Linking Tech-Support Scam Lure Tactics and Brand Targets From Cybercriminal Marketplaces To Victim Complaints

Raghavendra Cherupalli
The University of Tulsa
Tulsa, USA
rac8609@utulsa.edu

Yi Ting Chua
The University of Tulsa
Tulsa, USA
ytc2805@utulsa.edu

Tyler Moore
The University of Tulsa
Tulsa, USA
tyler-moore@utulsa.edu

Gary Warner
Dark Tower
Birmingham, USA
gary.warner@getdarktower.com

Abstract

In a tech-support scam (TSS), cybercriminals dupe victims into paying fees by impersonating common interactions customers have with companies. This often includes interacting by phone with a fraudulent call center. Here, we empirically investigate which brands are impersonated and the type of lure used to initiate contact, such as an email, phone call, or pop-up window. We examine two distinct datasets. On the cybercriminal supply side, we study hundreds of thousands of posts on public Facebook groups dedicated to discussing TSS activities among scam operators. On the victimization side, we examine thousands of reports made directly by victims to the Better Business Bureau (BBB) Scam Tracker. We observe differences and commonalities in the prevalence of brand mentions on the supply and victim side. We also find evidence of correlation between the services discussed by criminals and the lures reported by victims. Finally, we describe logistic regressions which compare the victimization rate to the brand impersonated and lure type observed.

ACM Reference Format:
Raghavendra Cherupalli, Yi Ting Chua, Tyler Moore, and Gary Warner. 2026. Linking Tech-Support Scam Lure Tactics and Brand Targets From Cybercriminal Marketplaces To Victim Complaints. In *1st Workshop on Technical Analysis and Knowledge Exchange on Disrupting Online Criminal Networks (TAKEDOWN '26), November 15–19, 2026, The Hague, Netherlands.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3847159.3848616

1 Introduction

Technical-support scams (TSSs), sometimes known as call-center or impersonation scams, remains a persistent and growing concern since first appearing in 2012 [15]. In the United States, FBI Internet Crime Complaint Center (IC3) statistics show that TSS has consistently ranked among the five highest-loss cybercrime categories since 2018 and among the top three since 2023. Notably, reported complaints have increased by only about three times, whereas reported financial losses have risen by nearly forty times, suggesting that modern TSS campaigns are becoming substantially more sophisticated, scalable, and economically destructive [7].

One factor contributing to the persistence of TSSs is offenders' evolving tactics. Traditionally associated with aggressive browser locking and phone-based impersonation [23, 25], modern TSS operations have expanded their reach, complexity, and scale to lure potential victims into calling scammers, often through the use of fake invoices and subscription renewals [12]. Brand impersonation, via recognizable logos, user interface, or toll-free numbers [18, 21], is routinely used by scammers to lure prospective victims into initiating communication with TSS operators. Brands such as Amazon [24] and Microsoft [20] are commonly targeted by scammers.

Although considerable research has predominantly focused on detecting scam websites and identifying scam campaigns, relatively less attention has been paid to provide a comprehensive view of the scam. Understanding the intersection between cybercriminal behavior and victim experiences is essential for gaining a broader picture of how scam campaigns evolve, affects victims, and can be disrupted. To address this, we correlate the behavior observed within these groups with victim complaints reported to BBB Scam Tracker [2]. This study is guided by the following research questions:

RQ1: To what extent do targeted brands differ between informal supply-side ecosystems and formal consumer reporting platforms?
RQ2: To what extent does supply-side market activity map to the incident frequencies observed in formal consumer reporting platforms?
RQ3: How does the victimization rate vary by brand target and communication channels?

1.1 Key Contributions

This paper offers important empirical insights and methodological contributions to the study of TSS marketplaces:

- **Mapping of Criminal Marketplace to Victim Complaints**: We provide the first direct empirical comparison between an informal supply-side TSS marketplace and real-world victim complaints.

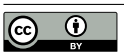

This work is licensed under a Creative Commons Attribution 4.0 International License.
TAKEDOWN '26, The Hague, Netherlands
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-3020-7/2026/11
https://doi.org/10.1145/3847159.3848616

- **Zero-shot Information Extraction Pipeline**: We design and validate a two-stage schema-based information extraction pipeline using the Gemma3:12B large language model (LLM) to identify brands and lures from victim reports.
- **Predictive Modeling of Victimization Risk**: Using logistic regression, we examined the relationship between commonly impersonated brands, lure methods and rate of victimization, demonstrating that some lure methods are more robust and consistent predictors of attack.

2 Related Work

We discuss prior work most relevant to the present study: measurement studies of tech-support scams, brand impersonation tactics used by cybercriminals, and cybercrime data extraction using LLMs.

2.1 Tech-Support Scams

The execution of TSS has undergone significant tactical evolution, transitioning from rudimentary cold-calling to complex, multi-vector ecosystems. Initially, TSS criminals mostly relied on cold-calling (outbound) techniques. Over time, the transformation adopted inbound calling techniques that leverage malvertising and counterfeit websites to dupe victims into initiating communication [21]. Building on earlier work, Srinivasan et al. [25] classified malicious TSS websites into two operational categories: aggressive and passive. Aggressive websites employ persistent pop-up messages and intimidating audio warnings to coerce user interaction, while passive websites mimic legitimate support portals through the use of corporate branding, trust seals, and black-hat search engine optimization (SEO) practices to increase their visibility and perceived legitimacy [25].

The infrastructure necessary to execute these diverse attack vectors is procured from thriving underground communities. Liu et al. [19] provided crucial findings by examining real-time communications within Facebook and WhatsApp groups, demonstrating how these platforms function as informal marketplaces for scammers to procure web development, toll-free numbers, and victim data. In response to this ecosystem, researchers have developed a range of automated detection pipelines, spanning machine learning systems capable of analyzing YouTube scam call recordings [28] to AI-driven platforms designed to identify TSS webpages [18]. Building on this work most recently, Cherupalli et al. [4] applied structured LLM prompting alongside few-shot in-context learning to accurately classify the illicit services being traded within these criminal Facebook groups. Furthermore, Cherupalli et al. analyzed the supply and demand signals for cybercrime services within these communities, shedding light on the underlying resource dependencies, and operational dynamics of the TSS marketplace [3].

