Q3 2019

DataVisor Fraud Index Report

DECONSTRUCTING CONTENT ABUSE PATTERNS IN USER-GENERATED CONTENT

DATAVISOR

Table of Contents

Foreword

> *If users can no longer trust the content they engage with on a particular platform, they will eventually cease to use the platform at all, and when customer churn increases, investors worry, advertisers depart, and businesses struggle.*

User-generated content (UGC) has come to play an increasingly important role in our digital economy. Its presence serves many purposes across industries, platforms, and use cases. From validating the relevance of a business through social proofing to humanizing a brand and establishing authenticity, incorporation of high-quality UGC can often make the difference between success and failure for an online business. With increasing pressure to streamline operational overhead while simultaneously growing profits, many organizations have increasingly come to rely on UGC for significant percentages of their entire content output. In some cases, UGC is baked right into the business model—social media platforms and review sites being two notable examples.

If the incorporation of high-quality user-generated content can add significant value for a brand, the opposite is regrettably also true. The proliferation of fake, abusive, fraudulent, deceptive, and toxic user-generated content can severely damage a brand. If users can no longer trust the content they engage with on a particular platform, they will eventually cease to use the platform at all, and when customer churn increases, investors worry, advertisers depart, and businesses struggle.

According to numbers released earlier this year by Stackla: "79 percent of people say user-generated content highly impacts their purchasing decisions," and "90 percent of consumers say authenticity is important when deciding which brands they like and support." Additionally, the report notes that "consumers are 2.4x more likely to say user-generated content (UGC) is authentic compared to brand-created content." These trends have been ongoing for some time. Research by the Pew Research Center from 2016, for example, indicates that "82% of U.S. adults say they at least sometimes read online customer ratings or reviews before purchasing items for the first time."

Given the importance of user-generated reviews to online shoppers, the levels of fake content are alarming. As but one example, widely circulated numbers recently provided by Fakespot suggest that more than 60% of electronics reviews on Amazon are fake.

Online content abuse is, of course, nothing new. In the 1990s, spam on instant messaging platforms was an ongoing concern. In the 2000s, spam content spread across a wider array of devices and platforms, infiltrating review sites, polluting mobile messaging services, and more. Throughout our current decade, the problem has become a bot-powered epidemic, with virtually no platform safe from toxic content.

New technologies have dramatically exacerbated the problem of content abuse. Bots make it possible for content abuse to take place at massive scale. The increasing proliferation of UGC-dependent sites—with all their easily accessible points-of-entry—have afforded fraudsters ample opportunity to flood the web with fake and abusive content. Moreover, fraudsters themselves have become far more sophisticated—at obfuscating their activities, and developing bespoke attack types that combine existing techniques like phishing and call center fraud with new approaches like formjacking.

The deadly combination of size, scope, and sophistication makes modern content abuse a formidable problem, and with so many organizations relying on user-generated content to grow their businesses, the world is in dire need of effective solutions to face down the challenge. Legacy solutions that depend only on supervised machine learning and rules-based approaches cannot keep pace with the agility and speed of modern fraud. Reactive strategies that do little more than offer after-the-fact damage control cannot act in time to prevent businesses from suffering under the weight of toxic content. Only with proactive, AI and unsupervised-machine learning approaches—informed by superior fraud domain expertise, and supported with vast amounts of relevant global intelligence—can we hope to protect businesses from content abuse that damages brands, drives away users, and results in significant financial loss.

- Ting-Fang Yen,
Director of Research, DataVisor

THE DATAVISOR GLOBAL INTELLIGENCE NETWORK

The DataVisor Global Intelligence Network (GIN) leverages deep learning technologies to provide real-time, comprehensive digital intelligence based on a vast set of data signals that include IP addresses, geographic locations, email domains, mobile device types, operating systems, browser agents, phone prefixes, and more. All told, the GIN aggregates anonymized signals across a global client database of more than four billion users. By analyzing the connections between these data points in context—not just in isolation—DataVisor provides fine-grained signals and reputation scores that can be consumed directly in detection, or used to enhance rules engines and machine learning (ML) solutions.

To produce this report, we processed and analyzed the following for the period of: April-June, 2019:

- **80 billion events**
- **758 million users**
- **368 million IP addresses**
- **4.69 million /24 IP subnets**
- **1.05 million email domains**
- **4.95 million user-agent strings**
- **229K device types**
- **458K phone number prefixes**

SECTION 1

Why Content Abuse Is An Ongoing Problem

> *The advancing democratization of online access brings with it unique challenges, in that there are now vastly more "entry points" that enable fraudsters to introduce malicious user-generated content into online ecosystems.*

Content abuse continues to be a major concern across the fraud landscape, and several factors are combining to intensify the problem. For one thing, modern fraudsters have an ever-broadening palette to choose from when it comes to attack types and techniques.

