Figure 1. Flowchart of the data selection process, filtering from the total corpus to the consensus-eligible subset.

3. Results

3.1. Model Calibration and Consensus

To ensure the reliability of moral labels generated by our dual-channel pipeline, we implemented a conservative selection funnel based on agreement, confidence, and ambiguity metrics (Figures 2–4). The distribution of the full corpus ($N = 758$) across the confidence–confidence plane is visualized in Figure 2. The upper-right quadrant isolates the consensus-eligible subset, where high-confidence convergence occurs between supervised probabilities ($S_{max} > 0.75$) and unsupervised semantic similarities ($U_{norm} > 0.75$). Observations falling outside this quadrant or marked as model disagreements underscore the significant lexical noise inherent in short-form advertising copy, justifying a two-tier depuration strategy to maintain high-precision moral signals.

Figure 3 quantifies the resulting classification outcomes. The vast majority of the corpus ($N = 733, 96.7\%$) results in channel disagreement, where the two models differ in their primary foundation assignment. Within the agreement cases ($N = 25, 3.3\%$), we distinguish between the high-precision Consensus Met core ($N = 8, 1.1\%$) and the Agree (Low Confidence) group ($N = 17, 2.2\%$), which converges on a label but fails one or both calibrated thresholds.

Finally, Figure 4 diagnoses classification ambiguity through winning margins (Δ). The analysis reveals that while the supervised lexical model (left, Δ_S) is highly decisive, the unsupervised semantic model (right, Δ_U^{norm}) acts as the primary driver of ambiguity due to a high frequency of near-ties ($\Delta < 0.05$). This divergence validates our requirement for dual-model agreement and stringent confidence markers. At the same time, the Consensus Met

subset provides the most reliable foundation for qualitative inspection; the low-confidence agreed cases serve only as secondary indicators for limited robustness checks.

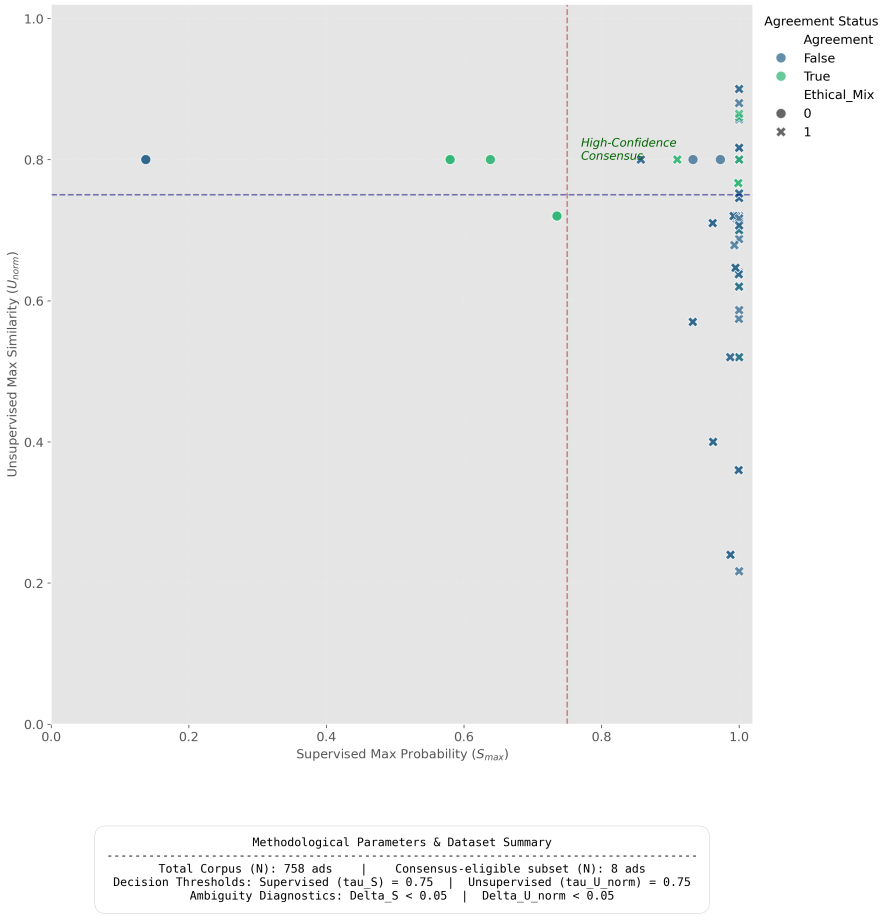


Figure 2. Model calibration and consensus quadrant selection ($N = 758$). Technical parameters and resulting subset sizes are summarized in the lower panel.

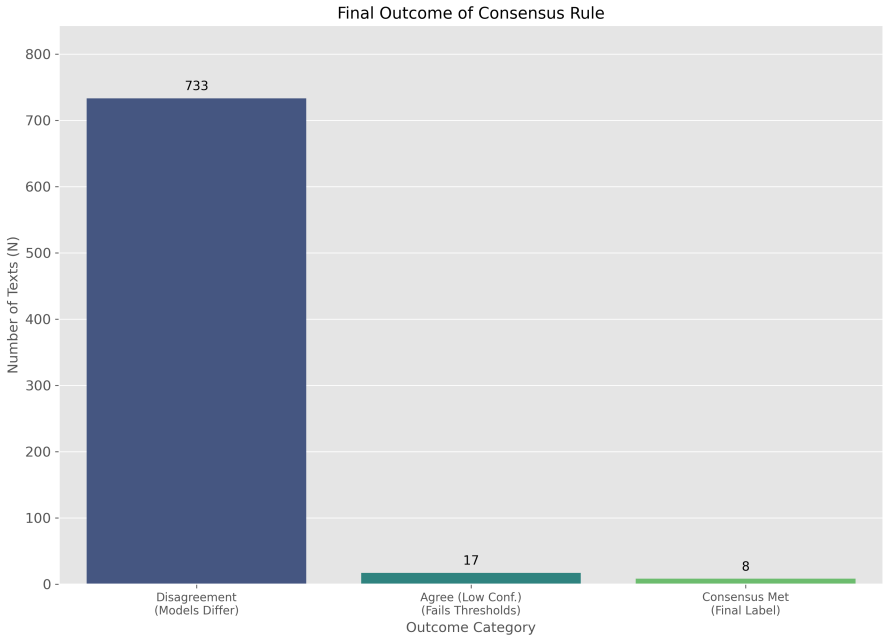


Figure 3. Final outcome distribution of the consensus rule ($N = 758$). The chart quantifies the transition from the total corpus to the isolated high-precision subsets.

