ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Signaling Authenticity: Verification Cues and Brand-Impersonation Scams in U.S. Beauty and Personal-Care Social Commerce

Ololade Zainab Adesokan

Digital Marketing and E-commerce Researcher, Institution and Address American National University, Salem VA 1813 E Main St, Salem, VA 24153, United States

https://orcid.org/0009-0006-6888-0552

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150900155

Received: 16 October 2026; Accepted: 21 October 2026; Published: 26 October 2026

ABSTRACT

Social commerce lets consumers buy from brands inside feeds, profiles and advertisements that impersonators can copy, and reported losses to scams that began on social media reached $2.1 billion in the United States in 2025, most often through purchases that never arrived. This study asks which U.S. beauty and personal-care brands publish verification cues that are costly to imitate, and whether those cues are associated with less brand-impersonation infrastructure. Drawing on signalling theory and information asymmetry, we measure nine publicly observable cues (SPF, DMARC enforcement, BIMI, DNSSEC, CAA, registry lock, corporate registrar, HSTS preload and TikTok verification) for 245 brands assembled from Wikidata, and count registered look-alike domains generated from each brand's name across seven top-level domains and six commerce keywords, distinguishing third-party from brand-controlled registrations. Negative binomial, hurdle and logistic models with robust standard errors show that third-party look-alike registrations rise steeply with web reach (incidence rate ratio 2.67 per log10 unit of Tranco reach) and domain age, that cue adoption rises just as steeply with reach, and that conditional on reach no cue or cue index is associated with fewer look-alikes. Two-thirds of brands have at least one third-party look-alike, and shop, store and official variants are almost always third-party. Verification cues therefore function as signals for consumers rather than deterrents for impersonators, and the brands least likely to signal are the small brands that social commerce multiplies. An Authenticity Signalling Framework assigns complementary roles to brands, platforms, regulators and consumers.

Keywords: Brand impersonation; signalling theory; verification cues; social commerce; combosquatting; email authentication; consumer protection; negative binomial regression

INTRODUCTION

Social commerce collapses the distance between discovering a product and paying for it: advertisement, creator post, brand profile and checkout sit inside one feed (Liang & Turban, 2011; Lin et al., 2017). The same design collapses the distance between a brand and its impersonator: a counterfeit storefront needs only a registered domain, a copied logo and an advertisement. The Federal Trade Commission reports that losses to scams that started on social media rose from $261 million in 2020 to $2.1 billion in 2025, that nearly 30% of fraud-loss reports began on social media, and that more than 40% of those losses began with an order for goods that mostly never arrived (Federal Trade Commission [FTC], 2026); business impersonation generated more than 330,000 reports in 2023 (FTC, 2024), prompting a trade regulation rule that prohibits materially and falsely posing as a business (Trade Regulation Rule on Impersonation of Government and Businesses, 2024).

Beauty and personal care is an exemplary setting: the category is creator-led and sold heavily through TikTok Shop, Instagram and Facebook, its products are small, branded and easily counterfeited, and purchase decisions rest on brand equity (Keller, 1993; Wilcox et al., 2009). Brands can respond with verification cues

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

that an impersonator cannot copy: cryptographic email authentication, registry-level domain protection and platform verification badges. Research on assurance seals and verification badges shows that consumers notice some cues and ignore others (Edgerly & Vraga, 2019; Hu et al., 2010; Vaidya et al., 2019), and research on typosquatting and combosquatting shows that look-alike domains are registered at scale against popular names (Agten et al., 2015; Kintis et al., 2017), but no study has measured both sides for a population of brands in one product market.

This paper does so for 245 U.S. beauty and personal-care brands, asking which brands publish verification cues (RQ1), how much impersonation infrastructure exists and in what form (RQ2), and whether, conditional on a brand's attractiveness as a target, cues are associated with less of it (RQ3). We contribute a reproducible open-source measurement of nine cues and of registered look-alike domains, descriptive evidence on the concentration and commercial orientation of impersonation infrastructure, regression evidence that separates a deterrence account from a target-attractiveness account (Figure 1), and a framework that assigns the production, transmission, enforcement and reception of authenticity signals to brands, platforms, regulators and consumers.

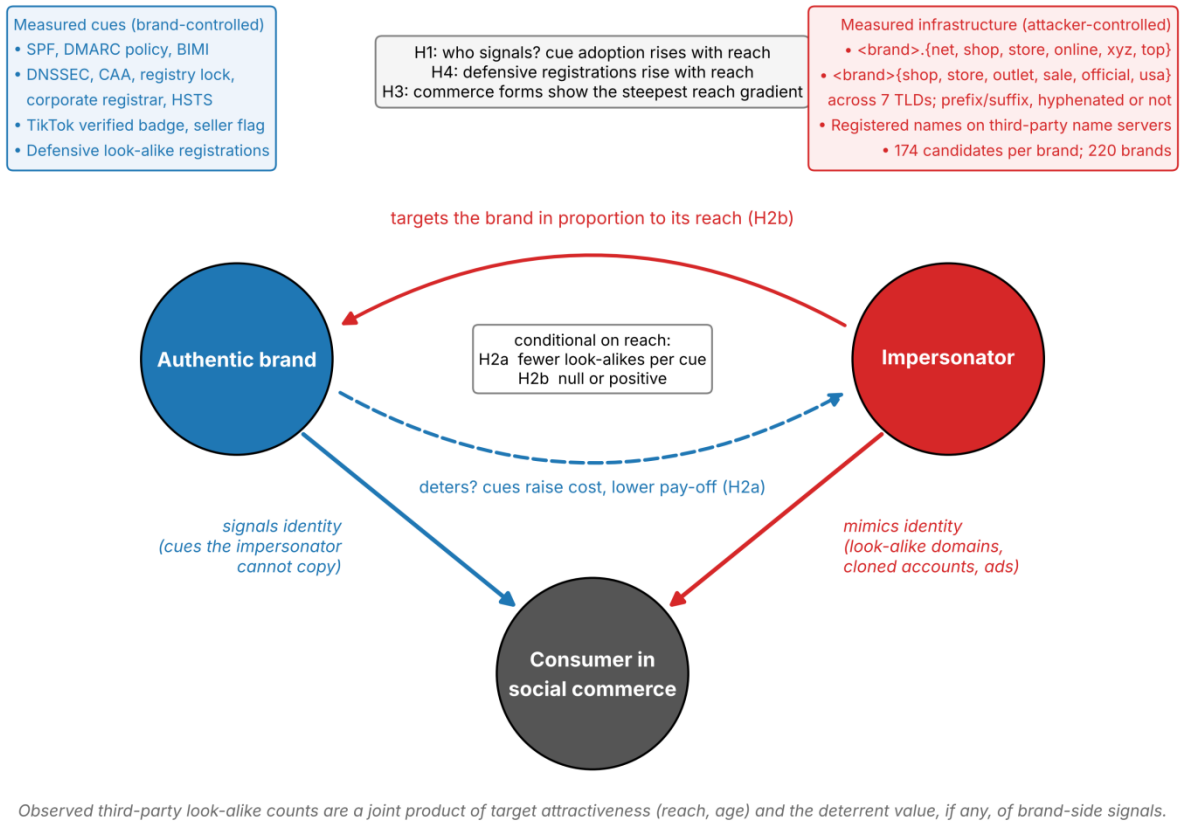


Figure 1. The signalling triad and the hypotheses tested.

THEORETICAL BACKGROUND AND HYPOTHESES

Social Commerce and the Authenticity Problem

Social commerce joins content, community and transaction on social media (Liang & Turban, 2011). Its trust mechanisms are social rather than institutional, which helps unknown sellers and impersonators alike, because the cues that build social trust can be copied (Gefen et al., 2003; McKnight et al., 2002). Impersonation takes three connected forms: look-alike domains that host fake storefronts or redirect to them (Bitaab et al., 2023; Bitaab et al., 2025), cloned social-media accounts that borrow a brand's name and imagery (Mariconti et al., 2017), and advertisements that route users from one to the other (Zeng et al., 2021). Fake e-commerce sites are numerous, short-lived and hard to distinguish from legitimate ones by content (Carpineto & Romano,

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

2017), and consumers judge legitimacy from surface features that attackers reproduce (Dhamija et al., 2006). Victims bear financial and emotional costs (Button et al., 2014) and complainants are a selected minority (Raval, 2020), so reported losses understate the problem.

Signalling Theory and Verification Cues

Information asymmetry depresses markets in which quality is unobservable (Akerlof, 1970); signalling theory explains how the informed party can credibly reveal quality through actions that are cheaper for the honest type than for the dishonest one (Connelly et al., 2011; Spence, 1973). A signal of unobservable quality works only if its cost is differentially borne (Kirmani & Rao, 2000). Verification cues are signals of identity rather than of product quality, and their credibility rests on who controls them. Email authentication (SPF, DMARC with an enforcing policy, BIMI with a certificate) is published in the brand's DNS and verified by receiving servers, so spoofed mail is rejected before the consumer sees it (Durumeric et al., 2015; Hu et al., 2018). DNSSEC, certificate-authority authorisation (CAA) and registry locks bind the domain to its owner at the registry and the certificate authority. Platform badges attest that an account belongs to the entity it names, and carry credibility for some users and purposes (Edgerly & Vraga, 2019; Morris et al., 2012; Vaidya et al., 2019), much as third-party assurance seals shape initial trust in unfamiliar web stores (Hu et al., 2010). Because these cues are costly for an impersonator but inexpensive for the brand, they satisfy the separating condition; what signalling theory does not say is whether signals change the behaviour of the party they exclude.