2.2 Brand Impersonation in Cybercrime

Brand impersonation has evolved from simple deceptive messaging into a sophisticated, semi-automated underground industry. In the earlier era of instant messaging worms, attackers distributed malicious links through the compromised accounts of a victim's social contacts [22]. Moore and Clayton [22] identified that in pursuit of higher click-through rates, threat actors subsequently abandoned generic URL shorteners in favor of domains that deliberately incorporated recognizable brand names. More recently, as legitimate brands transitioned their customer service operations to social platforms like X, Instagram, and Telegram, scammers followed suit to intercept users actively seeking assistance. Acharya et al. [1] found that fraudsters have advanced their fake profiles using increasingly sophisticated username-squatting techniques, from basic typo-squatting (minor spelling errors) to combo-squatting (pairing brand names with terms like "support" or "helpdesk"). A novel evolution of this tactic is fuzzy-squatting, which merges combo-squatting with alphanumeric substitutions (e.g., amaz0n_h3lp_ds3k) to evade automated platform detection systems. Previous studies suggest that attackers deliberately impersonate well-known brands because familiar identities serve as powerful trust cues that increase the effectiveness of phishing and impersonation attacks. According to the 2026 Hoxhunt Phishing Trends Report [16], cybercriminals increased the use of artificial intelligence to generate highly contextualized content, such as fake invoices or expiring subscription notices, that mimic routine business communications.

2.3 Information Extraction Using Large Language Models

Recent advancements in LLMs have transformed information extraction from requiring heavily fine-tuned, task-specific models to using flexible, zero-shot approaches. Extraction aims to extract structural knowledge from plain natural language text, and the strong text understanding and generation capabilities of modern generative LLMs have motivated a growing body of work that integrates LLMs into information extraction pipelines under this generative paradigm rather than relying on traditional discriminative extraction methods [30]. A key factor driving this shift is the ability of LLMs to perform zero-shot and few-shot learning, enabling a single model to generalize across diverse information extraction tasks without task-specific fine-tuning while also supporting data augmentation and the execution of more complex tasks [29].

The methodology adopted in this study closely aligned with prior work on schema-guided information extraction using LLMs. Kaddari et al. [17] evaluated ChatGPT's ability to extract predefined clinical attributes from anonymized medical reports using both zero-shot prompting and a schema-guided extraction approach. Their results demonstrated that while the model performed well in structuring narrative records, its extraction reliability diminished when faced with complex, ambiguous, or non-standardized text. Similar challenges arise in the present study, as BBB complaint narratives are likewise unstructured, informal, and highly variable in writing style and completeness, requiring accurate mapping to a standardized schema for subsequent quantitative analysis.

The use of open-weight LLMs in cybersecurity has expanded considerably in recent years, with Gemma [26, 27] emerging as a promising model for a range of classification tasks. Cherupalli et al. [4] employed Gemma and Gemma 3 to automatically classify over 380,000 cybercriminal marketplace posts into predefined scam service and product categories, reporting high agreement with manually annotated labels . Building on this emerging body of work,

the present study employs Gemma 3:12B to perform structured information extraction from BBB complaint narratives.

3 Data Collection Methodology

The current study examines the degree of correspondence between the brands targeted within the informal supply-side scam ecosystem and those documented in formal consumer reporting platforms. To achieve this objective, we focus on the ten most frequently impersonated brands and organizations to determine whether the entities discussed within scam-supporting communities match those reported by victims in real scam incidents.

3.1 Scam Contributing Groups

To analyze the supply-side scam marketplace, this work utilizes the dataset introduced by Cherupalli et al. [4]. The dataset consists of 96 public Facebook groups that were either established to facilitate TSS or demonstrated sustained tolerance for exchange of illicit products, services and knowledge related to TSS operations. In total, the dataset includes 186,936 unique users who contributed 381,843 posts published between 2015 to 2024.

3.2 BBB Scam Tracker

The BBB Scam Tracker is a public, open-source platform through which individuals can report their experiences with fraudulent actors [2]. During the reporting process, users classify the type of scam they experienced from a predefined list of categories including but not limited to TSS, government impersonation, and fake invoice scams. In addition, users can disclose monetary loss incurred as a result of the scam encounter. Within the accompanying narrative descriptions, users frequently explain the manner in which the scammers initiated the lure, as well as the nature of their own responses.

A Python-based pipeline was developed to systematically collect and organize relevant data from the BBB Scam Tracker. For this study, the analysis was restricted to two specific categories: TSS and fake invoice scams. All TSS reports submitted between 2015 and 2024 were included in the dataset to maintain consistency with the Facebook group postings (2015 - 2024). In contrast, complaints in the fake invoice category within the same time period were identified and compiled through targeted keyword searches, based on the names of frequently impersonated brands and organizations. The list of targeted brands and organizations (provided in Appendix A) were identified through a systematic review of consumer protection advisories and fraud reports published by the U.S. Federal Trade Commission (FTC) [8–14]. A total of 10,778 BBB complaints were collected and analyzed as a part of this study.

4 Methodology

The BBB Scam Tracker contains narratives describing victims' experience with scammers. These narratives offer rich context but remain unstructured, making them unusable directly for quantitative analysis. To address this limitation and enable systematic analysis, we developed a two-stage zero-shot automated classification and information extraction pipeline.

4.1 Information Extraction Schema

We frame the task as a schema-based information extraction problem. Instead of treating each narrative as a single classification instance, our goal is to convert the unstructured complaints into structured representation by extracting a set of scam-specific attributes that describe the interaction between victim and the scammer. These attributes capture the communication mechanism used to initiate the scam, the manner in which the victim responded, the brand or organization impersonated by the scammer, and the direction of any phone call communication involved. The structured format enables large scale quantitative analysis and allows direct comparison with behavior patterns observed in scam-contributing Facebook groups.

The extraction scheme consists of 4 attributes: *Lure*, *Victim Response*, *Phone Call Direction* and *Brand Targeted*. Three among these (*Lure*, *Victim Response*, and *Phone Call Direction*) use predefined label sets, whereas *Brand Targeted* is extracted as the name of the impersonated organization.

Lure This attribute identifies the contact channel or method the scammer used to initiate contact with the victim, because TSS frequently employs multiple communication channels during a single attack. For example, a fraudulent email may instruct the victim to call a phone number, or a malicious pop-up notification leading to a follow-up call. As a result, **Lure** is treated as multi-label task allowing multiple channels per complaint. The predefined labels include email, SMS, phone call, pop-up notification, website and unknown, with the latter applied when the complaint lacks sufficient detail to determine how contact was initiated.

Victim Response This attribute records the outcome of victim's engagement with the scam, with each complaint assigned a single response category. The label *ignored* indicates that the victim recognizes the scam and avoided any engagement or interaction. *Interacted and victimized* denotes instances in which the victim communicated with the scammer and the interaction resulted in a loss occurrence, including financial loss, disclosure of sensitive information, or granting remote access to their device. *Interacted and not victimized* captures the cases in which the victim interacted with the scammer and no loss occurred or resulted. Complaints lacking sufficient information to determine the outcome are labeled *unknown*.