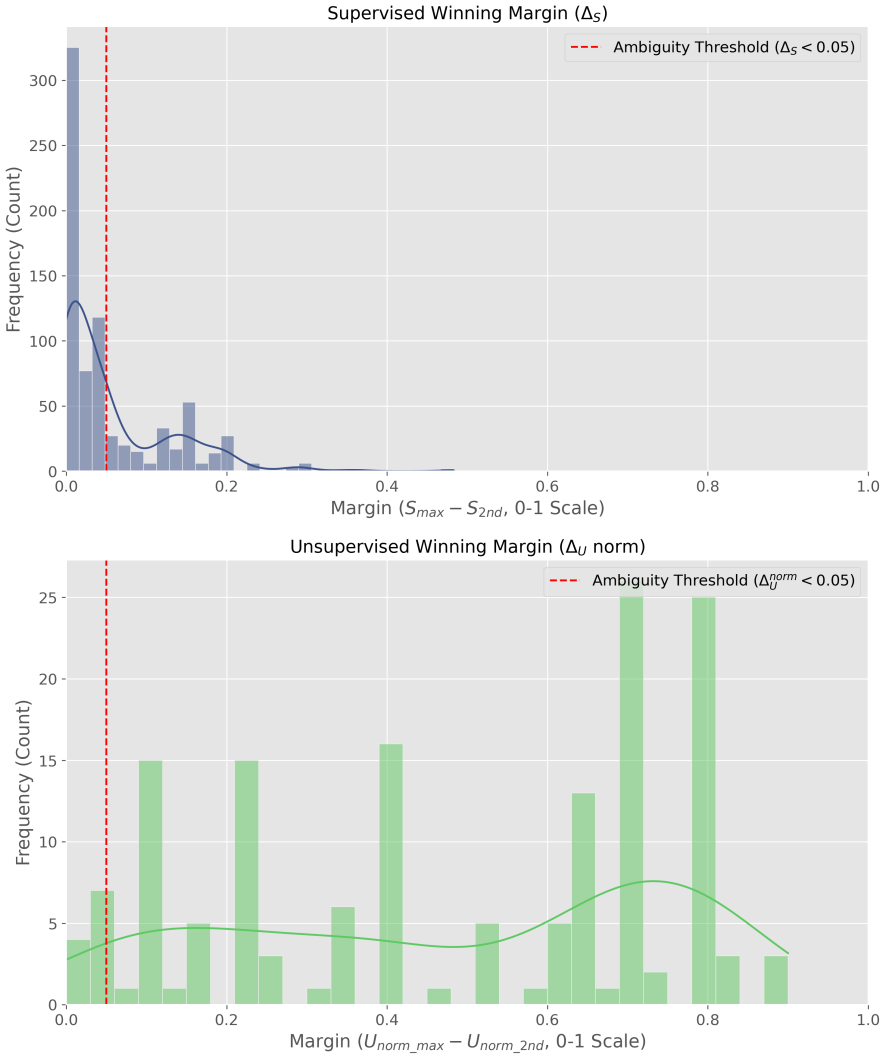


Figure 4. Distribution of winning margins (Δ) across channels ($N = 758$). The supervised channel is generally decisive, while the unsupervised channel exhibits a higher concentration of near-ties below the 0.05 threshold.

3.2. RQ1: Distribution of Moral Foundations in the Full Corpus (N = 758)

We first analyze the distribution of final moral labels across the full 758-item corpus. The heatmap in Figure 5 and the detailed percentages in Table 1 reveal distinct brand-level strategies.

Table 1. Moral Foundation Distribution by Brand in the Full Corpus.

Moral Label Final Lower Brand	Care	Fairness	Loyalty	Authority	Purity	Non-Moral	Mixed
Coca-Cola	0.76	15.27	8.40	3.05	0.76	4.58	67.18
Heineken	0.84	2.52	15.13	1.68	9.24	7.56	63.03
Hennessy	0.00	0.00	53.85	0.00	3.85	3.85	38.46
Pepsi	0.00	40.00	0.00	0.00	0.00	12.50	47.50
Red Bull	0.00	6.45	3.23	0.00	3.23	6.45	80.65
Unknown	0.26	12.11	8.68	1.32	2.37	23.42	51.84

Note. Share of Ad Copy % from N = 758.

Brands show clear preferences: Hennessy's messaging is overwhelmingly dominated by loyalty (53.85%), and Pepsi's by fairness (40.00%). Other brands show a more diverse moral palette; Coca-Cola utilizes both loyalty (20.61%) and fairness (15.27%), while

Heineken also leans on loyalty (16.81%) but adds significant use of purity (9.24%). The "Unknown" category, which contains ads where a brand could not be inferred, has the largest share of non-moral text (23.42%), as expected for ads lacking clear brand identifiers.

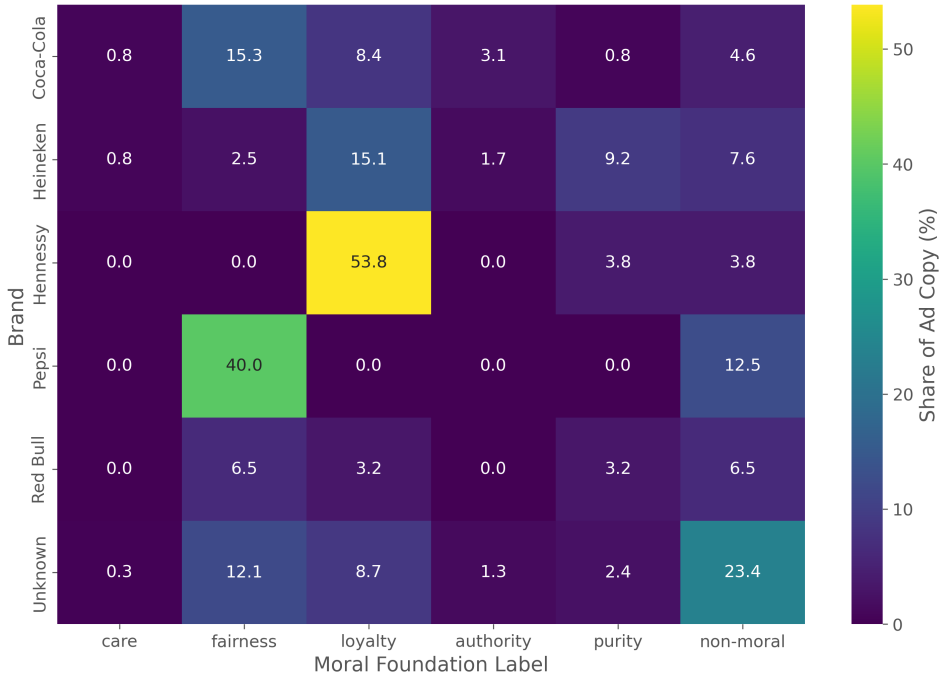


Figure 5. Heatmap of moral foundation distribution (share of ad copy %) by brand for the full corpus. Brands exhibit distinct moral foundation preferences.

3.3. RQ2: Category Contrast (Alcoholic vs. Non-Alcoholic) N = 758

Aggregating brands into 'Alcoholic' and 'Non-Alcoholic' categories reveals a clear divergence in moral framing, as shown in Figure 6 and Table 2. A stark contrast emerges: the Alcoholic category (e.g., Heineken, Hennessy) most frequently employs loyalty (23.45%) and purity (8.28%). This suggests a focus on in-group identity, heritage, and the quality or integrity of the product.