Signalling theory explains why a cue can separate an honest brand from an impersonator; dual-process theories explain whether consumers read it. The Elaboration Likelihood Model holds that, when motivation or ability to scrutinise is low, people rely on peripheral cues rather than on careful evaluation (Petty & Cacioppo, 1986), and the Heuristic-Systematic Model holds that a credibility heuristic, such as a source cue, can substitute for systematic processing (Chaiken, 1980). Scrolling a feed and buying on impulse are low-elaboration conditions, so a badge, a familiar logo or a domain name containing the brand works as a peripheral cue whether it is genuine or copied (Metzger & Flanagin, 2013).

The separating condition also differs by cue. DNS-based cues are published under the brand's own domain and cannot be obtained by a third party for that domain, whereas platform badges are weaker signals wherever verification can be bought by subscription or account names and imagery can be cloned (Mariconti et al., 2017); a cue that can be bought or copied shows what consumers trust without showing that the account is authentic. The two lenses are complementary: signalling theory governs the supply of credible cues by brands, and dual-process theory governs their reception by consumers. This study measures the supply side directly and treats reception as the mechanism through which cues could deter impersonators by lowering the pay-off of impersonation; Section 5 returns to the reception side.

Brand-Impersonation Infrastructure

Domain names are the most observable element of impersonation infrastructure. Typosquatting registers misspellings of popular names (Agten et al., 2015; Moore & Edelman, 2010); combosquatting combines the exact brand name with a keyword such as shop or official, is more prevalent, targets commerce and persists for years (Kintis et al., 2017). Look-alike names also appear in certificate logs and package registries (Roberts et al., 2019; Vu et al., 2020), and targets are selected by popularity, so impersonation infrastructure concentrates on high-reach brands (Kintis et al., 2017; Tian et al., 2018).

Brands respond with defensive registrations of their own name, observable in the same way, so a registered look-alike is informative only once brand-controlled and third-party registrations are distinguished. The economics of information security predicts that defences are adopted where the defender's loss is largest (Anderson & Moore, 2006), so the brands that invest in cues are those already under attack, an endogeneity any association between cues and impersonation must confront.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Hypotheses

Four hypotheses follow (Table 1). H1 concerns signal production: because cues are inexpensive but their pay-off scales with the reputation at stake, adoption rises with brand reach and equity (Anderson & Moore, 2006; Keller, 1993). H2 concerns the relationship between cues and impersonation conditional on reach and poses two competing accounts. Under the deterrence account (H2a), cues raise the cost and lower the pay-off of impersonation, so brands with more cues carry fewer third-party look-alikes. Under the target-attractiveness account (H2b), impersonators select targets by reach rather than by the defences they will meet, and cues are adopted in response to attack, so the conditional association is null or positive. H3 concerns form: because social commerce sells, impersonation infrastructure is commerce-oriented and the reach gradient is steepest

for shop, store, outlet and sale variants. H4 concerns bundling: defensive registration, the brand-side mirror of impersonation, rises with reach alongside the other cues, so brand protection is a bundle concentrated among large brands. Risk-management and cyber-resilience research frames such bundles as governance choices rather than technical ones (Gbadebo & Gbadebo, 2025; Olaniyi et al., 2023), evaluable only with open security data (Ejiofor, 2025).

Table 1 Hypotheses, Competing Predictions and Results

Hypothesis	Test	Deterrence account	Target-attractiveness account	Result
H1 Signal production	Cue index, DMARC enforcement, TikTok verification on reach (C1–C3)	+	+	Supported (+)
H2a Deterrence	Cue indices and individual cues in NB, hurdle and logistic models, conditional on reach	−	0 / +	Not supported
H2b Target attractiveness	Reach and age in the same models; sign of cue coefficients	+ (reach)	+ (reach); cues 0 / +	Supported
H3 Commerce orientation	Reach gradient for commerce-keyword vs exact-token names (R6, R9, R11)	n/a	steeper for commerce forms	Supported
H4 Bundled protection	Defensive registrations on reach (C4); correlation with cue index	+	+	Supported (+)

Note. + positive, − negative, 0 none; H2a and H2b are discriminated by the cue coefficients conditional on reach.

RESEARCH DESIGN AND METHOD

Design and Sampling Frame

The study is a cross-sectional analysis of public secondary data retrieved on 3 October 2026 (Table 2); every raw extract is archived with its retrieval timestamp. The unit of analysis is the brand. The frame was built from Wikidata by retrieving every brand, company or product-line item in cosmetics, fragrance, skin care, hair care, oral care, shaving, bath and body, or beauty retail linked to the United States, supplemented by the members of the corresponding English Wikipedia categories. Of 517 U.S.-linked items, 272 were excluded by explicit rules: parent companies, manufacturers and ingredient suppliers, people, salons and schools, household and over-the-counter products, feminine-hygiene and baby lines, fashion houses with licensed fragrances, non-U.S. brands and defunct brands (139); no official website (87); non-brand entity types (43); defunct (2); and a duplicate domain (1). The 245 remaining brands span fragrance (33.1%), cosmetics (22.9%), bath and body (9.8%), skin care (9.4%), hair care (8.6%), beauty retail (5.7%), direct selling (4.5%), shaving and grooming (4.5%) and oral care (1.6%). The frame favours brands with an encyclopaedic footprint, appropriate for a study of impersonation targets; selection is examined in Section 4.5.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

The sampling approach is a census of an enumerated frame rather than a probability or convenience sample: every item that met the inclusion rules was retained, so the 245 brands are the documented U.S. beauty and personal-care brands with an official website at the retrieval date, and the 220 distinctive-token brands form the analysis sample for the impersonation outcomes. The study has no human participants. It records no consumer demographics, attitudes or self-reports, so none of its measures depends on recall, social desirability or a respondent's willingness to admit having been deceived; the corresponding brand-level characteristics (category, web reach, TikTok audience, domain age and parent ownership) are reported in Tables 4 and 5. Only publicly available technical records about organisations were collected, and no personal data were processed. With 217 complete cases and at most 20 estimated parameters in the largest count model, the design retains more than ten complete cases per parameter.

Table 2 Data Sources (all retrieved 3 October 2026, UTC)

Source	Content used	Access	Records
Wikidata SPARQL; English Wikipedia API	Sampling frame: brand items, types, country, website, inception, parent, social handles	query.wikidata.org; en.wikipedia.org/w/api.php	517 U.S. items; 245 brands
Google Public DNS (DNS-over-HTTPS)	TXT (SPF, DMARC, BIMI), CAA, MX, A records and AD flag; NS/A of look-alike candidates	dns.google/resolve	245 domains; 41,798 candidates
RDAP (rdap.org)	Registration date, EPP status, registrar, name servers, DNSSEC delegation	rdap.org/domain	240 domains (5 registries without RDAP)
Chromium HSTS preload list	Preload status	hstspreload.org/api/v2/status	245 domains
TikTok	Account existence, verified badge, followers, seller flag	www.tiktok.com/@handle (oEmbed and profile data)	271 handles; 207 accounts
Tranco	Top-1M rank, list 94GG2	tranco-list.eu	114 ranked brands

Measures

Table 3 lists the variables. Verification cues were read from each brand's official domain with Google Public DNS and RDAP: an SPF record; a DMARC record and its policy, enforced when quarantine or reject; a BIMI record and whether it carries certificate evidence; CAA and MX records; DNSSEC; registry-level EPP locks; and the sponsoring registrar, coded as corporate when it is a brand-protection registrar such as MarkMonitor, CSC or Com Laude. HSTS preload status was read from the Chromium preload list, and TikTok verification, follower count and seller status from the brand's profile page, using the username recorded in Wikidata or, where none is recorded, the exact brand token. Two indices sum the email-authentication cues (SPF, DMARC, DMARC enforced, BIMI; 0–4) and the domain-protection cues (DNSSEC, CAA, registry lock, corporate registrar, HSTS; 0–5); a 0–9 index adds TikTok verification.

Impersonation infrastructure was measured by generating, for each brand, its brand token combined with the keywords shop, store, outlet, sale, official and usa as prefix or suffix, with and without hyphen, plus the bare token, across the com, net, shop, store, online, xyz and top top-level domains (174 candidates per brand after removing its own domain), and resolving each candidate's NS and A records, following generative typosquatting and combosquatting designs (Agten et al., 2015; Kintis et al., 2017). A candidate with name-server delegation is registered; a registered name sharing an authoritative name-server host or an IPv4 address with the brand's official domain is brand-controlled (defensive), and all others are third-party, the primary outcome, with a strict variant that also drops names sharing only a name-server provider. Tokens were taken

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

from the domain label with documented overrides; 20 brands with a generic token (a dictionary word or fewer than five letters, e.g. secret, ivory) were held out because registrations of generic strings cannot be attributed to impersonation, and five brands whose website is a parent or retailer domain have no candidate set, leaving 220.

