

Article

Not peer-reviewed version

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

[Mauren S. Cardenas-Fontecha](#), [Leonardo Hernan Talero-Sarmiento](#)^{*}, [Diego A. Vásquez-Caballero](#)

Posted Date: 2 December 2025

doi: 10.20944/preprints202512.0244.v1

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






Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

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 Universidad Autonoma de Bucaramanga, Colombia

² Universidad Industrial de Santander, Colombia

³ Universidad Cooperativa de Colombia, Colombia

* Correspondence: mcardenas894@unab.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, enforce a strict cross-channel consensus rule, and add 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) shows that moralized copy tends toward a negative valence that 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

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].

Moral Foundations Theory (MFT) provides a tractable framework for categorizing moral content into care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, and purity/degradation [7]. 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; these intuitions coevolve with culture, yielding plural moral “channels” that communities weight differently [6,8]. In marketing contexts, messages can trigger these channels by appealing to (or violating) specific foundations, shaping evaluations and intentions; evidence links morally charged communication to shifts in consumer responses, particularly when appeals align with audience moral identities [9–12].

MFT is not the only lens for moralized persuasion. Value-based accounts (e.g., Schwartz’s basic human values) emphasize universal motivational goals that segment audiences by prioritized ends [13,14]; the norm-activation model (NAM) and the value-belief-norm (VBN) framework highlight the role of personal norms, awareness of consequences, and responsibility appraisals in prosocial and sustainability decisions [15–17]. Dual-process work in moral judgment underscores the interaction of affective and controlled processes [18]. Alternative theories propose different primitives: Morality-as-Cooperation grounds moral rules in solutions to recurrent cooperation problems [19–21], whereas Dyadic Morality posits perceived harm as a central organizing construct [22]. Identity-based accounts examine how *moral identity* (internalization/symbolization) moderates responsiveness to moralized messages [23]. We situate our study within this broader landscape but adopt MFT pragmatically because (i) it offers *operationally distinct*, lexicon-mappable dimensions used in consumer research [6,8]; (ii) it generates audience-congruence predictions already validated in marketing and prosocial behavior [9,10]; and (iii) it aligns with available computational tools for large-scale text analysis [24].

A key reason moralized persuasion may look different in advertising is *multimodality*. Ad copy is brief and often subordinated to a central visual or audiovisual asset. Visual rhetoric suggests that text-image relations in ads are complementary, not redundant: text provides precision, while visuals activate emotions and cultural associations that amplify meaning [25]. Images possess an intrinsic persuasive force and often elicit stronger affective responses than text alone [26]; nevertheless, verbal information is indispensable for clarifying attributes and guiding inferences [27]. Eye-tracking shows a division of labor: visuals capture initial attention, whereas textual elements—once fixated—receive longer reading and support more specific inferences [28]. From a measurement standpoint, copy fields are standardized, exportable, and explicitly optimized by platforms and advertisers, making them a tractable locus for large-scale analysis even amid multimodality.

Platform governance further shapes moral expression. Meta’s review and policy systems restrict violent, discriminatory, or politically sensitive content [29]. Simultaneously, the ad-delivery stack encourages multiple variants of primary text, headlines, and descriptions that are algorithmically recombined with different creatives to optimize outcomes [30], and text quality predicts click-through and related metrics [31]. In practice, marketers tune language to “fit” algorithms and policies, embedding *implicit* moral cues that preserve resonance without triggering filters [32].

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 [33,34]. 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)? and (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 [8]) 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: United States, Great Britain (GB), Ireland (IE), Canada, Australia, New Zealand, South Africa, and Singapore.

Brand selection followed the 2025 *Forbes* Global 2000 list (Food, Drink, and Tobacco)¹ and a focus on multinational beverage firms with high consumer visibility and policy relevance. For RQ2 we contrasted *alcoholic* (*Heineken*, *Hennessy*, *Corona*) versus *non-alcoholic* (*Coca-Cola*, *Pepsi*, *Red Bull*, *Nescafé*) categories given 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.

¹ <https://www.forbes.com/global2000/list/>

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

We operationalized Moral Foundations Theory (care, fairness, loyalty, authority, purity) using the `moralstrength` library [24]. Two complementary estimators were employed: an unsupervised semantic similarity model (SIMON) and a lexicon-based supervised estimator. Where applicable, we pinned dictionary resources to their latest available release to stabilize vocabulary coverage.

2.3.1. Unsupervised Estimator (SIMON)

For each text t and foundation f , SIMON computes the cosine similarity

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

with $C_f(t) \in [-1, 1]$. For interpretability we rescale to a 0–10 metric:

$$U_f(t) = 5(C_f(t) + 1) \in [0, 10]. \quad (2)$$

We summarize SIMON by the maximum score and the winning margin

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

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

2.3.2. Supervised Estimator

The supervised channel is a lexicon-based unigram+count estimator that maps foundation-specific dictionary evidence to calibrated probabilities $S_f(t) \in [0, 1]$. Let $S_f^{(m)}(t)$ denote the score for foundation f under model m ; the primary model is m^* . We summarize by

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

and record $\text{top_SUP}(t) = \arg \max_f S_f^{(m^*)}(t)$. Probability scores are used only for ranking, thresholding, and correlation analyses; no downstream causal inference is attempted.