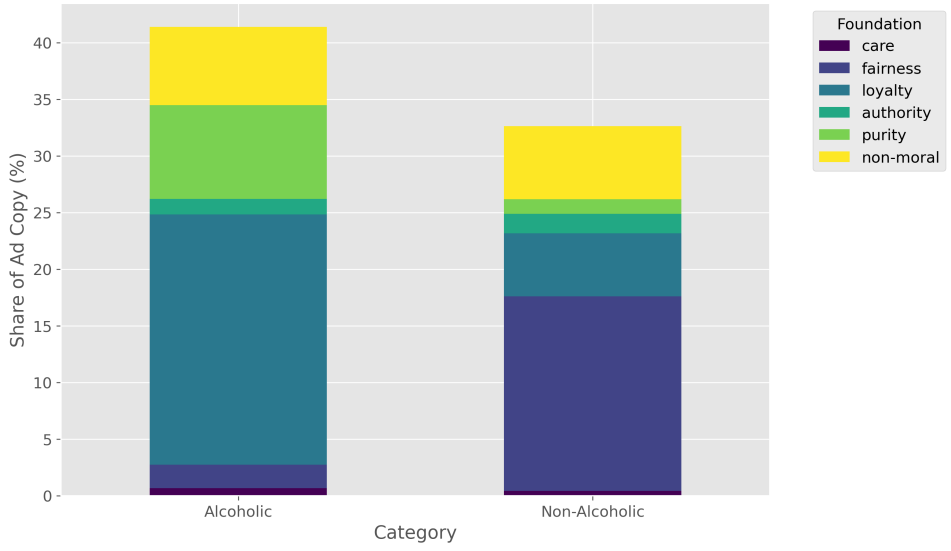


Figure 6. A stacked bar chart comparing the relative share of moral foundations used in ad copy for Alcoholic vs. Non-Alcoholic beverage categories. The 'mixed' category is excluded.

Table 2. Moral Foundation Distribution by Beverage Category (Share of Ad Copy %).

Moral Label Final Lower Category	Care	Fairness	Loyalty	Authority	Purity	Non-Moral	Mixed
Alcoholic	0.69	2.07	22.07	1.38	8.28	6.90	58.62
Non-Alcoholic	0.43	17.17	5.58	1.72	1.29	6.44	67.38

Conversely, the Non-Alcoholic category (e.g., Coca-Cola, Pepsi) heavily favors fairness (17.17%) and, to a lesser extent, loyalty (12.45%). This framing aligns with messages of social justice, equality, and community engagement, which are common themes in their corporate social responsibility campaigns.

3.4. RQ3: Moral–Sentiment Co-Occurrence N = 758

To answer RQ3, we assessed the co-occurrence of moral language with emotional sentiment. An initial pass using the NRC Emotion Lexicon (nrclex, [30]) yielded no matches, indicating its lexicon has poor coverage for our ad copy corpus. As a robust alternative, we used VADER to calculate sentiment polarity scores and computed the Spearman rank correlation (ρ) between these scores and our supervised moral foundation scores.

The results (Figure 7) reveal a distinct pattern of affective intensity. All five moral foundations show a moderate, positive correlation with VADER's compound sentiment score (e.g., Care $\rho = 0.38$; Authority $\rho = 0.36$). This suggests that moralized text is significantly more emotionally charged than non-moral text.

However, a decomposition of the sentiment channels reveals that this intensity is primarily driven by negative vocabulary. Contrary to expectation, the correlation with the negative sentiment score (e.g., Purity $\rho = 0.27$; Loyalty $\rho = 0.24$) is consistently stronger than the correlation with the positive score (e.g., Purity $\rho = 0.00$; Loyalty $\rho = 0.05$). This indicates that while the overall message (Compound) may be persuasive or constructive, the specific moral vocabulary employed is frequently rooted in the description of violations, problems, or threats.

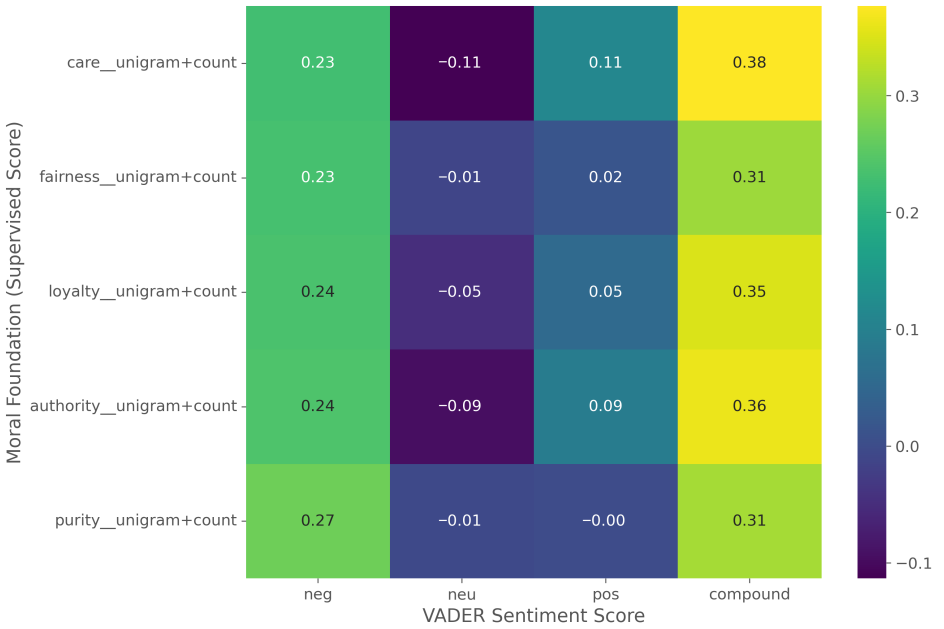


Figure 7. A correlation matrix (heatmap) showing the Spearman rho coefficient between supervised moral foundation scores and VADER sentiment polarity scores.

3.5. Illustrative Exemplars from the Consensus-Eligible Subset N = 12

To evaluate the operational validity of the labeling pipeline, we selected a target subset of twelve advertisements (t1–t12) for qualitative inspection (Tables 3 and 4). This selection represents the high-precision segment of the corpus, spanning both alcoholic and non-alcoholic categories across various creative styles, including promotional copy, corporate updates, and social responsibility messaging.

Table 3. Low-Confidence Agreement (LCA) advertisements (t1–t7) from the final curated set.