Impersonation exposure and victimisation are defined accordingly. A brand's impersonation exposure is the number of third-party registrations among its generated look-alike names, that is, the supply of infrastructure that could host a counterfeit storefront, receive advertising traffic or send mail in the brand's name; the binary outcome records whether any such name exists. Exposure is an attribute of the brand, not of a consumer: a registered look-alike is a capability for impersonation rather than evidence of a completed scam, and consumer victimisation (a purchase from, or a loss to, an impersonator) is not observed. Victimisation is documented instead at the population level by FTC Consumer Sentinel reports (FTC, 2024, 2026) and by the Better Business Bureau, whose Scam Tracker data show that online purchase scams made up more than 30% of scams reported in 2024 and that 87.5% of those reports involved a monetary loss (BBB Institute for Marketplace Trust, 2025).Reach was measured by Tranco rank (Le Pochat et al., 2019), transformed to log10(1,000,001) minus log10(rank) with unranked brands scored zero; by TikTok followers (log1p); and by the presence of a Wikipedia article. Controls are domain age from the RDAP registration date (brand inception, available for 197 brands, is a robustness check), a recorded parent company, token length and five category groups.

Measurement validity and reliability are established by design rather than by scale statistics. Each cue is read from the authoritative record that mail servers, browsers, registries and platforms themselves consult, so the measure is the cue as a consumer's software encounters it, and each is coded deterministically from that record rather than rated. The cue indices are formative: they count distinct controls that need not co-occur, so internal-consistency coefficients such as Cronbach's alpha, which assume reflective indicators of one latent trait, are not appropriate (Diamantopoulos & Winklhofer, 2001); for this reason the models are also estimated with individual cues. Attribution error in the outcome is bounded by the strict, all-registered and resolving-only variants, and every raw extract is archived with its retrieval timestamp so that the measurement can be repeated.

Table 3 Variables and Coding

Block	**Variables**	**Source**	**Coding**
Outcomes	n_third_party; any_third_party; n_third_party_strict; n_third_party_commerce; n_third_party_resolving; n_third_party_exact; n_registered	Generated candidates + DNS	counts; 0/1
Reach / attractiveness	reach_tranco; tranco_ranked; log_tiktok_followers; wikipedia_article	Tranco; TikTok; Wikipedia	$\log10(10^6+1) - \log10(\text{rank})$; 0/1; log(1+n); 0/1
Email-authentication cues	spf_present; dmarc_present; dmarc_enforced; bimi_present; email_auth_index	DNS TXT	0/1; index 0–4
Domain-protection cues	dnssec_signed; caa_present; registry_lock; corporate_registrar; hsts_preloaded; domain_protection_index	DNS; RDAP; HSTS list	0/1; index 0–5
Platform cue	tiktok_verified; tiktok_seller; tiktok_account_found	TikTok profile data	0/1
Brand-protection activity	n_defensive (brand-controlled look-alike registrations)	Generated candidates + DNS + RDAP	count
Controls	log_domain_age (RDAP); log_brand_age (inception); has_parent_company; token_length; category group (5)	RDAP; Wikidata	log years; 0/1; characters; fixed effects

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Analytical strategy

Because the outcome is an over-dispersed count with 32% zeros, the primary model is negative binomial regression with heteroskedasticity-consistent (HC0) standard errors (Cameron & Trivedi, 2013; White, 1980), reported as incidence rate ratios (IRR) with 95% confidence intervals. Nested specifications add reach and controls to category effects (M1), then the two cue indices and TikTok verification (M2); M3 replaces the indices with individual cues, and M4 adds defensive registrations. Cues are therefore tested both individually (M3) and in combination (the indices in M2). A logistic model of any third-party look-alike reports odds ratios, discrimination and calibration; hurdle and zero-inflated NB models separate whether a brand is targeted from how often and are compared with the NB model by information criteria, expected count frequencies and Vuong statistics, whose limits are noted (Wilson, 2015). H1 and H4 are tested with Poisson, logistic and NB models of cue adoption and defensive registration on reach. Multicollinearity is assessed with generalised variance inflation factors, influence with Cook's distance, leverage and DFBETAs, and robustness by excluding influential cases and extreme counts, adding generic-token brands, substituting binary reach and account measures, using alternative outcomes and substituting brand age. All estimates are associational. Code, data and outputs accompany the paper.

RESULTS

Who signals? Cue adoption across brands

Tables 4 and 5 describe the frame. The median domain is 16.5 years old; 46.5% of brands are in the Tranco top million (median rank 171,087 among those ranked), 84.5% have a TikTok account under their recorded or exact-name handle and 41.2% a seller flag. Basic email authentication is widespread (SPF 84.1%, a DMARC record 84.9%), but enforcement is not: only 47.3% of brands enforce DMARC and 28.2% reject spoofed mail outright, 4.9% publish BIMI and 2.9% a certificate-backed BIMI record. Registry-level protection is rare: 5.3% of domains are DNSSEC-signed, 3.7% publish CAA, 8.6% hold a registry lock, 1.2% are HSTS-preloaded, and 26.9% use a corporate registrar. Thirty-one percent of brands have a verified TikTok account. The mean brand publishes 2.13 of nine cues (SD 1.47).

Adoption is steeply graded by reach (Table 6, Figure 3): 28.2% of brands outside the Tranco top million enforce DMARC and 9.2% are TikTok-verified, against 96.3% and 66.7% of brands ranked within the top 100,000, with corporate registrars (11.5% against 48.1%) and registry locks (0.8% against 18.5%) following the same gradient (χ^2 tests $p < .05$ for every cue except DNSSEC, HSTS preload and account existence). The cue-adoption models (Table 8) support H1: each log10 unit of Tranco reach multiplies the expected cue count by 1.62 (95% CI [1.41, 1.86]), with independent contributions from TikTok followers (1.04 per log unit), domain age (1.20) and a parent company (1.14); the odds of DMARC enforcement rise 13.6-fold per reach unit and, among brands with a TikTok account, the odds of a verified badge 15.9-fold. Defensive registrations also rise with reach (IRR 2.62 [1.20, 5.70]) and correlate at 0.41 with the cue index, supporting H4.

Table 4 Descriptive Statistics (N = 245)

Variable	n	Mean	SD	Min	Median	Max	Skew
Brand age (years)	197	33.5	36.3	0.6	17.8	180.8	1.8
Domain age (years)	240	18	10.3	0.4	16.5	33	-0.1
Tranco rank (ranked brands)	114	229727.8	177509	5011	171086	746590	1.1
Web reach (Tranco, log10 scale)	245	0.36	0.47	0	0	2.3	1.1
TikTok followers (accounts found)	207	454802.7	3184838	0	2585	4.5E+07	13.1
Brand-token length	245	10.2	4.2	0	10	22	0.2

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Email-authentication index (0–4)	245	2.21	0.97	0	2	4	-0.6
Domain-protection index (0–5)	245	0.46	0.73	0	0	4	1.7
Verification-cue index (0–9)	245	2.13	1.47	0	2	7	0.7
Look-alike candidates generated	240	174.2	0.4	174	174	175	1.9
Registered look-alike names	240	4.5	5.4	0	3	39	2.1
Third-party look-alike names	240	4	5	0	2	39	2.5
Defensive (brand-controlled) names	240	0.5	1.4	0	0	12	5.2
Third-party, commerce keyword	240	1.6	3	0	0	29	4.4
Third-party, exact token in other TLD	240	1.8	1.7	0	1	6	0.5
Third-party with A record	240	3.6	4.7	0	2	38	2.7

Table 5 Verification-Cue Adoption and Impersonation Prevalence (N = 245)

Indicator	**All brands n**	**All brands %**	**Analysis sample %**
SPF record published	206	84.1	85.5
DMARC record published	208	84.9	85.5
DMARC policy quarantine or reject	116	47.3	47.7
DMARC policy reject	69	28.2	27.7
BIMI record	12	4.9	5
BIMI with certificate (VMC) evidence	7	2.9	2.7
CAA record	9	3.7	3.2
DNSSEC-signed domain	13	5.3	5.5
Registry lock (server-level EPP locks)	21	8.6	8.2
Transfer lock (any)	232	94.7	95
Brand-protection (corporate) registrar	66	26.9	26.8
HSTS preload list	3	1.2	1.4
Domain in Tranco top 1M	114	46.5	46.4
English Wikipedia article	105	42.9	42.7
Parent company recorded	87	35.5	35.5
TikTok account found	207	84.5	85.9
TikTok account under recorded handle	62	25.3	26.4
TikTok verified badge	76	31	32.7
TikTok seller / commerce flag	101	41.2	43.2
Any registered look-alike name	172	70.2	69.1
Any third-party look-alike name	170	69.4	68.2

Note. Analysis sample = distinctive-token brands.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

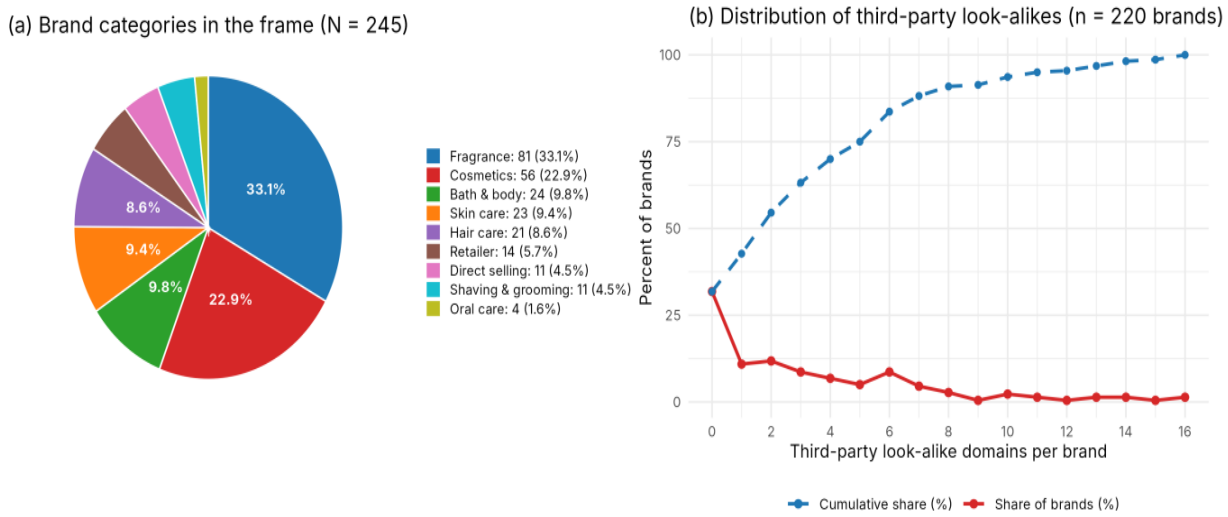


Figure 2. Brand categories (a) and third-party look-alikes per brand (b).