Phone Call Direction When a phone call communication is mentioned, this attribute captures whether the victim contacted the scammer (*inbound*) or the scammer contacted the victim (*outbound*) or cases of bi-directional call occurred (*inbound and outbound*). Complaints without telephone communication are labeled as not applicable.

Brand Targeted This attribute captures the legitimate company or organization impersonated during the scam. In contrast to the preceding categorical attributes, **Brand Targeted** is extracted as explicitly referenced in the complaint narrative such as Best Buy, Microsoft, Amazon and Norton. This supports analysis of brand impersonation trends and comparison with brands featured in the Facebook scam contributing groups.

4.2 Zero-Shot Structured Information Extraction

To automate the information extraction of the schema attributes, Gemma 3 (12 billion parameter) LLM was deployed using the Ollama framework, balancing computational constraints with advanced zero-shot reasoning. The decision to employ this specific language model builds directly upon recent literature [4], which has already proven its high performance in parsing complex text. To ensure deterministic outputs and mitigate model hallucination, generation temperature was strictly set to 0.0. The proposed framework was structured into two sequential stages.

Stage 1: Binary Inclusion-Exclusion Filter. Victim narratives contain both legitimate service complaints (such as warranty denials, shipping delays) and reports of fraudulent activity. An initial filtering process helped identify records associated with TSS-related activity. At this stage, the model was instructed to perform binary classification to distinguish authentic TSS, phishing or IT impersonation reports from non-TSS fraud and routine customer service issues. Only complaints that are classified as relevant scams were advanced to further stage.

Stage 2: Scam Attributes Extraction. For all complaints that passed the initial filter, the model was tasked with extracting scam-specific attributes into predefined JSON scheme developed for this study. In contrast to fine tuning a supervised model, we utilize a zero-shot approach relying exclusively on rigorous prompt engineering. Prompt instructions included explicit guardrails to prevent unsupported results, especially constraining the model from identifying the derived attributes unless such information was explicitly stated in the complaint.

4.3 Gold-Standard Annotation and Model Evaluation

We established a human-annotated benchmark to test the reliability and accuracy of the proposed information extraction pipeline before deploying it on the complete dataset. A random sample of 100 complaints was drawn from the larger corpus and manually annotated according to the predefined extraction schema. The resulting annotations serve as the ground-truth reference for evaluation. The structured JSON outputs generated by the locally hosted Gemma 3:12B model for this identical subset were subsequently evaluated systematically against the manual annotations. Standard performance metrics were calculated to evaluate the performance of both initial filtering and taxonomy extraction stage.

Validation of Stage 1: Binary Inclusion-Exclusion Filter. The binary inclusion-exclusion filter performed exceptionally well at distinguishing TSS scam related complaints from non-TSS complaints, achieving an overall accuracy of 94% and F1-score of 96.2%. The classifier correctly identified 75 of the 76 positive instances, achieving a recall of 98.7%, while incorrectly classifying only one positive case as negative. This filtration reflects a optimization strategy helping to retain all potential TSS scam reports for further processing rather than risking the exclusion of genuine fraud cases.

Validation of Stage 2: Scam Attributes Extraction. For complaints that advanced to the information extraction layer, performance of the pipeline was calculated across 4 principal dimensions:

- **Victim Response:** The proposed extraction framework demonstrated an overall accuracy of 86.6% with a Cohen's Kappa coefficient [5] of 0.79, explaining substantial agreement [6] between predicted and actual victim response class labels. The model missed only 1 of the 35 victimized complaints, showing a balanced accuracy 97.1% for the class *interacted and victimized*. The class *interacted and not victimized* was the most challenging class for the model to distinguish, as it mislabeled 6 out of 22 cases.
- **Lure:** Because lure information may include multiple communication channels within a single complaint, extraction performance was measured using token-level similarity measures. When compared with the ground truth labels, the model achieved a mean Jaccard similarity of 87% and a F1-score of 90%, elucidating the strong performance of pipeline in lure extraction.
- **Phone Call Direction:** The model achieved an overall accuracy of 86.8% with a Cohen's Kappa coefficient [5] of 0.81, indicating almost perfect agreement beyond chance [6]. Although the model performed accurately with not applicable and outbound calls, 6 among 25 inbound calls were marked as outbound calls. This pattern highlights a limitation of zero-shot prompting approach in resolving complex temporal relationships described in complaint narratives.
- **Brand Targeted:** The structured prompt zero-shot pipeline performed exceptionally well with the task of extracting targeted brand from the complaints. 61 out of 76 TSS scam-related complaints in the sample mentioned the brand that was targeted. The developed pipeline classified 58 out of 61 samples, yielding an overall accuracy of 95.1% and a Cohen's Kappa coefficient[5] of 0.944 indicating almost perfect agreement [6] between the model and benchmark labels.

These results confirm that the structured prompt zero-shot extraction framework was accurate enough to categorize and extract complex scam narratives and support quantitative analysis on the full corpus.

5 Results

Following the successful validation of the information extraction pipeline, we deployed the framework across the entire BBB dataset of 10,782 complaints. Of the entire dataset, 2,264 complaints were excluded on the grounds that they were classified as non-TSS reports. The remaining 8,518 confirmed TSS reports then advanced to the second stage, which extracted the four critical attributes: lure utilized, the victim's response, the direction of the phone call, and the brand targeted.

Of the 8,518 confirmed TSS reports, approximately 80% of them documented active engagement with the scammer. Of these interactions, 43.3% (3,696 reports) led to victimization (e.g., financial loss, data compromise, or unauthorized remote machine access), with more than half of the engagement resulting in victimization. Another 36.9% of victims interacted with the scammers but recognized the scam before any loss occurred. A minority of the reports (17%) indicated passive observations, where users ignored the initial scam attempt and submitted reports to share artifacts like phishing emails. The reported victimization rate shows that TSS is not

merely a nuisance, but a highly optimized social engineering vector that demands immediate structural intervention.

As noted earlier, a single report may contain multiple lure values. Quantitative analysis of the lure attribute revealed 25 distinct combinations derived from 6 categorical values. Among these, the five most prevalent combinations were: phone call alone (33.8%, 2,886 reports), email alone (19.8%, 1,690 reports), email combined with phone call (14.4%, 1,233 reports), pop-up notification combined with phone call (6.9%, 596 reports), and website combined with phone call (4.5%, 385 reports). When individual lure types were analyzed across all TSS reports, phone call emerged as the dominant communication channel in 68% of reports (5,796). Email was the second most common lure appearing in 38.2% of reports (3,255), while website and pop-up notification lures each accounted for roughly 10%.

Analysis of phone call direction yielded notable findings. Across all reports, 32% did not include phone call as a lure element. The remaining reports were quite uneven, with 44% of reports with outbound calls (where scammers directly call victims) and 17.3% with inbound calls (where victims call scammers' provided numbers). Bidirectional calls where both parties initiated calls during the scam life-cycle accounted for 6% of all reports. This distribution indicates the importance of phone as a communication channel in TSS.