ID	Approach	Brand	Ad Copy	Moral Label
t1	LCA	Heineken	As we continue to celebrate women in our business in recognition of International Women's Day we wanted to call out one special team at HEINEKEN UK who've been working tirelessly to make sure our ciders & beers are available for our customers online.	loyalty
t2	LCA	Heineken	From working at HEINEKEN by week to dancing their way through Manchester by weekend! Our colleagues showed up for the LGBTQIA+ community at @manchesterpride and brought the joy, the love, and some incredible dance moves with them! Brewing With Love.	loyalty
t3	LCA	Coca-Cola	You all know what time it is! This week's Digest is landing early with the latest highlights from the world of CCEP, including thoughts from our CEO on an important milestone, highlights from our Ramadan activations, and an escape to a digital world for football fans. Curious to learn more? Click the Weekly Digest News link in our bio for the inside scoop! We Are CCEP Innovation Immersive Marketing.	loyalty
t4	LCA	Coca-Cola	Enjoyed together since 1900. Try the new Bacardi & Coke premixed drink in a can.	loyalty
t5	LCA	Coca-Cola	Add a Coca-Cola to your order for a chance to win free flights across Europe! * Coca-Cola Tasty Celebrations.	authority
t6	LCA	Coca-Cola	Order now on Deliveroo.	authority
t7	LCA	Heineken	We're proud to continue our commitment to innovation and delivering exceptional drinks for our consumers. Please welcome the latest addition to the Old Mout family: Old Mout Cider Cocktails! Our new cider cocktails are a bold fusion of crisp, refreshing Old Mout cider and premium spirits. This launch reflects our ongoing dedication to expanding choice and elevating taste experiences, offering adventurous drinkers two vibrant new flavours—vodka with passionfruit and lime, and gin with raspberry and rhubarb. It's more than just cider; it's the spirit of innovation in every sip. Cheers to exploring new horizons together!	mixed

* LCA items agree on the assigned moral label but do not meet one or both high-confidence thresholds defined in Equation (5).

Table 4. Extrict advertisements (t8–t12) from the final curated set.

ID	Approach	Brand	Ad Copy	Moral Label
t8	Extrict	Coca-Cola	Step 1: Find a recycling machine at a Merlin Attraction. Step 2: Recycle your 500 mL plastic bottle during your visit. Step 3: Scan the QR code or visit the link on your device for a chance to WIN big! Good luck! UK/ROI 18+. Closes 07.09.25. Retain printed receipt for proof of entry. T&Cs apply, see https://merlinmagic.biz/coketerms/ (accessed on 15 October 2025). * Recycling machines are located inside participating attractions; entry ticket required.	purity
t9	Extrict	Heineken	Claim your free pint of Heineken® 0.0 today. "Redemption never tasted this good."	purity
t10	Extrict	Heineken	Is it still home advantage if you're on clean-up duty? Heineken 0.0. Great taste. Zero alcohol.	purity

Table 4. *Cont.*

id	Approach	Brand	Ad Copy	Moral Label
t11	Extrict	Heineken	Today, we celebrated how far we've come as a company with our commitment to fostering an open and inclusive workplace that values the contribution of every colleague. It was amazing to host our first Diversity, Equity, and Inclusion Townhall, "Brewing Inclusion," in Brighton just before the Pride weekend! We'll keep pushing forward together with our Colleague Networks to promote a culture of belonging where everyone feels safe, included, and valued for who they are. Happy Pride! Brew a Better World—Diversity, Equity, Inclusion, Pride.	care
t12	Extrict	Coca-Cola	To our colleagues in the Netherlands: our manufacturing site in Dongen has been awarded a platinum certification—the highest level—for sustainable water stewardship by the global Alliance for Water Stewardship (AWS), a global membership collaboration that drives, recognises, and rewards good water stewardship performance. The best bit? This is the first AWS certification for any site in the Netherlands. The site has worked hard over the years to optimise water usage, collaborating with local partners and adapting its bottle-washing line to reuse only clean, reclaimed water. These measures have improved water efficiency and contribute to protecting the health of the watershed and local ecosystem. Water is essential across our value chain, and we must treat it with the care it deserves. Good water stewardship is a key part of our sustainability strategy, focused on reducing water consumption and protecting local sources for future generations. The Dongen site is also part of a CCEP programme aiming for at least six sites to become carbon-neutral certified according to PAS 2060 by the end of 2023—a key part of our Net Zero 2040 ambition.	care

* The label Extrict identifies ad copies that satisfy the high-confidence rule defined in Equation (5).

For each exemplar, we extracted foundation scores from the unsupervised SIMON model and the supervised classifiers (Table 5). Final labels were assigned according to the dual-channel consensus logic defined in Equation (5). Table 6 summarizes the specific metrics and classification decisions for this subset.

Table 5. Moral foundation scores by channel for each advertisement (t1–t12) using the unsupervised SIMON (0–10) and supervised SUP (0–1) models.

		Text ID											
Foundation	**Metric**	**t1**	**t2**	**t3**	**t4**	**t5**	**t6**	**t7**	**t8**	**t9**	**t10**	**t11**	**t12**
Care	SIMON (0–10)	–	–	–	–	–	–	–	–	–	–	8.60	8.65
	SUP (0–1)	1.000	1.000	1.000	0.342	0.277	0.350	1.000	1.000	0.594	0.998	1.000	1.000
Fairness	SIMON (0–10)	–	–	–	–	–	–	–	–	–	–	4.21	4.50
	SUP (0–1)	1.000	1.000	1.000	0.375	0.080	0.062	1.000	1.000	0.881	0.991	1.000	1.000
Loyalty	SIMON (0–10)	5.20	6.20	6.38	7.20	–	–	7.00	–	–	–	6.80	5.40
	SUP (0–1)	1.000	1.000	1.000	0.735	0.194	0.173	1.000	1.000	0.207	0.804	1.000	1.000
Authority	SIMON (0–10)	–	–	–	–	8.00	8.00	6.40	–	–	7.17	–	–
	SUP (0–1)	1.000	0.999	1.000	0.196	0.638	0.580	1.000	0.998	0.464	0.994	1.000	1.000
Purity	SIMON (0–10)	–	–	–	–	–	–	–	8.00	8.00	7.67	–	7.83
	SUP (0–1)	1.000	1.000	1.000	0.581	0.226	0.223	1.000	1.000	0.910	0.999	1.000	1.000

Notes: SIMON is an unsupervised lexical model scored on a 0–10 scale; SUP is a supervised unigram+count model scored on a 0–1 scale. Dashes (–) indicate values not applicable or not detected.

Table 6. Summary of the label-decision metrics for each advertisement t_i based on Equation (5). All exemplars (t_1–t_{12}) are shown.

ID	sup_max	Δ_S	sim_max	Δ_U	$\min(\Delta_S, \Delta_U)$	Moral Label
t1	1.000	0.000	5.200	5.200	0.000	loyalty
t2	1.000	0.000	6.200	6.200	0.000	loyalty
t3	1.000	0.000	6.375	6.375	0.000	loyalty
t4	0.735	0.154	7.200	7.200	0.154	loyalty
t5	0.638	0.361	8.000	8.000	0.361	authority
t6	0.580	0.229	8.000	8.000	0.229	authority
t7	1.000	0.000	7.000	0.600	0.000	mixed
t8	1.000	0.000	8.000	8.000	0.000	purity
t9	0.910	0.029	8.000	8.000	0.029	purity
t10	0.999	0.001	7.667	0.500	0.001	purity
t11	1.000	0.000	8.600	1.800	0.000	care
t12	1.000	0.000	8.648	0.814	0.000	care

Within this illustrative set, all observations satisfied the high-confidence consensus criteria. The resulting distribution identifies loyalty and purity as the primary foundations, with care emerging in specialized contexts (t11, t12).