Table 6 Cue Adoption and Impersonation by Web Reach (% of brands unless stated)

Reach group	n	SPF	DMARC enforced	BIMI	DNSSEC	CAA	Registry lock	TikTok verified	Mean cue index	Mean defensive reg.	Mean third-party look-alikes
Not in Tranco top-1M	131	76.3	28.2	0.8	3.8	1.5	0.8	9.2	1.32	0.12	1.96
Tranco rank 100k-1M	87	92	60.9	9.2	6.9	4.6	17.2	52.9	2.9	0.6	6.35
Tranco rank <= 100k	27	96.3	96.3	11.1	7.4	11.1	18.5	66.7	3.59	1.65	6.85

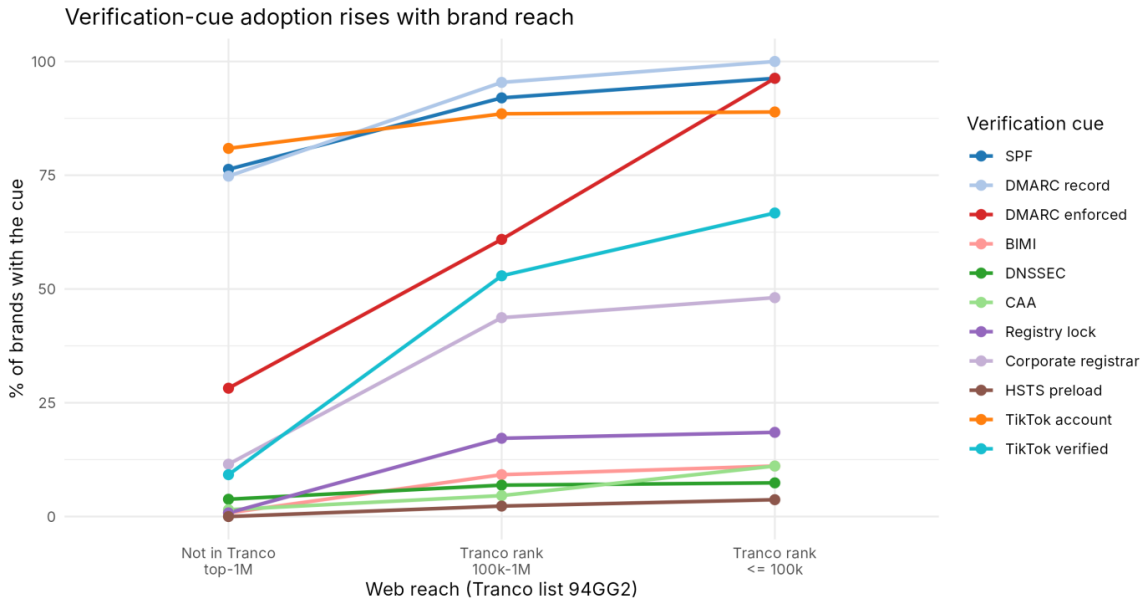


Figure 3. Verification-cue adoption by web reach.

Impersonation Infrastructure

For the 220 distinctive-token brands, 826 of 38,308 generated names were registered, 728 (88%) by third

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

parties and 98 by the brands. Two-thirds of brands (68.2%) have at least one third-party look-alike, the mean is 3.31 (SD 3.79; maximum 16; Figure 2b), the top tenth of brands hold 36.3% of third-party names, and the 46.4% of brands ranked by Tranco hold 78.8%. Table 7 and Figure 4 show their form. The exact token in another top-level domain is most common (425 registered, 382 third-party), followed by shop (156), store (77), usa (62), sale (46), official (34) and outlet (26); third parties hold 87–94% of the shop, store, usa and official names. By top-level domain, .com (332) leads, followed by .shop (140), .net (121), .xyz (80), .store (74), .online (60) and .top (19). Of the third-party names, 89.6% resolve to an address and 36.7% sit on parking or aftermarket name servers (GoDaddy's Afternic is the most common), so roughly a third are held for sale or monetised traffic and the rest serve content or redirect. The mean number of third-party look-alikes rises from 1.3 among brands outside the Tranco top million (46.6% with any) to 5.2 among brands ranked 100,000–1,000,000 (90.9%) and 6.9 among brands in the top 100,000 (100%; Figure 4c).

Table 7 Registered Look-Alike Names by Top-Level Domain and Name Form (analysis sample)

Dimension	Level	Registered	Third-party	Defensive	% third-party
TLD	com	332	283	49	85%
TLD	shop	140	129	11	92%
TLD	net	121	98	23	81%
TLD	xyz	80	76	4	95%
TLD	store	74	68	6	92%
TLD	online	60	56	4	93%
TLD	top	19	18	1	95%
Form	exact token (other TLD)	425	382	43	90%
Form	commerce: shop	156	137	19	88%
Form	commerce: store	77	67	10	87%
Form	authority: usa	62	55	7	89%
Form	commerce: sale	46	36	10	78%
Form	authority: official	34	32	2	94%
Form	commerce: outlet	26	19	7	73%

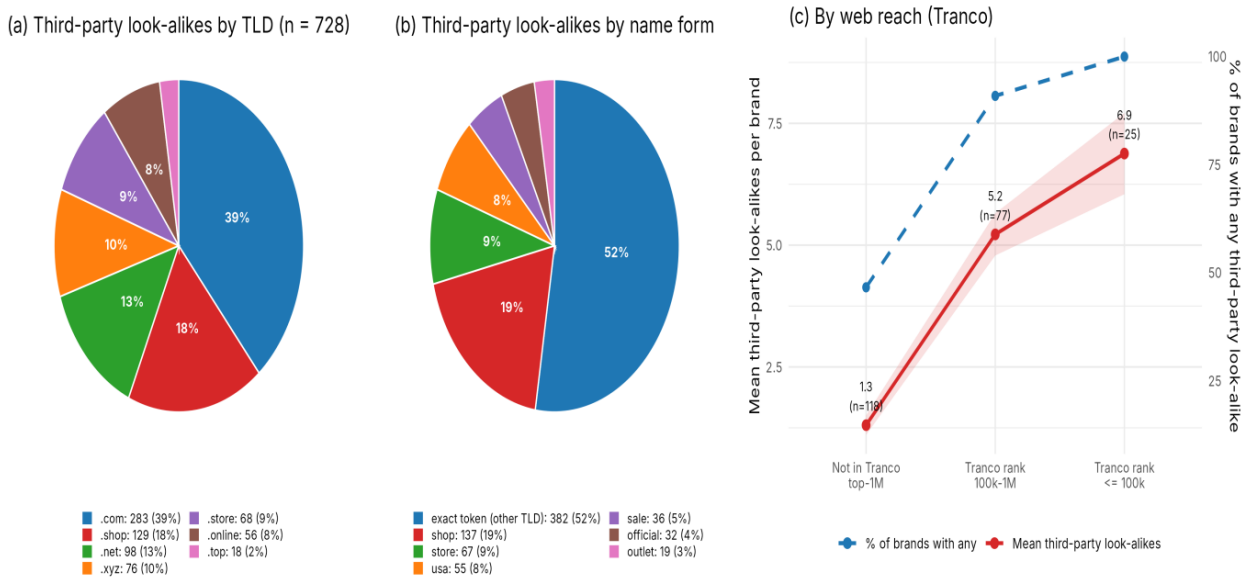


Figure 4. Third-party look-alikes by TLD (a), name form (b) and brand reach (c).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Model Selection

Table 9 compares specifications on the 217 complete cases. The Poisson model is over-dispersed (Pearson dispersion 2.84; z = 4.16, p < .001); the negative binomial raises the log-likelihood by 55.4 ($\chi^2(1) = 110.8$; $\theta = 2.01$) with Pearson dispersion 1.31. Adding reach and controls improves fit ($\chi^2(6) = 96.1$, p < .001; McFadden R^2 from 0.024 to 0.120); adding the cue block does not ($\chi^2(3) = 1.50$, p = .68). The hurdle and zero-inflated models obtain lower AIC (885.5 and 876.0 against 913.0) because the NB model under-predicts zeros (0.261 against 0.323 observed; Figure S2), but BIC-corrected Vuong statistics favour the NB model (z = 1.47 and 0.74) and the hurdle decomposition (Table 11) carries the same message.