5.1 Brand Concentration across Ecosystems

Table 1: Brand Mentions Across Datasets and Associated Victimization Rate

Brand	Number of Mentions Facebook	BBB	Victimized (%)
Cash App	6,672	84	94.05
Facebook	4,334	108	68.52
Comcast	7,442	81	59.26
Microsoft	25,946	1,101	51.68
Amazon	29,729	312	51.60
Apple	13,441	474	45.36
McAfee	6,030	356	35.39
Paypal	31,550	198	33.33
Best Buy	4,894	1,117	28.74
Norton	18,335	944	26.91
Average	14,837	477	49.48

5.1.1 ***Analysis of Targeted Brands in Informal Marketplace Postings.*** Our analysis pipeline builds upon the foundational taxonomy established by Cherupalli et al. [4], which maps informal marketplace posts to specific service/product vectors (provided in Appendix B) and classified their operational intent (supply side advertisements versus demand side requests). Three of the taxonomy categories – *Money Laundering Services, Scammer Warnings* and *Other* – were excluded from the analysis to avoid potential biases. For example, posts from *Money Laundering Services* frequently referenced gift card brands that are not associated with the impersonation attacks and therefore had the potential to bias the frequency analysis for targeted brands. The exclusion narrowed the total Facebook group posts from 381,843 posts to 232,979 posts.

After excluding the categories described above, we applied a keyword-matching technique to all Facebook group posts by cross-referencing them against the curated list of brand compiled from consumer protection advisories and fraud reports published by the FTC. Each post was examined for mentions of these targeted brands, allowing us to identify the frequency with which specific brands were referenced within the informal marketplace. Lexical analysis revealed that 148,373 (63.7%) of these posts referenced the mentioned brands. It should be noted that a single posting, whether an advertisement or request, can contain several brands. Within these posts, PayPal (31,550 posts), Amazon (29,729 posts) and Microsoft (25,964 posts) were most frequently mentioned, as shown in Table 1. Other technology-related brands such as Norton, Apple, Google, McAfee, Comcast, and Best Buy were also found in the top ten targeted brands. Furthermore, every brand in the top ten receives no fewer than 4,000 mentions.

5.1.2 ***Analysis of Targeted Brands in Consumer Reporting Platform.*** The scale of brand impersonation exposed by our information extraction pipeline highlights a highly asymmetric threat landscape. Across 8,518 public reports from the BBB Scam Tracker, we identified 1,583 unique brands impersonated across 7,407 reports, while 1,111 reports contained no brand reference. Analysis of the extracted brands reveals a highly skewed distribution of target selection, with a small number of brands appearing frequently and a pronounced long tail of infrequently targeted brands. Specifically, 1,325 brands were each impersonated only once across all victim reports. The three most frequently targeted brands, as reported by victims, were Best Buy (1,117 reports; 15.0%), Microsoft (1,101 reports; 14.8%), and Norton (944 reports; 12.7%) as shown in table 1. Apple, McAfee, Amazon, PayPal, Facebook, Cash App, and Comcast rounded out the remainder of the top ten most targeted brands. This finding highlights the dominance of technology-related brands across the top ten targeted brands mentioned by victims. By contrast, the long list of brands appearing only once likely reflects scattered, one-off impersonation attempts rather than a systematic threat.

5.1.3 ***Temporal Patterns of Top 10 Brands across Ecosystems.*** To better understand the relationship between supply-side threat actor operational focus and real-world victim impact, we analyzed and compared the temporal rank-order of the top 10 targeted brands across informal marketplaces and public scam-reporting platform between 2015 and 2024 as shown in Figure 1. During our analysis, we observed that legacy brands exhibited high relation across both datasets. Microsoft and Norton consistently ranked within the top five in both scam contributing groups and victim complaints. This suggests a persistent, high-volume nature of classic tech support and fake antivirus renewal scams, where the high supply of attacker infrastructure directly translates to sustained, measurable consumer victimization.

The data also exposes a profound operational dichotomy. Brands such as PayPal and Amazon dominate the informal marketplaces ranking, with both consistently appearing in between top four. Conversely, these brands barely sit in top three in BBB scam reports. Additionally, the COVID-19 pandemic (2020–2023) is associated

Table 2: Mapping Supply-Side Scam Infrastructure to Real-World Victim Encounters Across Leading Brands.

	Attack Delivery Vectors (%)											
	Email		*Website*		*Pop-up*		*SMS*		*Inbound*		*Outbound*	
Brand	**FB**	**BBB**	**FB**	**BBB**	**FB**	**BBB**	**FB**	**BBB**	**FB**	**BBB**	**FB**	**BBB**
Cash App	32.48	5.95	0.70	30.95	65.50	1.19	32.48	8.33	89.89	39.29	49.77	59.52
Facebook	44.03	5.56	2.86	26.85	47.14	2.78	44.03	2.78	67.90	42.59	47.00	50.00
Comcast	94.66	11.11	0.08	19.75	5.22	0.00	94.66	6.17	28.80	17.28	98.71	75.31
Microsoft	33.24	12.81	0.03	9.63	66.51	37.78	33.24	2.09	88.28	30.43	40.17	65.85
Amazon	56.03	22.12	0.33	18.27	43.28	5.45	56.03	8.65	72.74	22.12	67.11	65.06
Apple	48.27	9.49	0.49	17.72	48.55	10.34	48.27	4.01	73.07	21.94	58.60	73.00
McAfee	75.56	74.72	1.43	7.58	22.39	2.53	75.56	1.97	67.05	26.69	81.94	32.30
Paypal	65.21	81.31	0.64	5.56	31.45	0.51	65.21	3.54	68.24	23.23	72.80	31.31
Best Buy	76.84	71.80	1.08	3.94	21.32	1.16	76.84	4.30	76.03	21.67	85.96	30.17
Norton	69.63	81.04	0.69	3.28	29.45	0.85	69.63	3.28	75.36	28.07	74.49	30.61
Average	59.59	37.59	0.83	14.35	38.08	6.26	59.59	4.51	70.74	27.33	67.66	51.31
Correlation (*r*)	0.52		0.22		0.59		-0.05		0.51		-0.28	