Crucially, this subset demonstrates the robustness of the two-channel approach over single-model estimators. For instance, several cases (e.g., t1, t3, t4) exhibit supervised ties ($\Delta_S = 0.000$) where a standalone lexical classifier would fail to provide a decisive label. In these instances, the strong cross-channel consensus—where the unsupervised semantic matcher provides a clear tie-break—takes precedence, allowing for stable classification. This validates the depuration strategy as a means of ensuring interpretive reliability in the presence of lexical ambiguity.

4. Discussion

This study offered a descriptive map of moral language in digital advertising by applying a two-channel MFT pipeline to Meta ad copy. In line with intuitionist accounts that emphasize fast, affect-laden appraisals [8,9,31], we treated foundations as operationally distinct, lexicon-mappable dimensions [14] and refrained from causal claims. Our analyses addressed three questions—brand/market distributions (RQ1), category contrasts (RQ2), and co-occurrence/alignment patterns (RQ3)—and surfaced domain-specific challenges tied to multimodality, platform governance, and algorithmic optimization. As outlined before, this study is intentionally exploratory and descriptive, and all interpretations should be read in light of platform filtering, multimodality, and conservative labeling thresholds.

4.1. Interpreting Brand and Category Strategies

Moral language is not a uniform background feature of advertising copy; it is strategically deployed and patterned by brand identity and product category. Given strong platform constraints, this research design aimed to establish whether moral framing is detectable at all, rather than to estimate its prevalence or persuasive effectiveness. Accordingly, we primarily relied on the baseline classification of the full corpus (N = 758), as shown in Table 1, to interpret brand- and category-level strategies.

At the brand level, stark differences emerged. For example, Hennessy concentrated on loyalty (53.8%), consistent with heritage-based and in-group positioning, whereas Pepsi emphasized fairness (40.0%), aligning with community-oriented and social-good narratives. This pattern accords with prior evidence that persuasion improves when moral cues align with audience identities or ideological profiles [10,11], as well as with Moral Foundations Theory's prediction that communities differentially weight moral "channels" [8,31].

Notably, Coca-Cola and Pepsi, two closely competing brands, adopted different strategies in their use of moral language in advertising copy. Coca-Cola drew on multiple moral foundations, emphasizing loyalty and fairness, whereas Pepsi relied almost exclusively on fairness. However, Pepsi's use of fairness-related language should be interpreted with caution. Under the strict consensus criteria, no morally charged advertising copy from this brand was identified. This pattern suggests that the fairness signals attributed to Pepsi in the full corpus were likely driven by false positives rather than by substantive moral framing. Consequently, we performed a manual check of the fairness texts for Pepsi and identified 21 matches which, after removing duplicates, corresponded to five unique ad copies. These ads mostly referenced sports-related activities (e.g., *"Cheers to more games this Christmas," "It's time to cheer like a fan, @oceanlewis takes us through all the steps to support your team,"* and *"Football and fashion is a match made in heaven… if you do it correctly"*), which leads us to conclude that the MoralStrength package [14] is likely conflating fairness-related moral language with themes of competition and sportsmanship.

At the category level (RQ2), alcoholic placements showed a dual emphasis on loyalty (23.5%) and purity (8.3%). The first frames consumption via tradition and group belonging; the second invokes product integrity (e.g., low/no-alcohol variants), underscoring the need to distinguish moral purity from descriptive cleanliness or formulation claims. Our purity diagnostic was designed precisely for this disambiguation (see (8)). In contrast, *non-alcoholic* ads prioritized fairness (17.2%), consistent with broad CSR-style appeals that link brands to equality or community support [9].

Red Bull constitutes a distinctive case within our sample because, although it is not an alcoholic beverage, it is marketed primarily to adults. Our results show that its advertising stands out for the near absence of authority- and care-related moral language. This pattern aligns with the brand's longstanding communication strategy, which foregrounds autonomy, energy, and individual achievement rather than empathy or moral constraint. As noted by Rogers [32], Red Bull's identity is not built upon normative or paternalistic discourse, but instead on an aspirational narrative that emphasizes personal freedom and the pursuit of self-defined limits. From a Moral Foundations Theory perspective, the brand neither seeks to "care for" consumers nor to invoke hierarchical authority; rather, its messaging is oriented toward freedom. Its slogan "it gives you wings" encapsulates this ethos of empowerment, positioning the individual as the agent of their own vitality.

It is worth noting that, although recent work in MFT proposes Liberty/Oppression as a sixth foundation [33], our computational tool was not designed to detect liberty cues. It is therefore plausible that Red Bull's messaging deploys liberty-based moral framing that remained unclassified by our model.

Moving to aggregate patterns across firms, our analysis of the distribution of moral foundations indicates that brands show a lower propensity to employ authority-based moral language. Prior research proposes that such restraint is expected when firms target the "cosmopolitan consumer," who is open-minded toward products and receptive to diversity and multiculturalism messaging [34]. Because our sample comprises globally recognized brands, the paid ads we analyzed likely fit this category, and thus tend to avoid authority appeals. Conversely, it is noteworthy, given Wei et al. [35], that care was not used more often. Their work suggests that care is the most salient dimension consumers evaluate when judging a brand's morality, and it is associated with greater trust and purchase intention.

4.2. Strategic (Ad Copy) vs. Spontaneous (UGC) Moral Language

Our RQ3 analyses highlight a key domain distinction. Attempts to pair foundations with discrete emotion lexicon counts (e.g., NRC) yielded sparse, sometimes empty matrices

(less a coding error than an empirical feature of professionally produced, policy-constrained copy). By design, paid ads avoid overt anger/disgust terms and polarizing vocabulary in accordance with platform standards that restrict violent, discriminatory, or politically sensitive content [17] and with industry practice in which copy variants are tuned to fit algorithmic delivery while avoiding filters [18,20]. This pattern aligns with evidence that polarization can depress purchase intention by associating a brand with opposing ideological messages, even when initial preference is high [36].

However, the text remains affectively charged. A valence-oriented analysis using VADER [28] revealed positive associations between moral foundations and the compound sentiment score, alongside positive associations with negative valence. Despite the prevalence of negative unigrams, advertisers usually try to remain constructive [37]. The correlation between moral foundations and negative sentiment ($\rho_{neg} > \rho_{pos}$) aligns with the "negativity bias" [38], whereby negatively valenced language is more cognitively salient and can, in turn, be leveraged to increase the reach of marketing strategies and product adoption [39].

These observations also point to broader issues that involve multimodality, platform governance, and algorithmic optimization. Several structural features of paid social media advertising shape the way moral language and negative language appear in the copy. First, advertisements are inherently multimodal. Text–image relations tend to be complementary rather than redundant: textual elements provide semantic precision, whereas visuals amplify affect and cultural resonance [40]. Imagery and video often elicit stronger emotional responses than text alone [41], and viewers typically attend to textual content more deeply once initial attention has been captured [16]. Verbal information therefore remains crucial for clarifying product attributes and guiding consumer inferences [15]. However, because our analyses focus exclusively on text, any moral cues conveyed primarily through visual or narrative elements are necessarily underrepresented.