Table 8 Who Signals? Cue-Adoption Models (HC0 robust SE)

Predictor	**C1 Cue index (rate ratio)**	**C2 DMARC enforced (OR)**	**C3 TikTok verified (OR)**	**C4 Defensive registrations (IRR)**
Web reach (Tranco, log10 scale)	1.62 [1.41, 1.86]***	13.56 [5.11, 36.02]***	15.90 [4.79, 52.71]***	2.62 [1.20, 5.70]*
TikTok followers (log1p)	1.04 [1.03, 1.06]***	1.00 [0.93, 1.08]	—	1.05 [0.98, 1.12]
Wikipedia article (0/1)	1.14 [0.98, 1.33]	1.28 [0.57, 2.88]	1.76 [0.71, 4.31]	2.79 [1.21, 6.48]*
Brand age (log years)	—	—	—	—
Parent company recorded (0/1)	1.14 [1.01, 1.30]*	1.51 [0.75, 3.05]	2.06 [0.92, 4.65]	1.68 [0.81, 3.49]
Fragrance (vs cosmetics)	0.85 [0.70, 1.04]	0.36 [0.12, 1.07]	0.18 [0.05, 0.61]**	1.53 [0.60, 3.92]
Hair, grooming & oral care	1.05 [0.89, 1.23]	0.79 [0.26, 2.37]	0.49 [0.15, 1.56]	0.57 [0.22, 1.48]
Retail & direct selling	0.79 [0.66, 0.93]**	0.53 [0.15, 1.83]	0.07 [0.02, 0.33]***	0.89 [0.22, 3.54]
Skin, bath & body	1.25 [1.05, 1.50]*	1.64 [0.59, 4.53]	0.45 [0.13, 1.51]	1.53 [0.60, 3.92]
n	240	240	202	235

Note. C1 Poisson model of the 0–9 cue index; C3 brands with a TikTok account. Reference category cosmetics. * p < .05, ** p < .01, *** p < .001.

Reach, cues and impersonation

Table 10 and Figure 5 report the estimates; target attractiveness dominates. In M2, each log10 unit of Tranco reach multiplies the expected number of third-party look-alikes by 2.67 (95% CI [1.77, 4.03]) and each log unit of domain age by 1.47 [1.09, 1.98]; each character of token length reduces it by 7% (IRR 0.93 [0.90, 0.97]); followers, a Wikipedia article, a parent company and category are not independently associated. The unconditional association between cues and impersonation is strongly positive (Spearman $\rho = 0.60$ between the cue index and the third-party count; 0.67 for reach; Figure S3), yet conditional on reach it disappears: the email-authentication index (IRR 1.07 [0.90, 1.26]), the domain-protection index (0.92 [0.78, 1.10]) and TikTok verification (1.22 [0.69, 2.14]) are all null, and in M3 no individual cue is significant (SPF 1.64 [0.93, 2.90]; DMARC enforced 0.99; BIMI 1.09; DNSSEC 1.14; CAA 1.25; registry lock 0.87; corporate registrar 0.81 [0.57, 1.15]; verified badge 1.36). Defensive registrations are unrelated to third-party ones (M4: 0.97 [0.91, 1.05]). The pattern supports H2b over H2a: no cue, singly or bundled, is associated with fewer third-party look-alikes, and the email-authentication and platform-verification estimates lie above one.

The logistic model reproduces the pattern for whether a brand is targeted at all (reach OR 26.2 [2.41, 284.5]; cues not significant), discriminates well (AUC 0.901 [0.860, 0.943]; Nagelkerke $R^2 = 0.563$) and is calibrated across risk quintiles. The hurdle model (Table 11) assigns reach to both stages (zero hurdle OR 26.2; count

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

part IRR 2.31 [1.52, 3.51]) with null cue coefficients in both. Generalised VIFs are at most 4.61; diagnostics (Figure S1) show a moderate right tail (15 brands with Cook's D above 4/n; maximum 0.136).

Table 9 Model Comparison (n = 217)

Model	k	Log-lik.	AIC	BIC	McFadden R^2	Pearson disp.	Pred. zero share	θ
Poisson (M2 terms)	14	-496.9	1021.8	1069.1	0.29	2.84	0.157	—
M0 NB category only	6	-490.3	992.5	1012.8	0.024	1.09	0.281	0.869
M1 NB reach + controls	12	-442.2	908.5	949	0.12	1.27	0.261	1.982
M2 NB + cue indices	15	-441.5	913	963.7	0.121	1.31	0.261	2.015
M3 NB + individual cues	20	-438.5	917.1	984.7	0.127	1.29	0.263	2.082
M4 NB + defensive	16	-441.3	914.6	968.7	0.122	1.32	0.261	2.027
M5 ZINB (M2 terms)	29	-409	876	974	0.186	1.44	0.321	4.02
M6 Hurdle-NB (M2 terms)	29	-413.8	885.5	983.5	0.176	1.11	0.323	3.96

Note. Observed zero share 0.323. Vuong NB vs hurdle: raw z = −4.11, BIC-corrected 1.47; NB vs ZINB: raw −4.59, BIC-corrected 0.74 (positive favours NB).

Table 10 Negative Binomial and Logistic Models of Third-Party Look-Alike Registrations (n = 217; HC0 robust SE)

Predictor	M1 NB: IRR	M2 NB: IRR	M3 NB: IRR	M4 NB: IRR	L1 Logistic: OR
Web reach (Tranco, log10 scale)	2.91 [1.94, 4.36]***	2.67 [1.77, 4.03]***	2.51 [1.67, 3.79]***	2.69 [1.78, 4.07]***	26.20 [2.41, 284.49]**
TikTok followers (log1p)	1.05 [1.01, 1.09]*	1.03 [0.98, 1.09]	1.02 [0.97, 1.08]	1.03 [0.98, 1.09]	0.96 [0.85, 1.09]
Wikipedia article (0/1)	0.93 [0.65, 1.32]	0.96 [0.67, 1.38]	1.00 [0.69, 1.46]	0.98 [0.68, 1.40]	0.81 [0.29, 2.22]
Domain age (log years)	1.47 [1.09, 1.99]*	1.47 [1.09, 1.98]*	1.48 [1.08, 2.03]*	1.48 [1.09, 1.99]*	1.90 [1.02, 3.52]*
Parent company recorded (0/1)	0.97 [0.76, 1.24]	0.94 [0.74, 1.21]	1.01 [0.76, 1.33]	0.94 [0.74, 1.21]	1.18 [0.50, 2.79]
Brand-token length (characters)	0.93 [0.90, 0.97]***	0.93 [0.90, 0.97]***	0.93 [0.89, 0.97]***	0.93 [0.89, 0.97]***	0.88 [0.79, 0.97]*
Email-authentication index (0–4)	—	1.07 [0.90, 1.26]	—	1.07 [0.90, 1.27]	1.18 [0.77, 1.80]
Domain-protection index (0–5)	—	0.92 [0.78, 1.10]	—	0.92 [0.77, 1.09]	2.03 [0.64, 6.37]
TikTok verified badge (0/1)	—	1.22 [0.69, 2.14]	1.36 [0.75, 2.45]	1.23 [0.70, 2.15]	3.84 [0.63, 23.51]
SPF record (0/1)	—	—	1.64 [0.93, 2.90]	—	—
DMARC enforced (0/1)	—	—	0.99 [0.70, 1.40]	—	—
BIMI record (0/1)	—	—	1.09 [0.77, 1.54]	—	—

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

DNSSEC (0/1)	—	—	1.14 [0.50, 2.59]	—	—
CAA record (0/1)	—	—	1.25 [0.55, 2.83]	—	—
Registry lock (0/1)	—	—	0.87 [0.62, 1.21]	—	—
Corporate registrar (0/1)	—	—	0.81 [0.57, 1.15]	—	—
Defensive registrations (count)	—	—	—	0.97 [0.91, 1.05]	—
Fragrance (vs cosmetics)	0.78 [0.50, 1.22]	0.82 [0.51, 1.31]	0.80 [0.49, 1.30]	0.83 [0.51, 1.35]	0.54 [0.13, 2.27]
Hair, grooming & oral care	0.89 [0.60, 1.31]	0.90 [0.61, 1.32]	0.89 [0.62, 1.26]	0.89 [0.60, 1.32]	0.38 [0.08, 1.76]
Retail & direct selling	0.88 [0.53, 1.45]	0.90 [0.55, 1.48]	0.92 [0.55, 1.52]	0.90 [0.55, 1.47]	1.52 [0.15, 15.25]
Skin, bath & body	1.24 [0.89, 1.71]	1.25 [0.88, 1.78]	1.24 [0.82, 1.88]	1.26 [0.89, 1.80]	1.27 [0.28, 5.78]

Note. 95% CI in brackets; reference category cosmetics. * p < .05, ** p < .01, *** p < .001.

Table 11 Hurdle Negative Binomial Decomposition (M2 terms; n = 217)

Predictor	**Zero hurdle: any look-alike (OR)**	**Count part: how many (IRR)**
Web reach (Tranco, log10 scale)	26.20 [2.41, 284.49]**	2.31 [1.52, 3.51]***
TikTok followers (log1p)	0.96 [0.85, 1.09]	1.06 [1.01, 1.12]*
Wikipedia article (0/1)	0.81 [0.29, 2.22]	0.95 [0.69, 1.30]
Brand age (log years)	—	—
Parent company recorded (0/1)	1.18 [0.50, 2.79]	1.01 [0.80, 1.28]
Brand-token length (characters)	0.88 [0.79, 0.97]*	0.95 [0.92, 0.98]**
Email-authentication index (0–4)	1.18 [0.77, 1.80]	1.00 [0.84, 1.18]
Domain-protection index (0–5)	2.03 [0.64, 6.37]	0.92 [0.80, 1.05]
TikTok verified badge (0/1)	3.84 [0.63, 23.51]	0.76 [0.44, 1.32]

Note. Zero hurdle: logit for any look-alike; count part: truncated NB among brands with at least one.