2.4. Consensus Rules and Label Assignment

To prioritize precision, we assign a high-confidence label only when the two channels agree at sufficient confidence. With thresholds chosen a priori to favor precision over recall and varied in sensitivity checks,

$$\tau_S = 0.75, \quad \tau_U = 7.5, \quad (5)$$

the consensus indicator is

$$I_{\text{cons}}(t) = \begin{cases} 1, & \text{sup_max}(t) > \tau_S \wedge \text{sim_max}(t) > \tau_U \wedge \text{top_SUP}(t) = \text{top_SIMON}(t), \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

We flag potential ambiguity when either channel exhibits a small winning margin. Using a common 0–1 scale for comparison, let $\Delta_U^{\text{norm}}(t) = \Delta_U(t)/10$. With $\varepsilon = 0.05$,

$$\text{mixed_flag}(t) = \mathbb{1}[\Delta_S(t) < \varepsilon \vee \Delta_U^{\text{norm}}(t) < \varepsilon]. \quad (7)$$

A confidence index combines the channels on a common 0–1 metric:

$$\text{sim_max}^{\text{norm}}(t) = \text{sim_max}(t)/10, \quad c(t) = \frac{1}{2}(\text{sup_max}(t) + \text{sim_max}^{\text{norm}}(t)). \quad (8)$$

2.5. Purity Diagnostic

To detect potential misclassification of marketing claims related to cleanliness or formulation (for example, “clean,” “zero sugar”), we compute a non-overriding diagnostic

$$J_{\text{purity}}(t) = \mathbb{1}[\text{top_SUP}(t) = \text{purity} \wedge U_{\text{purity}}(t)/10 < 0.60 \wedge \mathcal{R}(t) = 1], \quad (9)$$

where $\mathcal{R}(t)$ matches a curated term list. This flag supports sensitivity analyses and does not alter labels.

2.6. Moral–Sentiment Co-Occurrence (RQ3)

We first evaluated the NRC Emotion Lexicon, which yielded no matches in this corpus. As an alternative suited to short, informal text, we used VADER [35]. For each text t , VADER outputs positive, negative, and neutral scores and a normalized compound score $s_{\text{compound}}(t) \in [-1, 1]$. These metrics are related to moral estimates through non-parametric association measures (Section 2.7).

2.7. Analytical Strategy, Weighting, and Inference

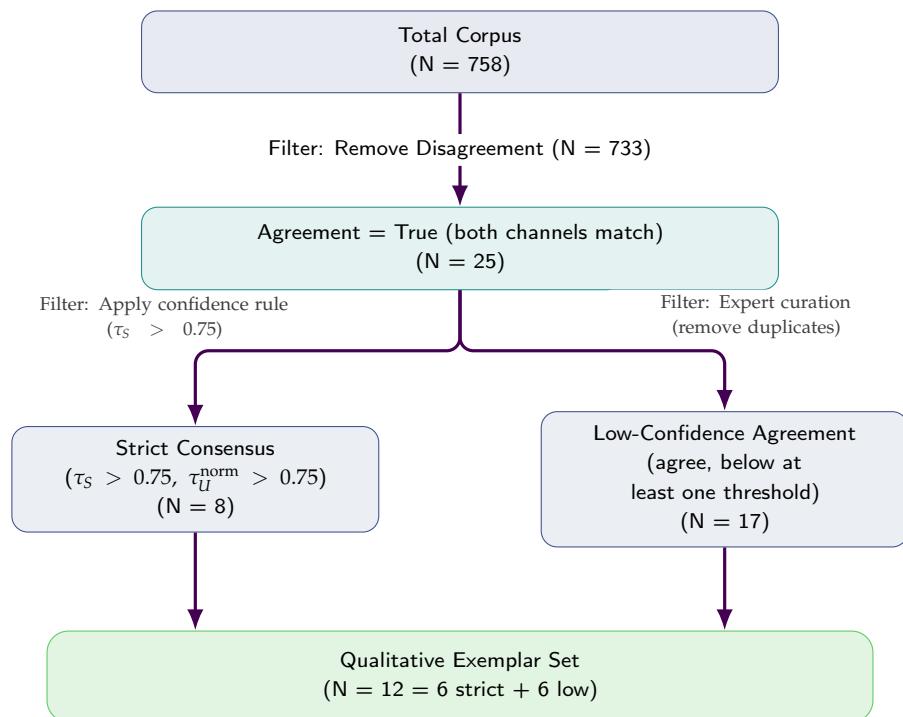
Analyses are descriptive and aligned with the three research questions. After exploding creatives, each text body is one observation linked to its parent ad. Unless stated otherwise, estimates use the full corpus of $N = 758$ text bodies pooled across markets; for interpretability we also report summaries on the high-confidence subset where $I_{\text{cons}}(t) = 1$.

For RQ1, we tabulate brand-level distributions over the label space care, fairness, loyalty, authority, purity, or non-moral using the final label $\ell(t)$. For RQ2, we aggregate labels to alcoholic versus non-alcoholic categories. For RQ3, we compute Spearman rank correlations between continuous foundation scores $S_f^{(m^*)}(t)$ and VADER polarity metrics (compound, positive, negative). Correlations are computed pairwise with listwise omission of non-finite values. Where p -values are reported, we control the false discovery rate across families of tests using Benjamini–Hochberg at $\alpha = 0.05$.

Weighting reflects data availability and platform constraints. Primary estimates are unweighted at the text-body level to maintain transparency given heterogeneous disclosure across markets. As a robustness check, where impression or spend bounds are exposed, we use midpoint values of the reported lower/upper bounds as analytic weights; observations without bounds receive unit weight. Because copy-creative mixing can decouple text from served impressions, weighted results are presented as sensitivity checks.