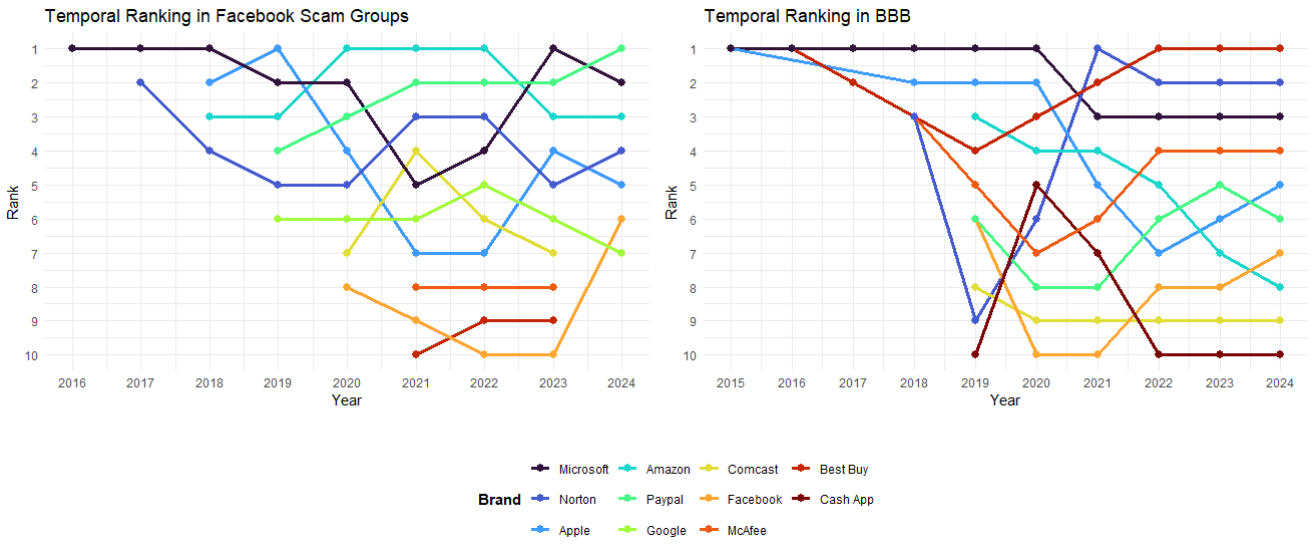

Figure 1: Temporal Rank-Order Of The Top 10 Targeted Brands Across Informal Marketplaces and BBB Reports

with a sharp rise in the prominence of e-commerce brands, which peaked in both datasets between 2020 and 2022, with Amazon leading in informal marketplaces and Best Buy in BBB complaints. This indicates a period of attackers focus and victim impact during the global surge in online shopping. This patterns also reveals a temporal transition from targeting the technical brands to targeting the e-commerce brands.

5.2 Mapping Supply-Side Threat Vectors to Consumer Victimization

To examine the extent to which supply-side market activity is reflected in real-world victimization, we performed a Pearson correlation analysis comparing the proportional distribution of attack vectors offered within informal marketplaces with those reported in victim complaints for the 10 most frequently impersonated brands as shown in Table 2.

To facilitate this comparison, informal marketplace products or services were mapped to their corresponding attack delivery vectors.

- **Email:** Derived by combining *Blasting Campaign Services* and *Victim Data Sales*.
- **SMS:** Derived by combining *Blasting Campaign Services* and *Victim Data Sales*.
- **Website:** Mapped exclusively to *Web Development Services*.
- **Pop-up:** Constructed by aggregating *PPC/Popups Calls* and *Toll-Free Number Providers*.
- **Inbound Phone Call:** Encompassed a broader aggregation of *PPC/Popups Calls*, *Blasting Campaign Services*, *Web Development Services* and *Toll-Free Number Providers*.
- **Outbound Phone Call:** Derived from the combination of *Victim Data Sales*, *Blasting Campaign Services* and *Toll-Free Number Providers*.

The correlation analysis revealed a vector-dependent relationship between attacker marketplace and victim reporting patterns. The strongest positive association observed was *Pop-up* vector (r=0.59), indicating a moderately strong, positive trend with the limited sample size (n=10). This finding suggests that the procurement of *PPC/Popups Calls* and *Toll-Free Number Providers* within the informal marketplace are generally reflected in real world pop-up scam reports, particularly for brands such as Microsoft, which accounted for 37.78% of victim reports and 66.51% of supply side postings.

The *Email* and *Inbound* vectors also demonstrated moderately strong positive correlations, although these relations were not statistically significant. Legacy security and retail brands, including Norton, McAfee and Best Buy, showed consistently high levels of email-based attack vectors across both ecosystems. This is consistent with the prominent role of *Blasting Campaign Services* in the informal marketplace, which is utilized for large-scale distribution of email and SMS-based lures [4].

In contrast, *SMS* and *Outbound* vectors displayed considerably weaker alignment across the two ecosystems, with both correlation factors exhibiting weak negative correlations. This divergence suggests that while these attack vectors may rely on similar underlying infrastructures (e.g., *Blasting Campaign Services* and *Victim Data Sales*), the prevalence of these resources in the marketplace does not necessarily correspond to their representation in victim complaints. This discrepancy is particularly apparent for SMS-based scams.

Across the BBB dataset, email-based lures accounted for 37.59% while SMS-based lures represented only 4.51% of complaints. The low volume of *SMS* reports may reflect both consumer reporting practices and the scope of the dataset. Since this study focused specifically on TSS and excluded broader BBB phishing reports, victims of SMS-based attacks may have reported these incidences under other scam categories. Overall, these findings suggest that supply-side marketplace activity provides a moderate indication of the prevalence of inbound and *Email*-based TSS operations, while offering substantially weaker predictive value for *SMS* and *Outbound* attack vectors.

Table 3: Logistic Regression (Odds Ratios) Predicting Victimization

	Dependent variable:	
	was_victimized	
	Model 1	**Model 2**
Brand Not in Top 10		1.461***
		(0.0.116)
Amazon		1.568***
		(0.221)
Apple		0.982***
		(0.119)
Best Buy		1.091***
		(0.115)
Cash App		20.825**
		(9.838)
Comcast		2.126***
		(0.528)
Facebook		2.690***
		(0.625)
McAfee		1.360***
		(0.193)
Norton		0.997***
		(0.111)
Paypal		1.411***
		(0.253)
has_email	0.495***	0.550***
	(0.028)	(0.035)
has_popup	3.057***	3.506***
	(0.271)	(0.330)
has_website	3.592***	3.459***
	(0.301)	(0.293)
has_sms	0.727***	0.697***
	(0.081)	(0.079)
has_inbound	2.444***	2.468***
	(0.147)	(0.151)
has_outbound	1.710***	1.734***
	(0.094)	(0.097)
Constant	0.481***	0.352***
	(0.027)	(0.031)
McFadden R-squared	0.12	0.13
AIC	10306.03	10187.32
BIC	10355.38	10307.17
Observations	8,518	8,518

Note: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

5.3 Regression Analysis of Victimization Rate Across Brands and Lure Vectors

To determine how attacker tactics such as impersonated brand and lure affect the likelihood of victimization, we require a statistical model capable of predicting binary outcomes. The dependent variable, *was_victimized*, is a binary indicator of whether a consumer reported falling victim to a scam (1) or simply reported the encounter (0). Because standard linear regression can produce invalid probabilities outside the [0,1] range when applied to binary data, it is mathematically unsuitable for this dataset. Therefore, we utilize logistic regression, which satisfies the conditions for our data by mapping the predictor variables to a valid probability curve.