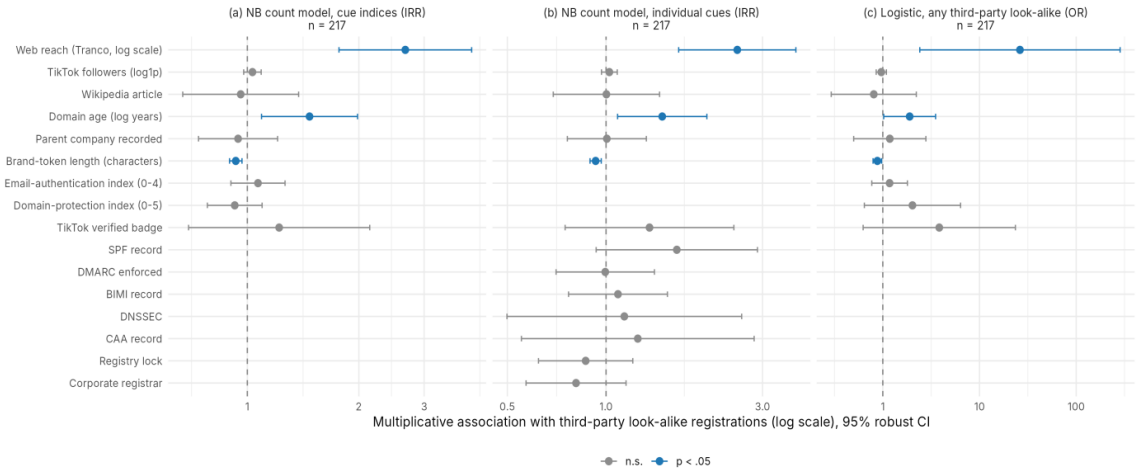


Figure 5. Associations with third-party look-alike registrations (95% robust CIs).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Robustness, Hypothesis 3 and Sample Selection

Table 12 summarises robustness checks. Reach remains the dominant association when influential brands or the top 1% of counts are excluded (IRR 3.15 and 2.48), when generic-token brands are added (2.27, n = 235), with a Tranco-ranked indicator (2.53) or a TikTok account indicator (2.73), for resolving names (2.54), for all registrations including defensive ones (2.60), for exact-token registrations (1.90) and with brand age in place of domain age (2.63). The cue indices are never negative and significant; excluding influential cases turns the email-authentication index positive (1.23 [1.05, 1.44]), and TikTok verification reaches the boundary of significance in the same direction with an account indicator (1.57 [1.00, 2.46]). Supporting H3, the reach gradient is steepest for commerce-keyword names (IRR 3.64 [2.12, 6.26]; logistic OR 9.99), where followers also become significant (1.09 [1.00, 1.18]), and shallowest for exact-token names (1.90). The 25 held-out brands do not differ from the 220 analysed on reach, age, parent ownership or cue adoption (Table 13; standardised mean differences below 0.4; all p > .10) except, by construction, on token length. Tranco coverage rises with followers (OR 1.17 per log unit), domain age (3.33) and retail category (10.4); unranked brands are treated as below the top million rather than as missing, and the ranked indicator yields the same conclusions.

Table 12 Robustness Checks (IRR or OR [95% CI])

Specification	n	Web reach	TikTok followers	Email-auth index	Domain-protection index	TikTok verified
R1 NB excluding Cook's D > 4/n	202	3.15 [2.20, 4.51]***	1.00 [0.96, 1.05]	1.23 [1.05, 1.44]*	0.92 [0.81, 1.05]	1.32 [0.80, 2.19]
R2 NB excluding top 1% counts	214	2.48 [1.64, 3.75]***	1.03 [0.97, 1.08]	1.06 [0.90, 1.25]	0.96 [0.82, 1.13]	1.21 [0.72, 2.01]
R3 NB including generic-token brands	235	2.27 [1.53, 3.38]***	1.02 [0.96, 1.08]	1.02 [0.88, 1.18]	0.95 [0.79, 1.15]	1.28 [0.70, 2.33]
R4 NB Tranco-ranked indicator instead of continuous reach	217	2.53 [1.70, 3.77]*** (ranked 0/1)	1.03 [0.98, 1.09]	1.13 [0.95, 1.33]	0.91 [0.74, 1.12]	1.24 [0.70, 2.17]
R5 NB TikTok account indicator instead of followers	217	2.73 [1.80, 4.12]***	1.21 [0.76, 1.92] (account 0/1)	1.07 [0.90, 1.26]	0.91 [0.77, 1.08]	1.57 [1.00, 2.46]
R6 NB commerce-keyword look-alikes only (shop/store/outlet/sale)	217	3.64 [2.12, 6.26]***	1.09 [1.00, 1.18]*	1.19 [0.89, 1.58]	0.90 [0.73, 1.12]	1.03 [0.43, 2.47]
R7 NB third-party look-alikes with an A record	217	2.54 [1.65, 3.91]***	1.03 [0.97, 1.09]	1.07 [0.90, 1.27]	0.87 [0.72, 1.06]	1.25 [0.70, 2.23]
R8 NB all registered look-alikes (incl. defensive)	217	2.60 [1.75, 3.86]***	1.03 [0.98, 1.09]	1.09 [0.93, 1.29]	0.91 [0.77, 1.08]	1.26 [0.73, 2.19]
R9 NB exact-token registrations in other TLDs	217	1.90 [1.33, 2.73]***	1.00 [0.96, 1.05]	1.08 [0.92, 1.26]	0.96 [0.82, 1.12]	1.09 [0.65, 1.83]
R10 NB brand age (inception) instead of domain age	174	2.63 [1.65, 4.19]***	1.04 [0.97, 1.11]	1.06 [0.88, 1.27]	0.89 [0.75, 1.05]	1.23 [0.60, 2.52]
R11 Logit any commerce-keyword look-alike	217	9.99 [2.60, 38.47]***	1.09 [0.95, 1.23]	1.59 [0.96, 2.63]	1.66 [0.84, 3.26]	1.14 [0.27, 4.80]

Note. * p < .05, ** p < .01, *** p < .001; — not in specification.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Table 13 Analysis Sample Versus Held-Out Brands

Variable	Analysis sample mean	Excluded brands mean	Std. mean diff.	p
Web reach (Tranco)	0.36	0.41	-0.1	0.924
Tranco-ranked	0.46	0.48	-0.03	1
TikTok followers	491183.91	72799.56	0.18	0.101
Wikipedia article	0.43	0.44	-0.03	1
Domain age (years)	17.9	19.18	-0.12	0.447
Brand age (years)	32.99	37.41	-0.11	0.881
Parent company	0.35	0.36	-0.01	1
Verification-cue index	2.16	1.88	0.2	0.394
DMARC enforced	0.48	0.44	0.07	0.887
DNSSEC	0.06	0.04	0.07	1
TikTok verified	0.33	0.16	0.39	0.137
Token length	10.89	4.32	1.99	< .001

Note. p: Wilcoxon or χ^2 tests.

DISCUSSION

Theoretical Contributions

Three findings organise the discussion. First, impersonation infrastructure tracks target attractiveness: reach and domain age predict whether and how often a brand's name is registered by others, commerce variants show the steepest gradient, and two-thirds of brands already carry such names, most of which resolve. This extends combosquatting measurement (Kintis et al., 2017; Tian et al., 2018) from the top sites of the web to the brands of one consumer market: the targets of social-commerce impersonation are the brands social commerce rewards. Second, authenticity signals are produced by the brands with the most reputation to defend: adoption of every costly cue rises steeply with reach, consistent with the economics of security investment (Anderson & Moore, 2006) and with signalling theory's prediction that signals are sent where the reputation at stake justifies their cost (Kirmani & Rao, 2000; Spence, 1973). Third, conditional on reach, no cue is associated with fewer third-party look-alikes, and the unconditional association is positive. The deterrence account (H2a) is not supported; the target-attractiveness account (H2b) is.

We read this as evidence that verification cues operate at the consumer-inference stage rather than the attacker-supply stage. DMARC enforcement stops spoofed mail, a registry lock stops a hijack, and a verified badge helps a consumer separate brand from clone (Hu et al., 2018; Vaidya et al., 2019), but none of them changes the cost of registering neutrogena-shop.com or the pay-off of an advertisement that points to it. Signals of identity are therefore necessary for consumers but insufficient against impersonation, and the volume of impersonation infrastructure does not measure their failure. Reverse causality and confounding remain plausible: brands under attack may adopt cues, and unobserved marketing intensity may drive both reach and adoption. The design cannot separate these mechanisms from a true null deterrent effect, and no association is causal; what it establishes is that cue adoption and impersonation share a common driver in reach, so any evaluation of cue effectiveness that ignores reach will find a spurious positive association. It follows that no single cue should be presented to consumers or policy makers as protective on its own: each closes one channel (spoofed mail, domain hijack or account confusion) and leaves the others open.