2.8. Qualitative Exemplar Selection ($N = 12$)

From the full corpus, we removed items where the channels disagreed on the top foundation ($N = 733$; Disagreement). The remaining $N = 25$ cases exhibited inter-model agreement and were partitioned into Strict Consensus ($N = 8$; Eq. 6) and Low-Confidence Agreement ($N = 17$). Manual review identified near-duplicates; a subject-matter expert curated a balanced exemplar set of $N = 12$ texts comprising 6 Strict-Consensus and 6 Low-Confidence items 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 ads table). Each run began by dropping prior unsupervised outputs, model-suffixed columns, and derived flags to prevent leakage across experiments. All thresholds ($\tau_S = 0.75$, $\tau_U = 7.5$, $\epsilon = 0.05$, $U_{\text{purity}}/10 < 0.60$) 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 subjects 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.

3. Results

3.1. Model Calibration and Consensus

To assess the reliability of moral labels produced by our two-channel pipeline (a supervised lexical classifier and an unsupervised semantic matcher), we calibrated agreement, confidence, and ambiguity, and used these diagnostics to define a conservative selection funnel (Figures 1–3).

Figure 1 plots the full corpus ($N=758$) in the confidence–confidence plane: the x-axis is the supervised model’s maximum predicted probability (sup_max) and the y-axis is the normalized unsupervised similarity to the best-matching foundation (sim_max_norm). The upper-right “Consensus Quadrant” encodes our gold-standard rule ($\tau_S > 0.75$ and $\tau_U^{\text{norm}} > 0.75$). Point color indicates whether both channels selected the same top foundation (yellow, $\text{Agreement}=\text{True}$) or not (purple, $\text{Agreement}=\text{False}$).

Figure 2 quantifies these outcomes. Most texts fall into *Disagreement* ($N=733$, 96.7%), while only $N=25$ texts (3.3%) exhibit channel agreement. Of the agreed cases, *Consensus Met* ($N=8$, 1.1%) satisfy

the high-confidence rule (inside the quadrant), and *Agree (Low Confidence)* (N=17, 2.2%) agree but fail one or both thresholds.




Figure 1. Methodology consensus: confidence scores (N=758). Each point represents a text positioned by supervised confidence (sup_max , 0–1) and normalized unsupervised confidence (sim_max_norm , 0–1). Color encodes channel agreement on the top foundation (yellow = agree; purple = disagree). The upper-right quadrant (N=8) satisfies the gold-standard rule ($\tau_S > 0.75$, $\tau_U^{norm} > 0.75$).

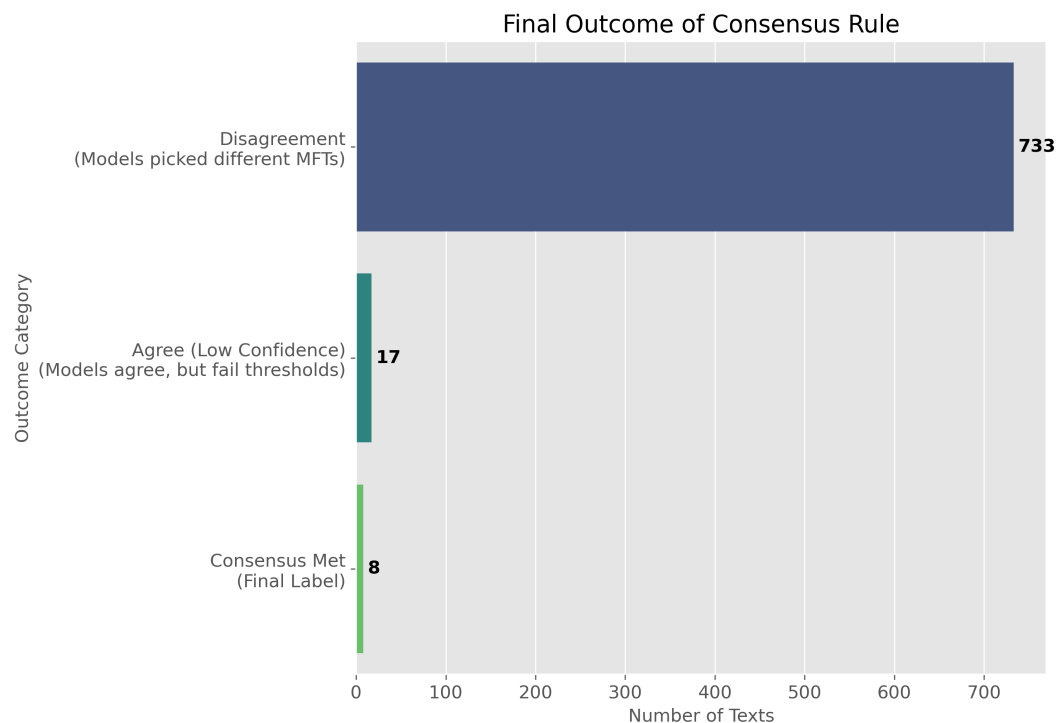


Figure 2. Final outcome of the consensus rule (N=758). Most texts are *Disagreement* (N=733); only N=25 show channel agreement, split into *Consensus Met* (N=8) and *Agree (Low Confidence)* (N=17).

Finally, Figure 3 diagnoses ambiguity via the “winning margin” Δ , defined as the gap between the highest and second-highest foundation score within each channel. The supervised lexical model (left, Δ_S) is highly decisive, with margins concentrated near 1.0. In contrast, the unsupervised semantic model (right, Δ_U^{norm}) is the main source of ambiguity, with a prominent cluster below $\Delta < 0.05$ (near-ties).

The agreement subset, consisting of 25 advertisements, represents the high-precision segment of the corpus used for qualitative inspection. Within this group, the most reliable references—termed the *Consensus Met* core—comprise eight cases where both the supervised and unsupervised estimators converge with strong confidence. Most observed disagreements arise from semantic near-ties between foundations, which validates our decision to require both estimator agreement and high confidence as conditions for final label assignment. The remaining cases categorized as *Agree (Low Confidence)* can be included in robustness checks to expand coverage, although their interpretive reliability is limited and they should therefore be treated with appropriate caution.