Spam, scams, phishing, promo abuse, ticketing fraud, fake reviews—these all require content, and so we see significant increases in the proliferation of toxic content accordingly.

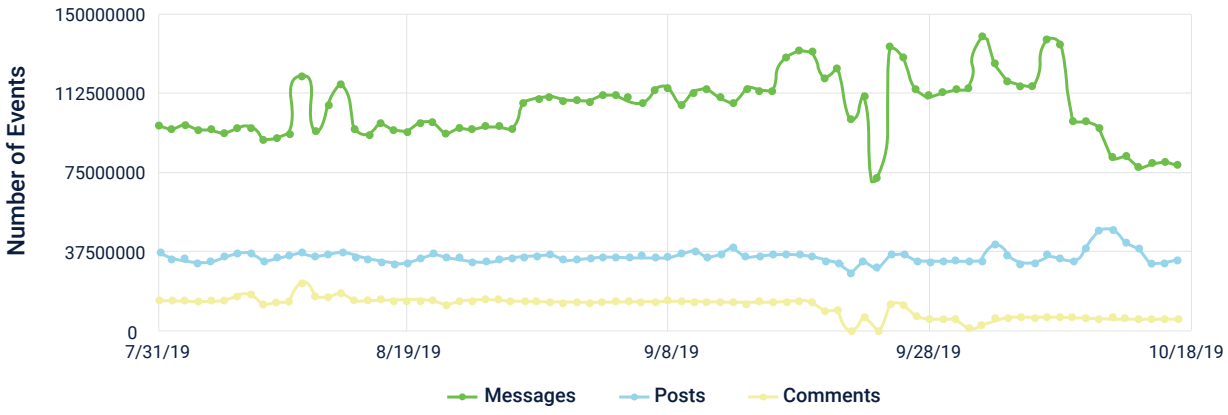

Figure 1: More than 140 million unique pieces of user-generated content are observed on average each day in the DataVisor Global Intelligence Network.

Additionally, the advancing democratization of online access—a positive in so many ways—also brings with it unique challenges, in that there are now vastly more "entry points" that enable fraudsters to introduce malicious user-generated content into online ecosystems. This content is almost uniformly volatile—easy to create, and inexpensive to manipulate. This makes it exceedingly difficult for fraud solutions to keep pace, and it is especially challenging to consistently and effectively cover the full scope of potential attack points. For example, fraud defenses may be monitoring the body copy of forum posts, but failing to track post titles; or, there may be a solution in place to monitor messaging content between users, but it may not be simultaneously reviewing the nicknames users are able to give themselves in their profiles.

Finally, modern fraudsters are innovating in new ways that are complicating prevention efforts. While "new" attacks types and techniques do continue to emerge, it is increasingly common for malicious actors to creatively assemble different combinations of existing attack types to advance their illicit agendas. By "mixing and matching" attack components available in the fraud-as-a-service underground economy, fraudsters can focus on their end goal, as opposed to getting bogged down in process. This gives them the leverage to launch highly specialized attacks targeted towards certain victims or platforms, and to have very specific goals and objectives for their campaigns.

As online platforms have continued to adopt ML-based content abuse solutions, fraudsters have learned how to manipulate their malicious content in ways that enable them to still reach real humans, while simultaneously bypassing machines. Examples of these kinds of "adversarial machine learning" strategies include fooling spam detectors by adding "noise" text from popular news articles, books, or text from a different language; fooling image classification/OCR by rotating, cropping, and altering images, and replacing characters with look-alikes or sound-alikes. Compared to the effort required to train a machine learning model, crafting malicious inputs to fool models is much quicker and easier.

SECTION 2

Content Abuse: Where and How

> *Fraudsters use bots to engage in content abuse techniques at massive scale, and bad actors have gotten highly sophisticated at obfuscating their intentions and impersonating legitimacy with their fake and malicious accounts.*

Everywhere across the internet sites and platforms are accepting user-generated content, and in virtually every instance where there is legitimate content being uploaded, there are fraudulent counterparts. Some examples include:

ACCOUNT PROFILES

Legitimate account profiles often include attributes such as names, nicknames, email addresses, websites, social handles, and more. Through means such as account takeover, identity theft, and first, third, and synthetic identity fraud, malicious accounts can be hijacked or created anew, and subsequently used to spread toxic content.

For example, a fraudster might engage in URL Shortener Spam by placing a spam link in an account profile's public-facing details and masking its true identity with a URL shortener, to try and trick unsuspecting viewers into clicking. Fraudsters can use bots to engage in these kinds of techniques at massive scale, and bad actors have gotten highly sophisticated at obfuscating their intentions and impersonating legitimacy with their fake and malicious accounts.