To formally define our approach, we utilized the standard logit function, where p is the probability of victimization. As a single report can contain multiple lure tactics simultaneously, we engineered a set of binary indicator variables for each observed lure feature. Each variable is set to 1 if the specific lure type is present in the report, and 0 otherwise. Model 1 (n=8,518) serves as our

baseline, evaluating the isolated impact of different lure methods across the entire dataset:

$$\ln\left(\frac{p}{1-p}\right) = \beta_0 + \beta_1 X_{\text{email}} + \beta_2 X_{\text{popup}} + \beta_3 X_{\text{web}} + \beta_4 X_{\text{sms}} + \beta_5 X_{\text{inbound}} + \beta_6 X_{\text{outbound}} \tag{1}$$

Model 2 (n=8,518) introduces the targeted brands into the regression to measure brand-specific effects alongside the lure methods. Rather than restricting the sample to a subset of brands, this model retains all observations by grouping the brand data. The top 10 most frequently targeted brands are preserved as distinct levels, whereas all other brands are collapsed into a single "Brand Not in Top 10":

$$\begin{aligned}\ln\left(\frac{p}{1-p}\right) =& \beta_0 + \beta_1 X_{\text{email}} + \beta_2 X_{\text{popup}} + \beta_3 X_{\text{web}} \\ &+ \beta_4 X_{\text{sms}} + \beta_5 X_{\text{inbound}} + \beta_6 X_{\text{outbound}} + \sum \beta_{\text{brand}} X_{\text{brand}}\end{aligned} \tag{2}$$

Table 3 presents the results of these regression models. For the ease of interpretation, coefficients for both models are reported as Odds Ratios (OR = e^{odds_ratio}) rather than raw log-odds. An odds ratio greater than 1.0 indicates an increased likelihood of victimization, while a value below 1.0 indicates a decreased likelihood relative to the baseline category. In model 2, Microsoft serves as the baseline reference category for the brand variable. In other words, all brand specific odds ratio represent the likelihood of victimization as compared to a Microsoft impersonated scam.

The efficacy of delivery vectors remains remarkably consistent across both models, revealing a distinct division between lure methods. Scams that force immediate user interaction or rely on consumer-initiated contact are the strongest predictors of victimization. Websites (Model 1 OR = 3.59; Model 2 OR = 3.45; both $p < .01$) pose the greatest risk, more than tripling the odds of victimization. The *Pop-up* attack vector similarly exhibited approximately three times greater odds of victimization (Model 1 OR = 3.06; Model 2 OR = 3.51; both $p < .01$). *Inbound* phone calls (OR = 2.44 and 2.47) and *Outbound* phone calls (OR = 1.71 and 1.73) are likewise associated with significantly elevated odds of victimization. Both *Email* (Model 1 OR = 0.50; Model 2 OR = 0.55; both $p < .01$) and *SMS* (Model 1 OR = 0.73; Model 2 OR = 0.69; both $p < .01$) exhibit odds ratios below 1.0. This indicates that while these vectors are heavily utilized for initial lure distribution, users are comparatively well-equipped to identify and report them without suffering financial or data loss.

Model 2 incorporates both impersonated brand and lure methods. The impersonated brand acts as a strong modifier of TSS victimization rate. Aside from Apple and Norton, which decrease the odds of victimization, all other brands increased the likelihood of victimization compared to Microsoft. Scams that targeted Amazon, McAfee, and PayPal had a slightly higher likelihood of victimization rate. In contrast, scams impersonating Cash App are associated with substantially elevated odds of victimization (OR = 20.82, $p < .05$), indicating that respondents targeted by Cash App impersonation are nearly twenty times more likely to fall for the scam compared to a Microsoft impersonation. Consumers encountering Facebook (OR = 2.69, $p < 0.01$) and Comcast (OR = 2.12, $p < 0.01$) impersonations are also roughly twice as likely to fall victim relative to the Microsoft baseline. Interestingly, Model 2 demonstrates that the odds of victimization were higher with the brands that are not in top 10 (OR = 1.46, $p < 0.01$), surpassing six out of top ten brands.

Overall, compared to Model 1, Model 2 emerges as the most robust specification, as the inclusion of brand substantially improves predictive performance over the baseline. In the context of predicting unpredictable outcomes like scam victimization, a McFadden R-squared of 0.13 denotes a highly credible and robust model fit. The statistical strength of Model 2 is best evidenced by its AIC score of 10187.32, a decrease of nearly 119 points from the base model, providing decisive evidence that brand inclusion is essential for accurate predictions.

6 Discussion

This research advances the methodological scope of current TSS literature by correlating informal supply-side market indicators with real-world victimization data aggregated via the BBB Scam Tracker. The integration of these disparate datasets yields a multidimensional threat model that captures operational structures, patterns of brand exploitation, lure-specific victimization risks, and the extent to which underground marketplace activities corresponds with victim-reported outcomes.

To ensure the empirical validity of downstream analyses, we rigorously evaluated the zero-shot LLM extraction pipeline's capacity to map unstructured qualitative text to a deterministic schema. The two-stage Gemma 3 pipeline performed strongly overall, particularly in Stage 1, with sensitivity of 98.7%. Stage 2 successfully extracted victim responses, lure methods and brands targeted but struggled slightly with identifying the direction of phone calls, often guessing the scammer called first even when the victim initiated contact. This suggests that directional reasoning is harder than categorical labeling for zero-shot LLMs.

Our victim response analysis shows that victims rarely ignored during initial interaction with scammers. Nearly 80% engaged with the scammers, and 43.3% suffered confirmed harm. This high victimization rate proves that once communication begins, the scammer is highly likely to succeed. Therefore, prevention efforts focused solely on public awareness of scam existence may be insufficient. Interventions that specifically interrupt the moment of initial contact, before sustained interaction begins, may offer a more effective point of leverage.

Across both the lure and phone call direction analysis, phone call communication emerged as the connective thread running through TSS operations. Phone calls appeared in 68% of all reports, either alone or paired with email, pop-ups, or websites, and the most common lure combinations consistently involved a phone call component. Interestingly, scammers initiated these calls (outbound) two-and-a-half times more often than victims did (inbound), proving that TSS is a highly aggressive, attacker-driven operation. This finding is confirmed by the logistic regression results: inbound calls, outbound calls, pop-ups, and websites were all associated with significantly elevated odds of victimization, while email and SMS were associated with reduced odds. These results align with evidence from the informal market place, where *Blasting Campaign Services* and *Victim Data Sale* categories showed similar prominence, highlighting the continued importance of direct engagement and victim targeting mechanisms in scam operations. Live conversation

simply gives scammers the real-time leverage needed to manipulate victims.