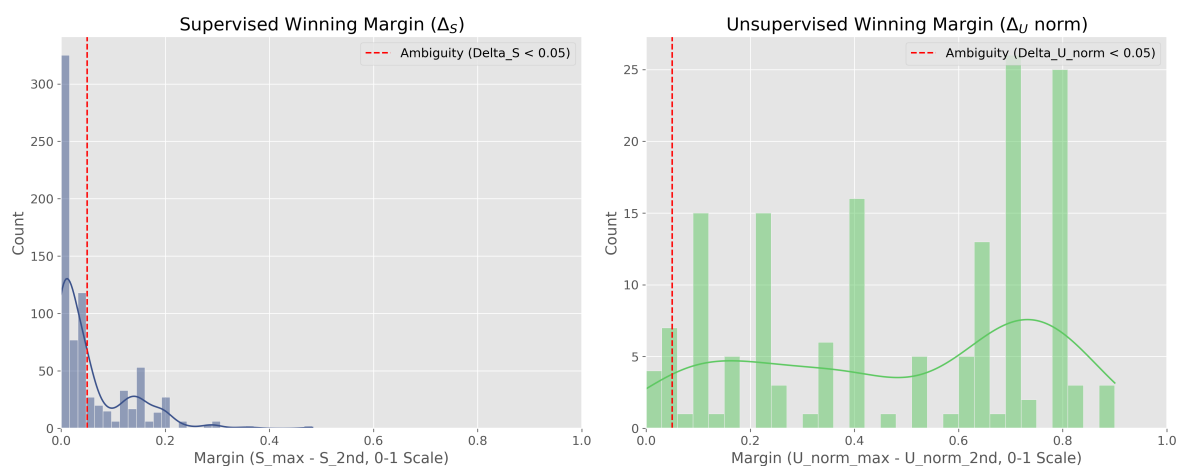


Figure 3. Distribution of winning margins (ambiguity analysis, $N=758$). The supervised lexical channel (left, Δ_S) is decisive (margins near 1.0). The unsupervised semantic channel (right, Δ_U^{norm}) drives ambiguity, with many near-ties below $\Delta < 0.05$.

3.2. RQ1: Distribution of Moral Foundations in the Full Corpus

We first analyze the distribution of final moral labels across the full 758-item corpus. The heatmap in Figure 4 and the detailed percentages in Table 1 reveal distinct brand-level strategies.




Figure 4. Heatmap of moral foundation distribution (share of ad copy %) by brand for the full corpus ($N=758$). Brands exhibit distinct moral foundation preferences.

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.

Table 1. Moral Foundation Distribution by Brand (Share of Ad Copy %).

moral label	final	lower	care	fairness	loyalty	authority	purity	non-moral	mixed
Brand									
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

3.3. Category Contrast (Alcoholic vs. Non-Alcoholic)

Aggregating brands into 'Alcoholic' and 'Non-Alcoholic' categories reveals a clear divergence in moral framing, as shown in Figure 5 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.

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.

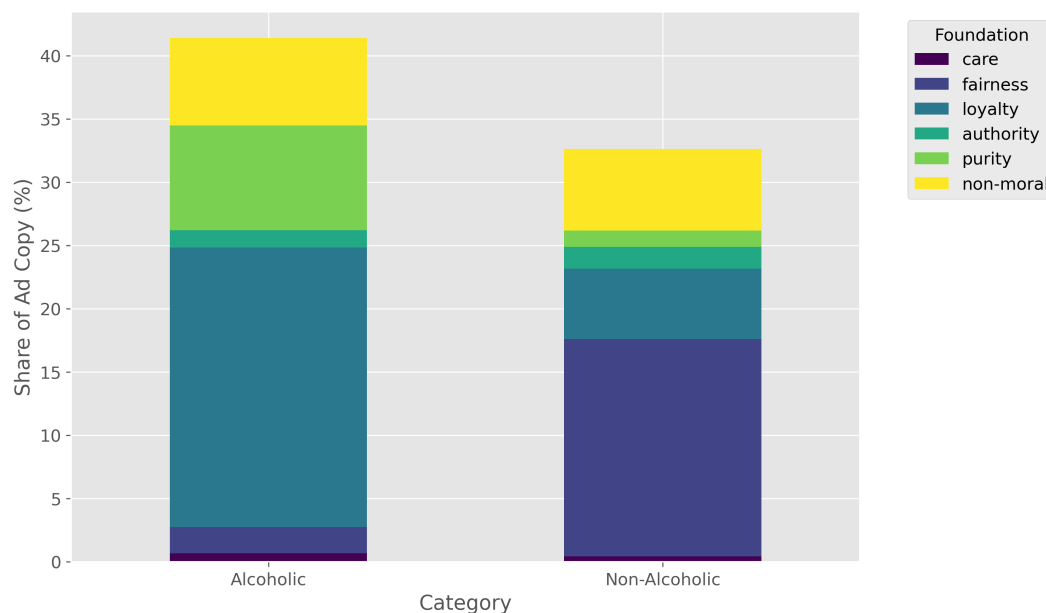


Figure 5. A stacked bar chart comparing the relative share of moral foundations used in ad copy for Alcoholic vs. Non-Alcoholic beverage categories. 'mixed' category is excluded.

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

moral label	final	lower	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

3.4. RQ3: Moral-Sentiment Co-Occurrence

To answer RQ3, we assessed the co-occurrence of moral language with emotional sentiment. An initial pass using the NRC Emotion Lexicon (nrclx) 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's rank correlation (ρ) between these scores and our supervised moral foundation scores.