The findings also expose a distributional problem. Social commerce multiplies small brands, and the majority of brands here, those outside the Tranco top million, enforce DMARC in fewer than three cases in ten and hold a verified badge in fewer than one in ten. For a consumer who meets an unfamiliar brand in a feed, the absence of a cue is therefore uninformative: it is equally consistent with an impersonator and with a legitimate

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

small brand that has not invested. A separating equilibrium requires that the honest type signal, and in this market it mostly does not. Under low elaboration, consumers may also substitute peripheral cues that impersonators can copy, such as logos, product imagery and a brand name in the domain, for the costly cues that they cannot (Chaiken, 1980; Petty & Cacioppo, 1986), so the signals that do separate the honest type are the ones least likely to be read at the moment of purchase. Research on data integrity in social platforms (Oladoyinbo et al., 2024), on fraud detection under adversarial adaptation (Adesokan et al., 2026) and on secure deployment of the generative tools that make convincing clones cheap (Ejoh et al., 2026) points the same way: authenticity is a property of the whole system, not of any single control.

Implications for Brands, Platforms, Regulators and Consumers

Figure 6 and Table 14 present the Authenticity Signalling Framework. The brand layer produces signals: DMARC at reject with BIMI, DNSSEC, CAA and a registry lock on the official domain, verified accounts on every commerce platform used, defensive registration of commerce look-alikes, and one canonical list of official domains and accounts. The platform layer transmits them: verification and seller status shown at the point of purchase and in advertisements, clear labels for brand-owned and authorised storefronts, identity verification before listings and advertisements go live, detection of handles and link domains that mimic verified brands, and rapid take-down channels, extending the seller-verification duties the INFORM Consumers Act imposes for high-volume sellers (INFORM Consumers Act, 2022). The regulator and registry layer enforces them through the Impersonation Rule (Trade Regulation Rule on Impersonation of Government and Businesses, 2024), registrar abuse reporting for brand-keyword registrations, which can be generated and checked automatically, and reporting data that separate impersonation from other scams. The consumer layer receives them: treat the badge, the official domain and the checkout as a bundle, distrust official, outlet and sale name variants, and navigate from the brand's canonical site to its shop rather than from the advertisement. Consumer education should teach this bundle rather than any single badge, because no single cue is protective on its own.



Figure 6. The Authenticity Signalling Framework.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

Table 14 Framework Layers, Controls and Empirical Anchors

Layer	Core controls	Empirical anchor (this study)
Brand: signal production	DMARC at reject with BIMI; DNSSEC, CAA, registry lock; verified commerce accounts; defensive commerce look-alikes; canonical list of official domains and accounts	47.3% enforce DMARC, 4.9% BIMI, 5.3% DNSSEC, 8.6% registry lock, 31.0% TikTok-verified; adoption rises 1.62× per reach unit
Platform: signal transmission	Verification and seller status at the point of purchase and in ads; brand-owned storefront labels; identity checks before listings; look-alike handle and link detection; take-down channels; INFORM Act disclosures	9.2% of brands outside the Tranco top million are TikTok-verified; 84.5% have an account
Regulator and registry: signal enforcement	Impersonation Rule enforcement; registrar abuse reporting for brand-keyword registrations; impersonation-specific reporting data	728 third-party look-alikes for 220 brands; 87–94% of shop, store, usa and official variants held by third parties
Consumer: signal reception	Treat badge, official domain and checkout as a bundle; distrust official, outlet and sale variants; navigate from the canonical site; report impersonation	No cue is associated with fewer look-alikes (IRR 0.92–1.22, n.s.)

LIMITATIONS AND FUTURE RESEARCH

The design is cross-sectional and observational. Impersonation infrastructure is measured by generated names in seven top-level domains with six keywords, which omits typosquats, homographs and other keywords and domains, and ownership attribution rests on name-server and address overlap, which can misclassify a defensive registration hosted elsewhere as third-party or a squatter on the brand's provider as defensive; the strict and all-registered variants bound this error. Registration is not fraud: parked and resolving names were distinguished, but content and use were not examined, and social-media clones and advertisements were not counted. TikTok attribution relies on recorded or exact-name handles, so verified accounts under other names were missed, and reach is measured only within the Tranco top million. Platform verification was observed on TikTok only, so differences across Instagram, Facebook and other commerce surfaces, including platforms on which verification can be bought by subscription, are not captured. Consumer-level factors that shape susceptibility, such as age, digital literacy, prior brand familiarity and price sensitivity, and the persuasive tactics that accompany impersonation, such as discounted prices, influencer endorsement and urgency, lie outside a brand-level design; web reach and a Wikipedia article proxy brand familiarity only at the population level. The data are a snapshot of 3 October 2026, and because scam tactics and platform verification policies change quickly, the levels reported here may date even if the reach gradient does not. The frame also favours brands with an encyclopaedic footprint, so very small and newly launched social-commerce brands, which the findings suggest are the least likely to signal, are under-represented, and the adoption rates reported here are likely upper bounds for the market as a whole. Future research should link look-alike registrations to advertising libraries, take-down records and consumer complaints, follow brands over time around cue adoption, and test experimentally which cues consumers use at the point of purchase.

Three designs would address the consumer side directly. First, experiments or choice-based conjoint studies could present genuine and counterfeit versions of each cue, individually and in combination and across platforms. Second, such studies should measure trust, perceived authenticity, purchase intention and, where possible, scam-detection accuracy or observed purchasing, using validated scales with reported reliability and validity (e.g., McKnight et al., 2002). Third, they should test moderation by age, digital literacy, brand familiarity and price sensitivity, with samples drawn beyond convenience online panels to include older, lower-income and non-English-speaking consumers. Linking brand-level exposure to FTC and BBB

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

complaints that name the impersonated brand would connect the supply of impersonation infrastructure to realised victimisation, and panel data collected around cue adoption would support the longitudinal, cross-platform analysis that a single snapshot cannot.

CONCLUSION

Social commerce has made the brand an interface anyone can copy. Across 245 U.S. beauty and personal-care brands, verification cues that impersonators cannot replicate are published mainly by the brands with the greatest reach, and two-thirds of brands already have third-party look-alike domains that follow the same gradient. Conditional on reach, no cue is associated with less impersonation infrastructure. These associations are brand-level and observational; they do not show that cues fail to protect consumers, and they give no grounds for treating any single cue as protective on its own. Authenticity signals protect the consumer who knows to read them, not the brand from being copied, and the brands least likely to send them are the small brands that social commerce creates. Reducing brand-impersonation scams therefore requires the layered framework proposed here: brands that signal, platforms that surface the signal where the purchase happens, regulators and registries that act on the look-alike supply, and consumers who check the domain rather than the logo. Longitudinal, cross-platform and experimental research is needed to establish which of these signals consumers actually read at the point of purchase.

REFERENCES

1. Adesokan, O. Z., Bamsa, A. O., Obioha-Val, O. A., Ejoh, C. I., & Akeke, M. A. (2026). Adversarial robustness of machine learning-based fraud detection systems: An empirical evaluation of attack impact and mitigation in fintech environments. Journal of Engineering Research and Reports, 28(7), 261–279. https://doi.org/10.9734/jerr/2026/v28i71957
2. Agten, P., Joosen, W., Piessens, F., & Nikiforakis, N. (2015). Seven months' worth of mistakes: A longitudinal study of typosquatting abuse. In Proceedings 2015 Network and Distributed System Security Symposium. Internet Society. https://doi.org/10.14722/ndss.2015.23058
3. Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. The Quarterly Journal of Economics, 84(3), 488–500. https://doi.org/10.2307/1879431
4. Anderson, R., & Moore, T. (2006). The economics of information security. Science, 314(5799), 610–613. https://doi.org/10.1126/science.1130992
5. BBB Institute for Marketplace Trust. (2025). 2024 BBB Scam Tracker risk report. https://BBBMarketplaceTrust.org/RiskReport
6. Bitaab, M., Cho, H., Oest, A., Lyu, Z., Wang, W., Abraham, J., Wang, R., Bao, T., Shoshitaishvili, Y., & Doupé, A. (2023). Beyond phish: Toward detecting fraudulent e-commerce websites at scale. In 2023 IEEE Symposium on Security and Privacy (SP) (pp. 2566–2583). IEEE. https://doi.org/10.1109/SP46215.2023.10179461
7. Bitaab, M., Karimi, A., Lyu, Z., Oest, A., Kuchhal, D., Saad, M., Ahn, G.-J., Wang, R., Bao, T., Shoshitaishvili, Y., & Doupé, A. (2025). ScamMagnifier: Piercing the veil of fraudulent shopping website campaigns. In Proceedings 2025 Network and Distributed System Security Symposium. Internet Society. https://doi.org/10.14722/ndss.2025.240763
8. Button, M., Lewis, C., & Tapley, J. (2014). Not a victimless crime: The impact of fraud on individual victims and their families. Security Journal, 27(1), 36–54. https://doi.org/10.1057/sj.2012.11
9. Cameron, A. C., & Trivedi, P. K. (2013). Regression analysis of count data (2nd ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139013567
10. Carpineto, C., & Romano, G. (2017). Learning to detect and measure fake ecommerce websites in search-engine results. In Proceedings of the International Conference on Web Intelligence (pp. 403–410). ACM. https://doi.org/10.1145/3106426.3106441
11. Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. Journal of Personality and Social Psychology, 39(5), 752–766. https://doi.org/10.1037/0022-3514.39.5.752