The brand-level analysis revealed that scammers do not impersonate brands at random. Out of 1,583 identified brands from the BBB Scam Tracker, the top 10 brands accounted for 64.4% of all targeted complaints, with Best Buy, Microsoft, and Norton dominating the landscape. Attackers clearly rely on the established trust of household names to bypass victim skepticism. Additionally, our regression analysis shows that the victimization can be better predicted with the inclusion of impersonated brands.

The Pearson correlation analysis comparing Facebook informal marketplace groups posting to BBB-reported attack vectors provides a novel bridge between the supply-side and real-world consequences. The strongest and most theoretically coherent relationship emerged for *Pop-up* (r = 0.58), suggesting that when underground actors advertise *PPC/Popups Calls* and *Toll-Free Number Services* [4], this activity does tend to materialize as real-world pop-up-based scams, a pattern clearly visible for Microsoft, which dominated both datasets. Moderate correlations were also observed for email and inbound vectors. Thus, supply-side measurement functions as a high-fidelity predictive indicator for specific attack vectors (pop-ups, email) but lacks the generalization required for comprehensive threat forecasting.

Our findings demonstrate TSS as a highly centralized ecosystem in which a small number of well-resourced, phone-centric operations impersonating a handful of trusted brands account for the overwhelming majority of consumer harm. While the correlation between informal market provisioning and formal incident reporting is vector-dependent, the statistically significant alignment within pop-up and email channels confirms that targeted monitoring of underground infrastructure can operationalize early-warning capabilities for specific classes of consumer fraud.

7 Limitations

Several limitations merit consideration. First, the scope of this study is confined to data collected from Facebook groups. Although similar scam-supporting communities are known to operate across multiple online platforms, the size, diversity, and sustained activity of the Facebook ecosystem provide a strong empirical foundation for investigating the characteristics of the informal marketplace.

Second, the use of a single LLM, Gemma 3, for automated information extraction. While evaluation against a human-annotated dataset demonstrated strong overall performance, the model's extraction accuracy may vary across different narrative styles present in the complete corpus. Specifically, the model demonstrated a measurable bias toward predicting outbound call flows in instances where semantic context indicated victim-initiated contact, highlighting the persistent limitations of zero-shot LLMs in resolving temporal and directional dependencies.

Finally,the BBB Scam Tracker captures only voluntary submissions, which naturally leads to under-reporting. Many targets ignore fraudulent overtures. If they do respond and are victimized, they may lack the awareness or incentive to report to third-party databases. Consequently, the victim narrative dataset reflects only self-reported complaints and may under-represent certain scam types (e.g., isolating smishing from TSS taxonomies) that victims categorize differently.

8 Future Work

This study can be further extended in several directions. First, investigating how supply-and-demand dynamics in informal market places evolve over time will improve on current knowledge on scam tactics and strategies, which may in turn affect the dynamics of the marketplace. Second, the dataset can be expanded by incorporating data from other platforms acting as informal market places and incorporating additional complaint sources, enabling a sophisticated understanding of cross-ecosystem corelation. Finally, future research should focus on developing empirical disruption strategies for TSS ecosystem on social media platforms.

9 Conclusion

This research advances the state of TSS literature by correlating informal supply-side market provisioning with real-world victim complaints. To achieve a scalable quantitative analysis, we engineered and validated a two-stage zero-shot LLM extraction pipeline, effectively translating 8,518 unstructured BBB complaints into a deterministic schema encompassing lure vectors, victim interaction states, phone call directionality, and brand impersonation. This methodological leap allowed us to move beyond simply detecting scams to mapping their entire operational life cycle.

The empirical results characterize TSS as a highly aggressive, phone-centric operation. Telephonic infrastructure mediates the vast majority of reports, with outbound adversarial calls demonstrating a dominant prevalence over inbound victim contact. The data exhibits a critical conversion threshold: once interaction is initiated, the probability of victimization escalates significantly, dictating that mitigation frameworks must interrupt the attack lifecycle prior to the establishment of synchronous communication. Brand targeting is highly concentrated around a small set of trusted, widely recognized names such as Best Buy, Microsoft, and Norton. However, our multivariate analysis demonstrates that evaluating the impersonated brand in tandem with the specific lure technique predicts victimization significantly better than a baseline model restricted to lure characteristics alone.

Finally, this research provides a blueprint for comprehensive disruption of tech support scam ecosystems. Regression analyses demonstrate that websites (OR=3.45) and pop-ups (OR=3.51) pose the highest odds of victimization, and our comparative findings confirm these vectors can be forecasted via supply-side dat. Consequently, advertising networks and search engines are uniquely positioned to ingest this actionable intelligence and initiate takedowns where appropriate. By proactively rejecting malvertising campaigns that impersonate high-risk entities such as Microsoft and Norton, these platforms can sever primary attack vectors.

Furthermore, social media networks hosting these informal syndicates can directly operationalize the zero-shot LLM extraction pipeline developed in this study. Integrating this pipeline into existing Trust and Safety architectures will enable platforms to autonomously detect, classify, and terminate the accounts of upstream infrastructure brokers, thereby dismantling the fraud supply chain at its source.

References

[1] Bhupendra Acharya, Dario Lazzaro, Efrén López-Morales, Adam Oest, Muhammad Saad, Antonio Emanuele Cinà, Lea Schönherr, and Thorsten Holz. 2024. The imitation game: exploring brand impersonation attacks on social media platforms. In *Proceedings of the 33rd USENIX Conference on Security Symposium* (Philadelphia, PA, USA) *(SEC '24)*. USENIX Association, USA, Article 248, 18 pages.

[2] Better Business Bureau. [n. d.]. BBB Scam Tracker. https://www.bbb.org/scamtracker.

[3] Raghavendra Cherupalli, Yi Ting Chua, Tyler Moore, and Gary Warner. 2026. Analyzing Supply and Demand Signals for Cybercrime Services in Technical-Support Scams. In *8th Workshop on Attackers and Cyber-Crime Operations (WACCO 2026)*. To appear.

[4] Raghavendra Cherupalli, Hawken Grubbs, Yi Ting Chua, Weiping Pei, Tyler Moore, and Gary Warner. 2025. Contextual Classification of Cybercriminal Posts Using Large Language Models: A Comprehensive Study on Tech Support Scam Marketplaces. In *2025 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 1–11.