The results (Figure 6) 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 6. A correlation matrix (heatmap) showing the Spearman's rho coefficient between supervised moral foundation scores and VADER sentiment polarity scores

3.5. Illustrative Exemplars and Decision Mechanics

To illustrate the labeling pipeline in practice, we selected a subset of twelve advertisements (t1–t12) for a qualitative "deep dive" (Table 3). These items span both alcoholic and non-alcoholic brands and include multiple creative styles (promotions, corporate news, 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 Eq. 6.

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

id	Approach	Brand	Ad Copy	Moral Label
t8	Extrinsic	Coca-Cola	Step 1: Find a recycling machine at a Merlin Attraction. Step 2: Recycle your 500ml 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/ . *Recycling machines are located inside participating attractions; entry ticket required.	purity
t9	Extrinsic	Heineken	Claim your free pint of Heineken® 0.0 today. “Redemption never tasted this good.”	purity
t10	Extrinsic	Heineken	Is it still home advantage if you’re on clean-up duty? Heineken 0.0. Great taste. Zero alcohol.	purity
t11	Extrinsic	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	Extrinsic	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 *Extrinsic* identifies ad copies that satisfy the high-confidence rule defined in Eq. 6.

For each text, we estimated foundation scores using both the unsupervised SIMON model and the supervised classifiers (Table 5). The final label was assigned using the consensus rule (Eq. 6) from our method. Table reports the metrics and final decisions for this subset.

Table 5. Moral foundation scores by channel for each advertisement (t_1 – t_{12}) using the unsupervised SIMON (0–10) and supervised SUP (0–1) models.

Foundation	Metric	Text ID											
		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 (6). 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

In this illustrative set, all twelve items satisfied the strong consensus rule. The resulting labels show *loyalty* and *purity* as dominant, with *care* appearing in two cases (t11, t12). This subset also allows us to see the diagnostic rules in action. For example, several rows (e.g., t1, t3, t4) exhibit supervised ties ($\Delta_S = 0.000$) but are not flagged as *mixed* because the strong cross-channel consensus rule takes precedence. This confirms the value of the two-channel approach over a simpler supervised-only model.

4. Discussion

This study offered a descriptive map of moral language in digital advertising by applying a two-channel Moral Foundations Theory (MFT) pipeline to Meta ad copy. In line with intuitionist accounts that emphasize fast, affect-laden appraisals [7,36,37], we treated foundations as operationally distinct, lexicon-mappable dimensions [24] 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 multi-modality, platform governance, and algorithmic optimization.

4.1. Interpreting Brand and Category Strategies

Moral language is not a uniform background feature of ad copy; it is strategically deployed and patterned by brand identity and product category. At the brand level, we observed stark differences,

for example, Hennessy concentrated on *loyalty* (53.8%), consistent with heritage/in-group positioning, whereas Pepsi emphasized *fairness* (40.0%), aligning with community-centric and social-good narratives. Such alignment is consistent with prior findings that persuasion improves when moral cues fit audience identities or ideological profiles [9,10] and with MFT's prediction that communities differentially weight moral "channels" [7,36].

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. As noted by prior work [38], the so-called "cola wars" have evolved into a more complex landscape now that soft-drink companies market diverse product lines, which makes traditional comparisons less straightforward. These differences may reflect distinct social-media marketing approaches tailored to each firm's specific challenges. Experimental evidence also shows that moral and political framing can shape consumer preferences, comparing Coca-Cola and Pepsi directly [39]. That work finds that ingroup versus outgroup language can shift preference toward one brand or the other, which could help explain Coca-Cola's greater emphasis on loyalty as a core foundation.

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 9). In contrast, *non-alcoholic* ads prioritized *fairness* (17.2%), consistent with broad CSR-style appeals that link brands to equality or community support, again reflecting that moral rhetoric can scaffold prosocial positioning beyond generic affect [37]. These descriptive fingerprints illustrate MFT's utility for mapping how brands "speak morality" differently across markets and categories without presuming effectiveness.

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 long-standing communication strategy, which foregrounds autonomy, energy, and individual achievement rather than empathy or moral constraint. As noted by Rogers [40], 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.

A similar interpretation emerges from scholarship on contemporary branding strategies. Sanchis-Roca et al. [41] argue that Red Bull exemplifies a paradigmatic shift in digital communication, in which brands abandon top-down, authority-driven models in favor of horizontal, participatory relationships with consumers. Through its sponsorship of extreme sports, festivals, and collective experiential events, Red Bull promotes values of adventure, challenge, and self-expression. Within this communicative ecosystem, moral languages associated with care and authority are largely incongruent with the brand's narrative. Instead, emotional appeals related to excitement, risk, and communal experience operate as key mechanisms for consumer engagement, a dynamic consistent with work on emotional versus rational appeals in advertising [42].

It is worth noting that, although recent work in MFT proposes Liberty/Oppression as a sixth foundation [43], 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.

These brand-specific patterns underscore how moral language can diverge sharply when grounded in a distinctive narrative architecture. 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 [44]. Because our sample comprises globally recognized brands, the paid ads we analyzed likely fit this category, and thus tend to avoid authority appeals. In contrast, loyalty was invoked, at least in part, by nearly all brands (with the exception of Pepsi). Relatedly, loyalty correlates positively with both consumer ethnocentrism and cosmopolitanism, a combination that can broaden audience reach and may therefore encourage more frequent use of loyalty framing [44].

Conversely, it is noteworthy, given Wei et al. [45], 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. Experimental evidence further shows that care-framed messages can increase consumer endorsement of public policies and industry regulations [46]. Nonetheless, Moral Foundations Theory holds that morality encompasses more than care [47]. For example, prior MFT research indicates that binding foundations (authority, loyalty, purity) tend to resonate more with conservative-leaning consumers, whereas liberal-leaning audiences typically prefer individualizing foundations (care, fairness) [48].