Second, governance and optimization matter. Meta's ad policies shape what can be said [17], and the delivery stack encourages multiple copy/headline variants that are algorithmically recombined with creatives to optimize performance [18]. Text quality itself predicts engagement [19]. In this tuned environment, implicit moral cues (subtle, non-polarizing, policy-safe) are more likely than explicit moral claims [20]. These dynamics explain why our discrete-emotion lexicon failed in many cells while valence and foundations still registered signals.

4.3. Qualitative Insights from Copy Analysis

Our qualitative reading of the exemplar advertisements (Table 4) reveals that the pipeline's primary label captures the dominant moral foundation but does not exhaust the moral content present in each copy. For instance, the message we label as loyalty in `t1` celebrates collective effort toward a shared goal, yet it simultaneously invokes fairness and care when it calls for respect for women and the LGBTQIA+ community (`t2`). Such blending of signals underscores that marketers often craft multilingual moral narratives to resonate with heterogeneous audiences. Moral Foundations Theory explicitly treats morality as multi-dimensional: in addition to harm and fairness, cultures and individuals also moralize loyalty, authority and purity [8]. Because communities weight these "channels" differently, appeals that touch multiple foundations can broaden the audience, aligning with evidence that liberals emphasize care and fairness, whereas conservatives draw on the full palette of foundations [42]. Our findings suggest that brands exploit this pluralism (embedding secondary cues alongside the primary frame) to engage consumers with diverse sensitivities.

The authority messages (e.g., `t5`, `t6`) illustrate another strategic nuance. We propose that, in these examples, authority is not expressed through hierarchical or punitive rhetoric but rather through directive language that gently instructs the audience ("check this out", "join us", or "try now"). By couching authority in the form of friendly guidance, brands can motivate behavior without appearing coercive or paternalistic, aligning with contemporary content-creation practices [43,44]. However, prior research on authority-related language in calls to action has typically operationalized authority using lexemes more directly associated with obedience to authority figures, social hierarchy, or tradition (e.g., "respect," "obey," or phrases such as "follow proud leaders") [45,46]. From this perspective, the examples discussed here may also reflect false positives arising from the limitations of the proposed model, rather than unequivocal instances of moral authority framing.

Our inspection of `t8`–`t10` also reveals that purity rhetoric extends beyond religious sanctity. While MFT links purity concerns to spiritual and bodily cleanliness [8], advertisers invoke these intuitions in secular forms. Several ads frame low- or zero-sugar products as ways to avoid contaminating one's body with harmful additives, or highlight environmental stewardship through appeals to cleaning and protecting nature. Such messaging resonates with audiences for whom purity entails maintaining health or preserving the environment rather than observing sacred rituals. Research using moral foundations to frame environmental issues finds that purity cues emphasizing contamination and cleansing can increase pro-environmental concern among otherwise sceptical audiences (see [34], for evidence on purity-based environmental appeals).

Finally, the single care exemplar (`t11`) speaks directly to a cosmopolitan, liberal audience. It stresses empathy, community support and social justice, hallmarks of the individualising foundations of care and fairness. This targeting is consistent with MFT findings that care and fairness resonate most strongly with liberal audiences, whereas binding foundations such as loyalty, authority and purity appeal more to conservatives [42]. These qualitative observations complement our quantitative results by showing how advertisers flexibly combine moral registers to craft inclusive narratives and by illustrating how purity and authority frames are adapted to contemporary consumer sensibilities.

4.4. Methodological Implications and Limits

Methodologically, our dual-channel consensus rule (supervised & unsupervised) trades coverage for precision. Only 102 of 663 unique ads (15.4%) cleared the strong-agreement threshold, despite a larger set of text bodies (758), indicating that much ad copy is either non-moral or too ambiguous for confident labeling. This justifies a two-tier reporting strategy: broad, unweighted distributions for scope, and a transparently curated high-consensus subset for interpretability. The purity diagnostic further reduces false-positive risk in categories where cleanliness/formulation vocabulary is common. However, our strict consensus threshold yielded 25 high-confidence ads (3.3% of the corpus), consistent with the general scarcity of moral rhetoric [7]. Throughout the paper, we therefore frame our analysis as exploratory and caution that the patterns we report are indicative rather than definitive.

Important limitations follow from design choices. First, our models are text-only; multimodal integration is a priority for future work given known divisions of labor between visuals and text [15,16,40,41]. Second, data access is constrained: outside GB/IE, API settings (`POLITICAL_AND_ISSUE_ADS`) can bias retrieval toward issue-oriented content, and our focus on large, English-language beverage brands narrows generalizability. Third, the `moralstrength` tools were developed largely on spontaneous, user-generated discourse [14]; domain shift to curated, policy-filtered copy increases ambiguity. Fourth,

impression/spend weighting was not primary—though feasible where bounds are available—and should be explored as sensitivity.

4.5. Implications and Next Steps

Substantively, brands appear to mobilize distinct moral channels in ways that fit their identities and categories—loyalty for heritage and in-group belonging, fairness for CSR-style appeals, and purity where product integrity or health is salient. These patterns are consistent with MFT's pluralism and audience-congruence effects in persuasion [9–11,31]. Methodologically, domain-specific pipelines are needed. Three concrete directions follow from our findings: (i) integrate ad-creative imagery with copy via multimodal encoders so that visual moral cues (e.g., sanctity/purity iconography) are modeled jointly with text; (ii) align analysis with platform variant optimization (A/B asset mixing), tracking copy–creative pairings rather than text in isolation [18]; and (iii) incorporate text-strength signals and delivery-side outcomes to calibrate descriptive labels against engagement proxies [19].

Overall, mapping moral language in paid social contexts is feasible and informative, but it demands methods tuned to multimodal persuasion and platform governance. Advancing this agenda will require combining intuitionist theory [8] with multimodal measurement and policy-aware sampling so that moral communication in the marketplace can be studied at scale, with nuance and with appropriate caution. Indeed, our strict consensus rule and focus on a single product category produced a small high-confidence sample. While this enhances precision, it limits generalisability. Future work should collect larger, more diverse corpora (across categories, languages, and markets) and conduct human coding to validate automated labels. Researchers could also link moral language intensity to campaign effectiveness measures, where data are available, to assess whether moral framing affects engagement or sales.

5. Conclusions

This paper mapped how brands mobilize moral language in Meta ad copy using a two-channel Moral Foundations framework. Treating foundations as operational, lexicon-mappable dimensions, we combined an unsupervised semantic estimator (SIMON) with supervised classifiers and enforced a strict consensus rule, complemented by a purity diagnostic. The resulting labels reveal patterned, brand- and category-specific use of moral cues—most visibly in loyalty, fairness, and purity frames—while also showing that much ad text is either non-moral or too ambiguous for confident classification. Together with the valence-oriented sentiment pass, these findings support an intuitionist view of moralized persuasion in paid social contexts while remaining descriptive and non-causal by design.