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

12. Connelly, B. L., Certo, S. T., Ireland, R. D., & Reutzel, C. R. (2011). Signaling theory: A review and assessment. Journal of Management, 37(1), 39–67. https://doi.org/10.1177/0149206310388419
13. Dhamija, R., Tygar, J. D., & Hearst, M. (2006). Why phishing works. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 581–590). ACM. https://doi.org/10.1145/1124772.1124861
14. Diamantopoulos, A., & Winklhofer, H. M. (2001). Index construction with formative indicators: An alternative to scale development. Journal of Marketing Research, 38(2), 269–277. https://doi.org/10.1509/jmkr.38.2.269.18845
15. Durumeric, Z., Adrian, D., Mirian, A., Kasten, J., Bursztein, E., Lidzborski, N., Thomas, K., Eranti, V., Bailey, M., & Halderman, J. A. (2015). Neither snow nor rain nor MITM…: An empirical analysis of email delivery security. In Proceedings of the 2015 Internet Measurement Conference (pp. 27–39). ACM. https://doi.org/10.1145/2815675.2815695
16. Edgerly, S., & Vraga, E. K. (2019). The blue check of credibility: Does account verification matter when evaluating news on Twitter? Cyberpsychology, Behavior, and Social Networking, 22(4), 283–287. https://doi.org/10.1089/cyber.2018.0475
17. Ejiofor, V. O. (2025). Open data for cyber resilience: An analysis of public-private collaboration in AI-supported threat intelligence sharing. Asian Journal of Research in Computer Science, 18(6), 403–423. https://doi.org/10.9734/ajrcos/2025/v18i6710
18. Ejoh, C. I., Akeke, C. U., Oladoyinbo, O. B., Henry, O., & Archibong, U. N. (2026). Ethical and secure deployment of generative AI: Balancing innovation, data privacy, and enterprise risk governance. Journal of Engineering Research and Reports, 28(7), 48–68. https://doi.org/10.9734/jerr/2026/v28i71943
19. Federal Trade Commission. (2024, April 1). Impersonation scams: Not what they used to be [Consumer Protection Data Spotlight]. https://www.ftc.gov/news-events/data-visualizations/data-spotlight/2024/04/impersonation-scams-not-what-they-used-be
20. Federal Trade Commission. (2026, April 27). Reported losses to scams on social media eight times higher than in 2020 [Consumer Protection Data Spotlight]. https://www.ftc.gov/news-events/data-visualizations/data-spotlight/2026/04/reported-losses-scams-social-media-eight-times-higher-2020
21. Gbadebo, O. A., & Gbadebo, M. O. (2025). Strengthening cyber resilience: Policy and governance responses to global cyber threats. Asian Journal of Research in Computer Science, 18(10), 183–198. https://doi.org/10.9734/ajrcos/2025/v18i10772
22. Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust and TAM in online shopping: An integrated model. MIS Quarterly, 27(1), 51–90. https://doi.org/10.2307/30036519
23. Hu, H., Peng, P., & Wang, G. (2018). Towards understanding the adoption of anti-spoofing protocols in email systems. In 2018 IEEE Cybersecurity Development (SecDev) (pp. 94–101). IEEE. https://doi.org/10.1109/SecDev.2018.00020
24. Hu, X., Wu, G., Wu, Y., & Zhang, H. (2010). The effects of Web assurance seals on consumers' initial trust in an online vendor: A functional perspective. Decision Support Systems, 48(2), 407–418. https://doi.org/10.1016/j.dss.2009.10.004
25. Integrity, Notification, and Fairness in Online Retail Marketplaces for Consumers Act, Pub. L. No. 117-328, div. BB, tit. III, 15 U.S.C. § 45f (2022). https://www.ftc.gov/legal-library/browse/statutes/inform-consumers-act
26. Keller, K. L. (1993). Conceptualizing, measuring, and managing customer-based brand equity. Journal of Marketing, 57(1), 1–22. https://doi.org/10.1177/002224299305700101
27. Kintis, P., Miramirkhani, N., Lever, C., Chen, Y., Romero-Gómez, R., Pitropakis, N., Nikiforakis, N., & Antonakakis, M. (2017). Hiding in plain sight: A longitudinal study of combosquatting abuse. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 569–586). ACM. https://doi.org/10.1145/3133956.3134002
28. Kirmani, A., & Rao, A. R. (2000). No pain, no gain: A critical review of the literature on signaling unobservable product quality. Journal of Marketing, 64(2), 66–79. https://doi.org/10.1509/jmkg.64.2.66.18000

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

29. Le Pochat, V., Van Goethem, T., Tajalizadehkhoob, S., Korczyński, M., & Joosen, W. (2019). Tranco: A research-oriented top sites ranking hardened against manipulation. In Proceedings 2019 Network and Distributed System Security Symposium. Internet Society. https://doi.org/10.14722/ndss.2019.23386
30. Liang, T.-P., & Turban, E. (2011). Introduction to the special issue social commerce: A research framework for social commerce. International Journal of Electronic Commerce, 16(2), 5–14. https://doi.org/10.2753/JEC1086-4415160201
31. Lin, X., Li, Y., & Wang, X. (2017). Social commerce research: Definition, research themes and the trends. International Journal of Information Management, 37(3), 190–201. https://doi.org/10.1016/j.ijinfomgt.2016.06.006
32. Mariconti, E., Onaolapo, J., Ahmad, S. S., Nikiforou, N., Egele, M., Nikiforakis, N., & Stringhini, G. (2017). What's in a name? Understanding profile name reuse on Twitter. In Proceedings of the 26th International Conference on World Wide Web (pp. 1161–1170). ACM. https://doi.org/10.1145/3038912.3052589
33. McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. Information Systems Research, 13(3), 334–359. https://doi.org/10.1287/isre.13.3.334.81
34. Metzger, M. J., & Flanagin, A. J. (2013). Credibility and trust of information in online environments: The use of cognitive heuristics. Journal of Pragmatics, 59, 210–220. https://doi.org/10.1016/j.pragma.2013.07.012
35. Moore, T., & Edelman, B. (2010). Measuring the perpetrators and funders of typosquatting. In R. Sion (Ed.), Financial cryptography and data security (Lecture Notes in Computer Science, Vol. 6052, pp. 175–191). Springer. https://doi.org/10.1007/978-3-642-14577-3_15
36. Morris, M. R., Counts, S., Roseway, A., Hoff, A., & Schwarz, J. (2012). Tweeting is believing? Understanding microblog credibility perceptions. In Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (pp. 441–450). ACM. https://doi.org/10.1145/2145204.2145274
37. Oladoyinbo, T. O., Olabanji, S. O., Olaniyi, O. O., Adebiyi, O. O., Okunleye, O. J., & Alao, A. I. (2024). Exploring the challenges of artificial intelligence in data integrity and its influence on social dynamics. Asian Journal of Advanced Research and Reports, 18(2), 1–23. https://doi.org/10.9734/ajarr/2024/v18i2601
38. Olaniyi, O. O., Olabanji, S. O., & Abalaka, A. I. (2023). Navigating risk in the modern business landscape: Strategies and insights for enterprise risk management implementation. Journal of Scientific Research and Reports, 29(9), 103–109. https://doi.org/10.9734/jsrr/2023/v29i91789
39. Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. Advances in Experimental Social Psychology, 19, 123–205. https://doi.org/10.1016/S0065-2601(08)60214-2
40. Raval, D. (2020). Whose voice do we hear in the marketplace? Evidence from consumer complaining behavior. Marketing Science, 39(1), 168–187. https://doi.org/10.1287/mksc.2018.1140
41. Roberts, R., Goldschlag, Y., Walter, R., Chung, T., Mislove, A., & Levin, D. (2019). You are who you appear to be: A longitudinal study of domain impersonation in TLS certificates. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (pp. 2489–2504). ACM. https://doi.org/10.1145/3319535.3363188
42. Spence, M. (1973). Job market signaling. The Quarterly Journal of Economics, 87(3), 355–374. https://doi.org/10.2307/1882010
43. Tian, K., Jan, S. T. K., Hu, H., Yao, D., & Wang, G. (2018). Needle in a haystack: Tracking down elite phishing domains in the wild. In Proceedings of the Internet Measurement Conference 2018 (pp. 429–442). ACM. https://doi.org/10.1145/3278532.3278569
44. Trade Regulation Rule on Impersonation of Government and Businesses, 16 C.F.R. Part 461, 89 Fed. Reg. 15017 (March 1, 2024) (effective April 1, 2024). https://www.govinfo.gov/content/pkg/FR-2024-03-01/html/2024-04335.htm
45. Vaidya, T., Votipka, D., Mazurek, M. L., & Sherr, M. (2019). Does being verified make you more credible? Account verification's effect on tweet credibility. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (pp. 1–13). ACM. https://doi.org/10.1145/3290605.3300755

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IX, September 2026

46. Vu, D. L., Pashchenko, I., Massacci, F., Plate, H., & Sabetta, A. (2020). Typosquatting and combosquatting attacks on the Python ecosystem. In 2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW) (pp. 509–514). IEEE. https://doi.org/10.1109/EuroSPW51379.2020.00074
47. White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica, 48(4), 817–838. https://doi.org/10.2307/1912934
48. Wilcox, K., Kim, H. M., & Sen, S. (2009). Why do consumers buy counterfeit luxury brands? Journal of Marketing Research, 46(2), 247–259. https://doi.org/10.1509/jmkr.46.2.247
49. Wilson, P. (2015). The misuse of the Vuong test for non-nested models to test for zero-inflation. Economics Letters, 127, 51–53. https://doi.org/10.1016/j.econlet.2014.12.029
50. Zeng, E., Kohno, T., & Roesner, F. (2021). What makes a "bad" ad? User perceptions of problematic online advertising. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (pp. 1–24). ACM. https://doi.org/10.1145/3411764.3445459