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 [29] and with industry practice in which copy variants are tuned to fit algorithmic delivery while avoiding filters [30,32]. 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 [39]. Within MFT, related effects are well documented, as people tend to reject products related to their moral outgroup [49]. Consequently, polarizing frames may backfire and alienate potential customers if they perceive the product as tied to opposing political or ideological views.

However, the text remains affectively charged. A valence-oriented analysis using VADER [35] 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 [50]. A likely explanation for this paradox is the structure of persuasive copy: high-intensity moral words are used to establish the stakes (driving the neg and compound scores up via intensity), while the surrounding non-moral text offers the solution. In other words, brands address problem-oriented topics (e.g., social inequity, environmental harms) while positioning themselves as part of the solution. This pattern is consistent with classic evidence that emotional appeals shape the consumer response [2,3] and with intuitionist models in which moral cues and affect co-travel [36]. The correlation between moral foundations and negative sentiment ($\rho_{neg} > \rho_{pos}$) aligns with the ‘negativity bias’ [51], whereby negatively valenced language is more cognitively salient and can, in turn, be leveraged to increase the reach of marketing strategies and product adoption [52]. These results show that moral language acts as an emotional amplifier; to speak morally is to abandon neutrality, utilizing a vocabulary of crisis and resolution that VADER detects as highly intense.

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 [25]. Imagery and video often elicit stronger emotional responses than text alone [26], and viewers typically attend to textual content more deeply once initial attention has been captured [28]. Verbal information therefore remains crucial for clarifying product attributes and guiding consumer inferences [27]. 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 [29], and the delivery stack encourages multiple copy/headline variants that are algorithmically recombined with creatives to optimize performance [30]. Text quality itself predicts engagement [31]. In this tuned environment, implicit moral cues (subtle, non-polarizing, policy-safe) are more likely than explicit moral claims [32]. These dynamics explain why our discrete-emotion lexicon failed in many cells while valence and foundations still registered signal.

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 τ_1 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 (τ_2). 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 [7]. 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 [48]. 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., τ_5 , τ_6) illustrate another strategic nuance. 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"). Such non-violent commands function as calls to action and are ubiquitous in social-media marketing, leveraging consumers' responsiveness to clear, actionable cues. By couching authority in the form of friendly guidance, brands can motivate behavior without appearing coercive or paternalistic, aligning with contemporary content-creation practices [53,54].

Our inspection of τ_8 – τ_{10} also reveals that *purity* rhetoric extends beyond religious sanctity. While MFT links purity concerns to spiritual and bodily cleanliness [7], 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 44, for evidence on purity-based environmental appeals].

Finally, the single *care* exemplar (τ_{11}) 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 [48]. 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.

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 [25–28]. 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 [24]; 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 [9,10,36,37]. 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 [30]; and (iii) incorporate *text-strength* signals and delivery-side outcomes to calibrate descriptive labels against engagement proxies [31].

Overall, mapping moral language in paid social is feasible and informative, but it demands methods tuned to multimodal persuasion and platform governance. Advancing this agenda will require combining intuitionist theory [7] 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.

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 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.C-F. (marketing framing and brand/category selection), D.V-C. (psychological theory and constructs), and L.T-S. (computational study design); methodology, L.T-S. (pipeline design, thresholds, consensus rules) and D.V-C. (construct operationalization and validity checks); software, L.T-S.; validation, D.V-C. (content/construct validity) and L.T-S. (technical validation); formal analysis, L.T-