Methodologically, the study underscores that ad copy is strategic, multimodal, and platform-governed: visuals carry a substantial share of affect; policies suppress overtly polarizing language; and algorithmic recombination pairs copy with creatives in ways a text-only pipeline cannot fully capture. Future work should therefore (i) integrate visuals with text via multimodal encoders, (ii) align analyses to copy–creative variant mixing and text-strength signals, (iii) incorporate impression/spend weighting where available, and (iv) extend sampling beyond English and beverages to test cross-cultural and sectoral generality. By releasing transparent rules and emphasizing reproducibility, this paper offers a tractable baseline for studying moral communication in digital advertising and a roadmap for building domain-sensitive, multimodal measurement that links moral cues, platform constraints, and audience-congruent messaging at scale.

Author Contributions: Conceptualization, M.S.C.-F. (marketing framing and brand/category selection), D.A.V.-C. (psychological theory and constructs), and L.H.T.-S. (computational study design); methodology, L.H.T.-S. (pipeline design, thresholds, consensus rules) and D.A.V.-C. (construct operationalization and validity checks); software, L.H.T.-S.; validation, D.A.V.-C. (content/construct validity) and L.H.T.-S. (technical validation); formal analysis, L.H.T.-S. (primary) with interpretive support from D.A.V.-C.; investigation, M.S.C.-F. (industry/platform context), L.H.T.-S. (data collection via API and preprocessing), and D.A.V.-C. (case review and coding guidance); resources, M.S.C.-F. (brand/market inputs) and L.H.T.-S. (API access and tooling); data curation, L.H.T.-S.; writing—original draft preparation, L.H.T.-S. (Methods/Results), M.S.C.-F. (Introduction/Discussion—marketing perspective), and D.A.V.-C. (Theoretical background/Discussion—psychology perspective); writing—review and editing, M.S.C.-F., L.H.T.-S., and D.A.V.-C.; visualization, L.H.T.-S.; supervision, M.S.C.-F. (lead) and D.A.V.-C. (theoretical oversight); project administration, M.S.C.-F. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data and code supporting the findings of this study are openly available on Zenodo at https://doi.org/10.5281/zenodo.17738960 (accessed on 27 November 2025). The repository includes the cleaned dataset retrieved from the Meta Ads Library API and the Python scripts used for moral foundation estimation necessary for reproducing the analyses.

Acknowledgments: The authors acknowledge the use of the Meta Ads Library API as the primary data source for this research and express their gratitude to their respective institutions for their academic and technical support throughout this study. The authors also acknowledge that during the preparation of this manuscript, they used OpenAI's ChatGPT (GPT-5, 2025) and Grammarly exclusively for language editing and structural refinement. The authors have reviewed and edited the generated text and take full responsibility for the content of this publication. Finally, the authors express their gratitude with their institutions for supporting this research: Universidad Autónoma de Bucaramanga, Universidad Industrial de Santander, and Universidad Cooperativa de Colombia.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

A/B	Split (variant) testing
API	Application Programming Interface
BH	Benjamini–Hochberg procedure (FDR control)
CI	Confidence interval
CSR	Corporate social responsibility
FDR	False discovery rate
GB	Great Britain
IE	Ireland
IRB	Institutional Review Board
MFT	Moral Foundations Theory
NFKC	Unicode Normalization Form KC
NRC	NRC Emotion Lexicon
RQ	Research question
SDG	Sustainable Development Goal
SIMON	Unsupervised semantic similarity estimator (moralstrength)
UGC	User-generated content
URL	Uniform Resource Locator
VADER	Valence Aware Dictionary and sEntiment Reasoner

References

1. Garg, P.; Raj, R.; Kumar, V.; Singh, S.; Pahuja, S.; Sehrawat, N. Elucidating the role of consumer decision making style on consumers' purchase intention: The mediating role of emotional advertising using PLS-SEM. *J. Econ. Technol.* **2023**, *1*, 108–118. [CrossRef]
2. Holbrook, M.B.; Batra, R. Assessing the Role of Emotions as Mediators of Consumer Responses to Advertising. *J. Consum. Res.* **1987**, *14*, 404–420. [CrossRef]
3. Bagozzi, R.P.; Gopinath, M.; Nyer, P.U. The Role of Emotions in Marketing. *J. Acad. Mark. Sci.* **1999**, *27*, 184–206. [CrossRef]
4. Kang, M.J.; Camerer, C.F. fMRI evidence of a hot-cold empathy gap in hypothetical and real aversive choices. *Front. Neurosci.* **2013**, *7*, 104. [CrossRef]
5. Vrtăna, M.; Križanová, A. The Role of Emotions in Consumer Behaviour and Impulse Buying: A Systematic Review. *Sustainability* **2023**, *15*, 13337. [CrossRef]
6. Ramos, G.; Goldfarb, A.; Misra, S.; Shafir, E.; Small, D.A.; Sussman, A.B.; Teeny, J.D.; Tomaino, G.; Williams, E.F. Moral Psychology for Consumer Research. *J. Consum. Psychol.* 2024, *target article, in press.*
7. Atari, M.; Mehl, M.R.; Graham, J.; Doris, J.M.; Schwarz, N.; Davani, A.M.; Omrani, A.; Kennedy, B.; Gonzalez, E.; Jafarzadeh, N.; et al. The paucity of morality in everyday talk. *Sci. Rep.* **2023**, *13*, 5967. [CrossRef]
8. Haidt, J. The New Synthesis in Moral Psychology. *Science* **2007**, *316*, 998–1002. [CrossRef]
9. Goenka, S.; Thomas, M. Moral Foundations Theory and Consumer Behavior. *J. Consum. Psychol.* **2024**, *34*, 536–540. [CrossRef]
10. Winterich, K.P.; Zhang, Y.; Mittal, V. How political identity and charity positioning increase donations: Insights from Moral Foundations Theory. *Int. J. Res. Mark.* **2012**, *29*, 346–354. [CrossRef]
11. Kidwell, B.; Farmer, A.; Hardesty, D.M. Getting Liberals and Conservatives to Go Green: Political Ideology and Congruent Appeals. *J. Consum. Res.* **2013**, *40*, 350–367. [CrossRef]
12. Yang, S.; Maloney, E.K.; Tan, A.S.L.; Cappella, J.N. When Visual Cues Activate Moral Foundations: Unintended Effects of Visual Portrayals of Vaping within Electronic Cigarette Video Advertisements. *Hum. Commun. Res.* **2018**, *44*, 223–246. [CrossRef]
13. Wannow, S.; Haupt, M.; Ohlwein, M. Is brand activism an emotional affair? The role of moral emotions in consumer responses to brand activism. *J. Brand Manag.* **2024**, *31*, 168–192. [CrossRef]
14. Araque, O.; Gatti, L.; Kalimeri, K. MoralStrength: Exploiting a Moral Lexicon and Embedding Similarity for Moral Foundations Prediction. *Knowl.-Based Syst.* **2020**, *191*, 105184. [CrossRef]
15. Smith, R.A. The Effects of Visual and Verbal Advertising Information on Consumers' Inferences. *J. Advert.* **1991**, *20*, 13–24. [CrossRef]
16. Hernández-Méndez, J.; Muñoz-Leiva, F. What type of online advertising is most effective for eTourism 2.0? An eye tracking study based on the characteristics of tourists. *Comput. Hum. Behav.* **2015**, *50*, 618–625. [CrossRef]
17. Meta Platforms, Inc. Meta Advertising Standards. 2025. Available online: https://transparency.meta.com/policies/ad-standards/ (accessed on 25 October 2025).
18. Loomer, J. The Evolution of Meta Ads Copy and Creative Control. 2025. Available online: https://www.jonloomer.com/meta-ads-copy-and-creative/ (accessed on 25 October 2025).
19. Mishra, S.; Hu, C.; Verma, M.; Yen, K.; Hu, Y.; Sviridenko, M. TSI: An Ad Text Strength Indicator using Text-to-CTR and Semantic-Ad-Similarity. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management, Virtual Event, Queensland, Australia, 1–5 November 2021; pp. 4036–4045. [CrossRef]
20. Brown, M.G.; Carah, N.; Robards, B.; Dobson, A.; Rangiah, L.; Lazzari, C.D. No Targets, Just Vibes: Tuned Advertising and the Algorithmic Flow of Social Media. *Soc. Media + Soc.* **2024**, *10*, 1–15. [CrossRef]
21. World Health Organization. *Global Status Report on Alcohol and Health and Treatment of Substance Use Disorders*; World Health Organization: Geneva, Switzerland, 2024.
22. United Nations. *Transforming Our World: The 2030 Agenda for Sustainable Development*; Resolution adopted by the General Assembly on 25 September 2015, A/RES/70/1; United Nations: New York, NY, USA, 2015.
23. Vanmassenhove, E.; Shterionov, D.; Gwilliam, M. Machine Translationese: Effects of Algorithmic Bias on Linguistic Complexity in Machine Translation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*; Merlo, P., Tiedemann, J., Tsarfaty, R., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2021; pp. 2203–2213. [CrossRef]
24. Lee, M.T.; Raschke, R.L.; Krishen, A.S. Understanding ESG scores and firm performance: Are high-performing firms E, S, and G-balanced? *Technol. Forecast. Soc. Chang.* **2023**, *195*, 122779. [CrossRef]
25. AlHares, A.; Elamer, A.A.; Alshbili, I.; Moustafa, M.W. Board structure and corporate R&D intensity: Evidence from Forbes global 2000. *Int. J. Account. Inf. Manag.* **2020**, *28*, 445–463. [CrossRef]
26. Keller, K.L. Conceptualizing, Measuring, and Managing Customer-Based Brand Equity. *J. Mark.* **1993**, *57*, 1–22. [CrossRef]
27. Alfano, M.; Cheong, M.; Curry, O.S. Moral universals: A machine-reading analysis of 256 societies. *Heliyon* **2024**, *10*, e25940. [CrossRef]