[5] Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20, 1 (1960), 37–46.

[6] Juliet M Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology* 13, 1 (1990), 3–21.

[7] Federal Bureau of Investigation. 2026. *2025 Internet Crime Report*. Annual Report. Internet Crime Complaint Center (IC3). https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf

[8] Federal Trade Commission. 2024. *New FTC Data Shed Light on Companies Most Frequently Impersonated by Scammers*. https://www.ftc.gov/news-events/news/press-releases/2024/05/new-ftc-data-shed-light-companies-most-frequently-impersonated-scammers

[9] Federal Trade Commission. 2024. Got a Barcode from Your Utility Company to Make a Payment? That's a Scam. https://consumer.ftc.gov/consumer-alerts/2024/07/got-barcode-your-utility-company-make-payment-thats-scam. Consumer Alert, Accessed: 2026-07-10.

[10] Federal Trade Commission. 2024. Impersonation Scams: They're Not What They Used To Be. https://www.ftc.gov/news-events/data-visualizations/data-spotlight/2024/04/impersonation-scams-not-what-they-used-be. Data Spotlight, Accessed: 2026-07-10.

[11] Federal Trade Commission. 2025. Are You Really Out of Cloud Storage? Or Is It a Message Scam? https://consumer.ftc.gov/consumer-alerts/2025/07/are-you-really-out-cloud-storage-or-message-scam. Consumer Alert, Accessed: 2026-07-10.

[12] Federal Trade Commission. 2025. How To Spot, Avoid, and Report Tech Support Scams. https://consumer.ftc.gov/articles/how-spot-avoid-and-report-tech-support-scams

[13] Federal Trade Commission. 2025. Seemingly Urgent Security Messages Could Lead to Tech Support Scams. https://consumer.ftc.gov/consumer-alerts/2025/04/seemingly-urgent-security-messages-could-lead-tech-support-scams. Consumer Alert, Accessed: 2026-07-10.

[14] Federal Trade Commission. 2025. Think That Text Message Is from USPS? It Could Be a Scam. https://consumer.ftc.gov/consumer-alerts/2025/04/think-text-message-usps-it-could-be-scam. Consumer Alert, Accessed: 2026-07-10.

[15] David Harley, Martijn Grooten, Steven Burn, Craig Johnston, et al. 2012. My PC has 32,539 errors: How telephone support scams really work. *Virus Bulletin* (2012).

[16] Hoxhunt. 2026. *Phishing Trends Report 2026*. Technical Report. Hoxhunt. https://hoxhunt.com/guide/phishing-trends-report

[17] Zakaria Kaddari, Ikram El Hachmi, Jamal Berrich, Rim Amrani, and Toumi Bouchentouf. 2024. Evaluating large language models for medical information extraction: A comparative study of zero-shot and schema-based methods. *Applied Computer Science* 20, 4 (2024).

[18] Jonathan Larson, Bryan Tower, Duane Hadfield, Darren Edge, and Christopher White. 2018. Using web-scale graph analytics to counter technical support scams. In *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 3968–3971.

[19] Jienan Liu, Pooja Pun, Phani Vadrevu, and Roberto Perdisci. 2023. Understanding, measuring, and detecting modern technical support scams. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. IEEE, 18–38.

[20] Microsoft CEE Multi-Country News Center. 2024. Microsoft Digital Defense Report: 600 Million Cyberattacks per Day Around the Globe. https://news.microsoft.com/en-cee/2024/11/29/microsoft-digital-defense-report-600-million-cyberattacks-per-day-around-the-globe/.

[21] Najmeh Miramirkhani, Oleksii Starov, and Nick Nikiforakis. 2017. Dial One for Scam: A Large-Scale Analysis of Technical Support Scams. In *Proceedings 2017 Network and Distributed System Security Symposium*. Internet Society.

[22] Tyler Moore and Richard Clayton. 2015. Which malware lures work best? Measurements from a large instant messaging worm. In *2015 APWG Symposium on Electronic Crime Research (eCrime)*. 110–. doi:10.1109/ECRIME.2015.7120801

[23] Sampsa Rauti and Ville Leppänen. 2017. "You have a potential hacker's infection": A study on technical support scams. In *2017 IEEE International Conference on Computer and Information Technology (CIT)*. IEEE, 197–203.

[24] Dawna Roberts. 2025. Amazon Scam: How It Works and How to Avoid It. https://moonlock.com/amazon-scam

[25] Bharat Srinivasan, Athanasios Kountouras, Najmeh Miramirkhani, Monjur Alam, Nick Nikiforakis, Manos Antonakakis, and Mustaque Ahamad. 2017. By hook or by crook: Exposing the diverse abuse tactics of technical support scammers. *arXiv preprint arXiv:1709.08331* (2017).

[26] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786* (2025).

[27] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295* (2024).

[28] Ian Wood, Michal Kepkowski, Leron Zinatullin, Travis Darnley, and Mohamed Ali Kaafar. 2023. An analysis of scam baiting calls: Identifying and extracting scam stages and scripts. *arXiv preprint arXiv:2307.01965* (2023).

[29] Yongqin Xian, Bernt Schiele, and Zeynep Akata. 2017. Zero-shot learning-the good, the bad and the ugly. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4582–4591.

[30] Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. 2024. Large language models for generative information extraction: A survey. *Frontiers of Computer Science* 18, 6 (2024), 186357.

A List of brands and organizations extracted from news articles, and federal organizations reports

Airbnb, Amazon, Apple, Bank of America, Best Buy, Binance, CVS, Comcast, Facebook, FedEx, Google, Microsoft, Netflix, Norton, PayPal, Publishers Clearing House, Toyota, UPS, USPS, Utility Company, Walgreens, Walmart, and Wells Fargo.

B Descriptions of Product/Service Categories provided by Cherupalli et al. [4]

Table 4: Descriptions of Product/Service Categories

Category	Definition
Money-Laundering Services	Services that cash out or transfer stolen payments
Fake/Illicit Document Services	Impersonating the documents of legitimate organizations, as well as fake IDs.
Victim Data Sales	Data of victim contact or personal details.
Blasting Campaign Services	Services that transmit mass email or SMS messages to prospective victims using the technique called email/SMS blasting, by which inbound calls are generated.
PPC/Popup Calls	Services that assist scammers in running malware advertising campaigns, such as PPC and pop-ups to generate inbound calls.
Job Offerings	Hiring workers that take phone calls at call centers.
Criminal IT Infrastructure Operations	Scripts and tools required to operate criminal IT infrastructure, excluding toll free number providers, website development services and remote access services.
Toll-Free Number Providers	Toll-free number (TFN) numbers used to contact victim.
Remote Access Services	Remote access software used to connect to victim devices.
Web Development Services	Website services, including hosting and promotion.