S. (primary) with interpretive support from D.V-C.; investigation, M.C-F. (industry/platform context), L.T-S. (data collection via API and preprocessing), and D.V-C. (case review and coding guidance); resources, M.C-F. (brand/market inputs) and L.T-S. (API access and tooling); data curation, L.T-S.; writing—original draft preparation, L.T-S. (Methods/Results), M.C-F. (Introduction/Discussion—marketing perspective), and D.V-C. (Theoretical background/Discussion—psychology perspective); writing—review and editing, M.C-F., L.T-S., and D.V-C.; visualization, L.T-S.; supervision, M.C-F. (lead) and D.V-C. (theoretical oversight); project administration, M.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, the Python scripts used for moral foundation estimation, and supplementary materials 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
MDPI	Multidisciplinary Digital Publishing Institute
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. *Journal of Economy and Technology* **2023**, *1*, 108–118. <https://doi.org/10.1016/j.ject.2023.10.001>.
2. Holbrook, M.B.; Batra, R. Assessing the Role of Emotions as Mediators of Consumer Responses to Advertising. *Journal of Consumer Research* **1987**, *14*, 404–420. <https://doi.org/10.1086/209123>.
3. Bagozzi, R.P.; Gopinath, M.; Nyer, P.U. The Role of Emotions in Marketing. *Journal of the Academy of Marketing Science* **1999**, *27*, 184–206. <https://doi.org/10.1177/0092070399272005>.
4. Kang, M.J.; Camerer, C.F. fMRI evidence of a hot-cold empathy gap in hypothetical and real aversive choices. *Frontiers in Neuroscience* **2013**, *7*. <https://doi.org/10.3389/fnins.2013.00104>.
5. Vrtána, M.; Križanová, A. The Role of Emotions in Consumer Behaviour and Impulse Buying: A Systematic Review. *Sustainability* **2023**, *15*, 13337. <https://doi.org/10.3390/su151813337>.
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. *Journal of Consumer Psychology* **2024**. Target article, in press, <https://doi.org/10.1002/jcpy.1492>.
7. Haidt, J. The New Synthesis in Moral Psychology. *Science* **2007**, *316*, 998–1002. <https://doi.org/10.1126/science.1137651>.
8. Goenka, S.; Thomas, M. Moral Foundations Theory and Consumer Behavior. Technical Report SSRN 4979057, SSRN, 2024.
9. Winterich, K.P.; Zhang, Y.; Mittal, V. How political identity and charity positioning increase donations: Insights from Moral Foundations Theory. *International Journal of Research in Marketing* **2012**, *29*, 346–354. <https://doi.org/10.1016/j.ijresmar.2012.05.002>.
10. 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. <https://doi.org/10.1086/670610>.
11. Yang, A.X.; McFerran, B.; Soman, D. When Visual Cues Activate Moral Foundations: How Visual Cues of Purity Influence Consumer Judgments and Choices. *Journal of Advertising* **2018**, *47*, 269–281. <https://doi.org/10.1080/00913367.2018.1462642>.
12. 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. <https://doi.org/10.1057/s41262-023-00326-9>.
13. Schwartz, S.H. Universals in the Content and Structure of Values: Theoretical Advances and Empirical Tests in 20 Countries. In *Advances in Experimental Social Psychology*; Zanna, M.P., Ed.; Academic Press, 1992; Vol. 25, pp. 1–65. [https://doi.org/10.1016/S0065-2601\(08\)60281-6](https://doi.org/10.1016/S0065-2601(08)60281-6).
14. Schwartz, S.H. An Overview of the Schwartz Theory of Basic Values. *Online Readings in Psychology and Culture* **2012**, *2*, 1–20. <https://doi.org/10.9707/2307-0919.1116>.
15. Schwartz, S.H. Normative Influences on Altruism. In *Advances in Experimental Social Psychology*; Berkowitz, L., Ed.; Academic Press, 1977; Vol. 10, pp. 221–279. [https://doi.org/10.1016/S0065-2601\(08\)60358-5](https://doi.org/10.1016/S0065-2601(08)60358-5).
16. Stern, P.C.; Dietz, T.; Abel, T.; Guagnano, G.A.; Kalof, L. A Value-Belief-Norm Theory of Support for Social Movements: The Case of Environmentalism. *Human Ecology Review* **1999**, *6*, 81–97.
17. Stern, P.C. Toward a Coherent Theory of Environmentally Significant Behavior. *Journal of Social Issues* **2000**, *56*, 407–424. <https://doi.org/10.1111/0022-4537.00175>.
18. Greene, J.D.; Sommerville, R.B.; Nystrom, L.E.; Darley, J.M.; Cohen, J.D. An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science* **2001**, *293*, 2105–2108. <https://doi.org/10.1126/science.1062872>.
19. Curry, O.S. Morality as Cooperation: A Problem-Centred Approach. In *The Evolution of Morality*; Shackelford, T.K.; Hansen, R.D., Eds.; Springer: Cham, 2016; pp. 27–51. https://doi.org/10.1007/978-3-319-19671-8_2.
20. Curry, O.S.; Mullins, D.A.; Whitehouse, H. Is It Good to Cooperate? Testing the Theory of Morality-as-Cooperation in 60 Societies. *Current Anthropology* **2019**, *60*, 47–69. <https://doi.org/10.1086/701478>.
21. Curry, O.S.; Jones Chesters, M.; Van Lissa, C.J. Mapping Morality with a Compass: Testing the Theory of 'Morality-as-Cooperation' with a New Questionnaire. *Journal of Research in Personality* **2019**, *78*, 106–124. <https://doi.org/10.1016/j.jrp.2018.10.008>.
22. Schein, C.; Gray, K. The Theory of Dyadic Morality: Reinventing Moral Judgment by Redefining Harm. *Personality and Social Psychology Review* **2018**, *22*, 32–70. <https://doi.org/10.1177/1088868317698288>.
23. Aquino, K.; Reed, A.I. The Self-Importance of Moral Identity. *Journal of Personality and Social Psychology* **2002**, *83*, 1423–1440. <https://doi.org/10.1037/0022-3514.83.6.1423>.