28. Hutto, C.J.; Gilbert, E. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proc. Int. AAAI Conf. Web Soc. Media (ICWSM)* **2014**, *8*, 216–225. [CrossRef]
29. Benjamini, Y.; Hochberg, Y. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *J. R. Stat. Soc. Ser. B (Methodol.)* **2018**, *57*, 289–300. [CrossRef]
30. Mohammad, S.M.; Turney, P.D. Crowdsourcing a Word-Emotion Association Lexicon. *Comput. Intell.* **2013**, *29*, 436–465. [CrossRef]
31. Ramos, G.A.; Johnson, W.; VanEpps, E.M.; Graham, J. When consumer decisions are moral decisions: Moral Foundations Theory and its implications for consumer psychology. *J. Consum. Psychol.* **2024**, *34*, 519–535. [CrossRef]
32. Rogers, G.C. Red Bull-It Gives You Wings! An Examination of the Emotional Experiences That Drive the Brand for the Popular Energy Drink. *J. Integr. Stud.* **2014**, *5*, 1–9.
33. Harper, C.A.; Rhodes, D. Reanalysing the factor structure of the moral foundations questionnaire. *Br. J. Soc. Psychol.* **2021**, *60*, 1303–1329. [CrossRef] [PubMed]
34. Prince, M.; Yaprak, A.N.; Palihawadana, D. The moral bases of consumer ethnocentrism and consumer cosmopolitanism as purchase dispositions. *J. Consum. Mark.* **2019**, *36*, 429–438. [CrossRef]
35. Wei, Y.; Ekinci, Y.; Sit, K.J. More than law-abiding: A multi-staged consumer study on brand morality. *Psychol. Mark.* **2025**, *42*, 600–614. [CrossRef]
36. Hoewe, J.; Hatemi, P.K. Brand Loyalty Is Influenced by the Activation of Political Orientations. *Media Psychol.* **2017**, *20*, 428–449. [CrossRef]
37. Loh, D.W.; Parker, A.; Jefferies, L. Relationship Between Brand Presence and Emotions on Overall Acceptance and Purchase Intent of Commercial Chicken Noodle Soup. *Foods* **2025**, *14*, 3505. [CrossRef]
38. Rozin, P.; Royzman, E.B. Negativity Bias, Negativity Dominance, and Contagion. *Pers. Soc. Psychol. Rev.* **2001**, *5*, 296–320. [CrossRef]
39. Frank, D.A.; Chrysochou, P.; Mitkidis, P. The paradox of technology: Negativity bias in consumer adoption of innovative technologies. *Psychol. Mark.* **2023**, *40*, 554–566. [CrossRef]
40. El-Sayed, Y.M.E.S. Text-Image Relations in Print Advertisement: A Multimodal Discourse Analysis. *Egypt. J. Engl. Lang. Lit. Stud.* **2018**, *9*, 543–605. [CrossRef]
41. Otamendi, F.J.; Martín, D.L.S. The Emotional Effectiveness of Advertisement. *Front. Psychol.* **2020**, *11*, 2088. [CrossRef] [PubMed]
42. Graham, J.; Haidt, J.; Nosek, B.A. Liberals and conservatives rely on different sets of moral foundations. *J. Personal. Soc. Psychol.* **2009**, *96*, 1029–1046. [CrossRef]
43. Bünzli, F.; Eppler, M.J. Calls to action and user engagement: The role of visual and verbal requests in nonprofit advocacy on Facebook. *Comput. Hum. Behav. Rep.* **2025**, *18*, 100611. [CrossRef]
44. Larsson, A.O.; Tønnesen, H.; Magin, M.; Skogerbø, E. Calls to (what kind of?) action: A framework for comparing political actors' campaign strategies across social media platforms. *New Media Soc.* **2025**, *27*, 3807–3828. [CrossRef]
45. Hoover, J.; Johnson, K.; Boghrati, R.; Graham, J.; Dehghani, M. Moral Framing and Charitable Donation: Integrating Exploratory Social Media Analyses and Confirmatory Experimentation. *Collabra Psychol.* **2018**, *4*, 9. [CrossRef]
46. Erceg, N.; Burghart, M.; Cottone, A.; Lorimer, J.; Manku, K.; Pütz, H.; Vlašiček, D.; Willems, M. The Effect of Moral Congruence of Calls to Action and Salient Social Norms on Online Charitable Donations: A Protocol Study. *Front. Psychol.* **2018**, *9*, 1913. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.