24. Araque, O.; Gatti, L.; Kalimeri, K. MoralStrength: Exploiting a Moral Lexicon and Embedding Similarity for Moral Foundations Prediction. *Knowledge-Based Systems* **2020**, *191*, 105184. <https://doi.org/10.1016/j.knsys.2019.105184>.
25. El-Sayed, Y.M.E.S. Text-Image Relations in Print Advertisement: A Multimodal Discourse Analysis. *Egyptian Journal of English Language and Literature Studies* **2018**, *9*, 543–605. <https://doi.org/10.21608/ejels.2018.134103>.
26. Otamendi, F.J.; Martín, D.L.S. The Emotional Effectiveness of Advertisement. *Frontiers in Psychology* **2020**, *11*. <https://doi.org/10.3389/fpsyg.2020.02088>.
27. Smith, R.A. The Effects of Visual and Verbal Advertising Information on Consumers' Inferences. *Journal of Advertising* **1991**, *20*, 13–24. <https://doi.org/10.1080/00913367.1991.10673351>.
28. 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. *Computers in Human Behavior* **2015**, *50*, 618–625. <https://doi.org/10.1016/j.chb.2015.03.017>.
29. Meta Platforms, I. Meta Advertising Standards. <https://transparency.meta.com/policies/ad-standards/>, 2025. Accessed: 2025-10-25.
30. Loomer, J. The Evolution of Meta Ads Copy and Creative Control, 2025. Accessed: 2025-10-25.
31. Mishra, S.; Hu, C.; Verma, M.; Yen, K.; Hu, Y.; Sviridenko, M. TSI. In Proceedings of the Proceedings of the 30th ACM International Conference on Information & Knowledge Management. ACM, 10 2021, pp. 4036–4045. <https://doi.org/10.1145/3459637.3481957>.
32. 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. *Social Media + Society* **2024**, *10*. <https://doi.org/10.1177/20563051241234691>.
33. World Health Organization. *Global Status Report on Alcohol and Health and Treatment of Substance Use Disorders*; World Health Organization: Geneva, 2024. Licence: CC BY-NC-SA 3.0 IGO.
34. Nations, U. Transforming Our World: The 2030 Agenda for Sustainable Development. Resolution adopted by the General Assembly on 25 September 2015, A/RES/70/1, 2015. United Nations, New York.
35. Hutto, C.J.; Gilbert, E. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. In Proceedings of the Proceedings of the International AAAI Conference on Web and Social Media (ICWSM). AAAI Press, 2014, Vol. 8, pp. 216–225. <https://doi.org/10.1609/icwsm.v8i1.14550>.
36. 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. *Journal of Consumer Psychology* **2024**, *34*, 519–535. <https://doi.org/10.1002/jcpy.1427>.
37. Goenka, S.; Thomas, M. Moral foundations theory and consumer behavior. *Journal of Consumer Psychology* **2024**, *34*, 536–540. <https://doi.org/10.1002/jcpy.1429>.
38. Gertner, D.; Gertner, R.; Guthery, D. Coca-Cola's marketing challenges in Brazil: The tubaínas war. *Thunderbird International Business Review* **2005**, *47*, 231–254. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/tie.20050>, <https://doi.org/10.1002/tie.20050>.
39. Hoewe, J.; Hatemi, P.K. Brand Loyalty Is Influenced by the Activation of Political Orientations. *Media Psychology* **2017**, *20*, 428–449. publisher: Routledge _eprint: <https://doi.org/10.1080/15213269.2016.1202839>, <https://doi.org/10.1080/15213269.2016.1202839>.
40. Rogers, G.C. Red Bull-It Gives You Wings! An Examination of the Emotional Experiences That Drive the Brand for the Popular Energy Drink. *Journal of Integrated Studies* **2014**, *5*, 1–9.
41. Sanchis-Roca, G.; Canós-Cerdá, E.; Maestro-Cano, S. Red Bull, a paradigmatic example of the new communication strategies of brands in the digital environment. *Revista Latina de Comunicación Social* **2016**, *71*, 373–397. <https://doi.org/https://doi.org/10.4185/RLCS-2016-1100>.
42. Kim, C.; Jeon, H.G.; Lee, K.C. Discovering the Role of Emotional and Rational Appeals and Hidden Heterogeneity of Consumers in Advertising Copies for Sustainable Marketing. *Sustainability* **2020**, *12*, 5189. <https://doi.org/10.3390/su12125189>.
43. Harper, C.A.; Rhodes, D. Reanalysing the factor structure of the moral foundations questionnaire. *British Journal of Social Psychology* **2021**, *60*, 1303–1329. _eprint: <https://bpspsychub.onlinelibrary.wiley.com/doi/pdf/10.1111/bjso.12452>, <https://doi.org/10.1111/bjso.12452>.
44. Prince, M.; Yaprak, A.N.; Palihawadana, D. The moral bases of consumer ethnocentrism and consumer cosmopolitanism as purchase dispositions. *Journal of Consumer Marketing* **2019**, *36*, 429–438. <https://doi.org/10.1108/JCM-11-2017-2432>.

45. Wei, Y.; Ekinci, Y.; Sit, K.J. More than law-abiding: A multi-staged consumer study on brand morality. *Psychology & Marketing* **2025**, *42*, 600–614. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mar.22139>, <https://doi.org/10.1002/mar.22139>.
46. Yang, F.E.; Yang, S. Effects of Moral Frames Within Vaping Prevention Messages on Current smokers' Support for Electronic Cigarette Regulations. *Journal of Health Communication* **2023**, *28*, 412–424. publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/10810730.2023.2217104> PMID: 37246839, <https://doi.org/10.1080/10810730.2023.2217104>.
47. Haidt, J. The New Synthesis in Moral Psychology. *Science* **2007**, *316*, 998–1002. <https://doi.org/10.1126/science.1137651>.
48. Graham, J.; Haidt, J.; Nosek, B.A. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology* **2009**, *96*, 1029–1046. PMID: 19379034, <https://doi.org/10.1037/a0015141>.
49. Wannow, S.; Haupt, M.; Ohlwein, M. Is brand activism an emotional affair? The role of moral emotions in consumer responses to brand activism. *Journal of Brand Management* **2024**, *31*, 168–192. <https://doi.org/10.1057/s41262-023-00326-9>.
50. 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. Publisher: Multidisciplinary Digital Publishing Institute, <https://doi.org/10.3390/foods14203505>.
51. Rozin, P.; Royzman, E.B. Negativity Bias, Negativity Dominance, and Contagion. *Pers Soc Psychol Rev* **2001**, *5*, 296–320. Publisher: SAGE Publications Inc, https://doi.org/10.1207/S15327957PSPR0504_2.
52. Frank, D.A.; Chrysochou, P.; Mitkidis, P. The paradox of technology: Negativity bias in consumer adoption of innovative technologies. *Psychology & Marketing* **2023**, *40*, 554–566. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mar.21740>, <https://doi.org/10.1002/mar.21740>.
53. Bünzli, F.; Eppler, M.J. Calls to action and user engagement: The role of visual and verbal requests in nonprofit advocacy on Facebook. *Computers in Human Behavior Reports* **2025**, *18*, 100611. <https://doi.org/10.1016/j.chbr.2025.100611>.
54. 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 & Society* **2025**, *27*, 3807–3828. Publisher: SAGE Publications, <https://doi.org/10.1177/14614448241229156>.

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.