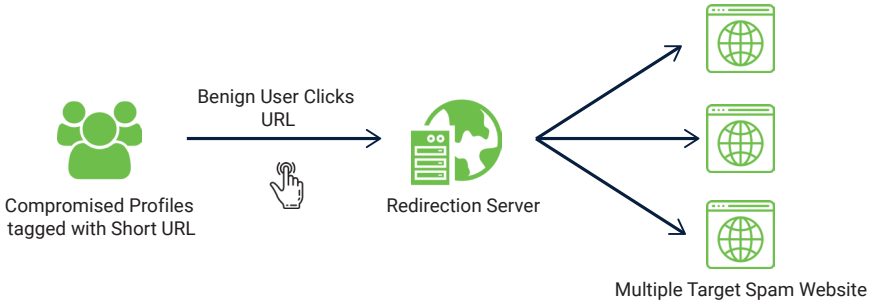

Figure 2: An example of a URL redirection attack where the landing server is reused to serve multiple spam campaigns.

POSTS, MESSAGES, AND COMMENTS

Today there are so many sites that depend heavily on authentic user-generated content as part of their business models—from social networks and ecommerce platforms, to content aggregators and online marketplaces. For all these businesses to succeed, users need to believe the content they interact with is real and trustworthy, and that the products or services they're considering are legitimate and truthfully represented.

In all these instances, users also need to feel safe to engage, ask questions, and leave comments, without fear of spam or abuse. Businesses need to provide seamless access for users to upload and post their content. However, this access—while necessary—also represents an exploitable vulnerability to fraudsters engaged in creating and disseminating spam, scams, fake listings, phishing, malware, and other types of fraudulent content.

Just how bad is this problem? On average, 7% of all posts, listings, comments, and messages contain malicious or fraudulent content. This number varies by the type of platform and the type of user-generated content. Social media and email service platforms—which are simultaneously mediums for disseminating information and channels for online communication—have a noticeably higher rate of content abuse than marketplaces or review sites whose content is typically better curated.

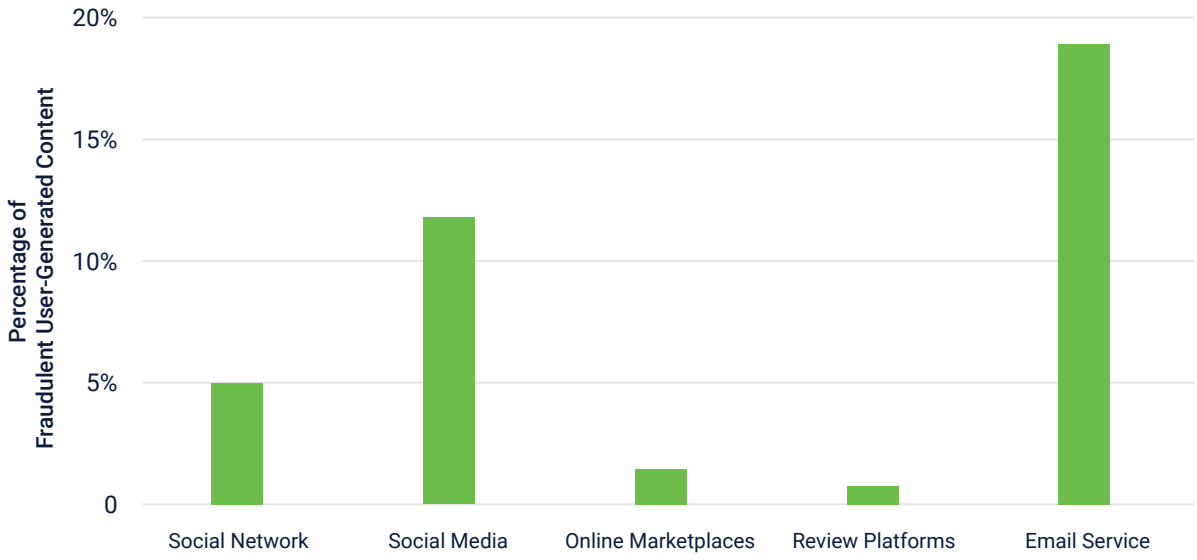

Figure 3: Fraction of user-generated content that is fraudulent (e.g., spam, scams, malware, phishing, fake products, fake reviews) on different types of online platforms.

There are also differences in how the three most common types of user-generated content—posts, comments, and direct user-to-user messages and chats—are used by fraudsters. Messages and chats make up around two-thirds of all user-generated content, but they have the lowest rate of abuse, at 3%. By contrast, the rates of abuse found in posts and comments are roughly 4x-5x higher at 13% and 15%, respectively. One reason for this is that the content of posts and comments can be directed to many users indiscriminately, unlike messages and chats, which are usually transferred between a sender and a recipient in a one-to-one manner. For fraudsters looking to spread malicious content at scale, posts and comments offer the most convenient option to reach many users simultaneously without having to compile a specific list of possible recipients.

In the next section, we'll explore in detail some of the ways content abusers use account features and platform formats to execute their schemes, including: Fake Names, Fake Purchase Information, Malicious URLs, and Fake Listings.

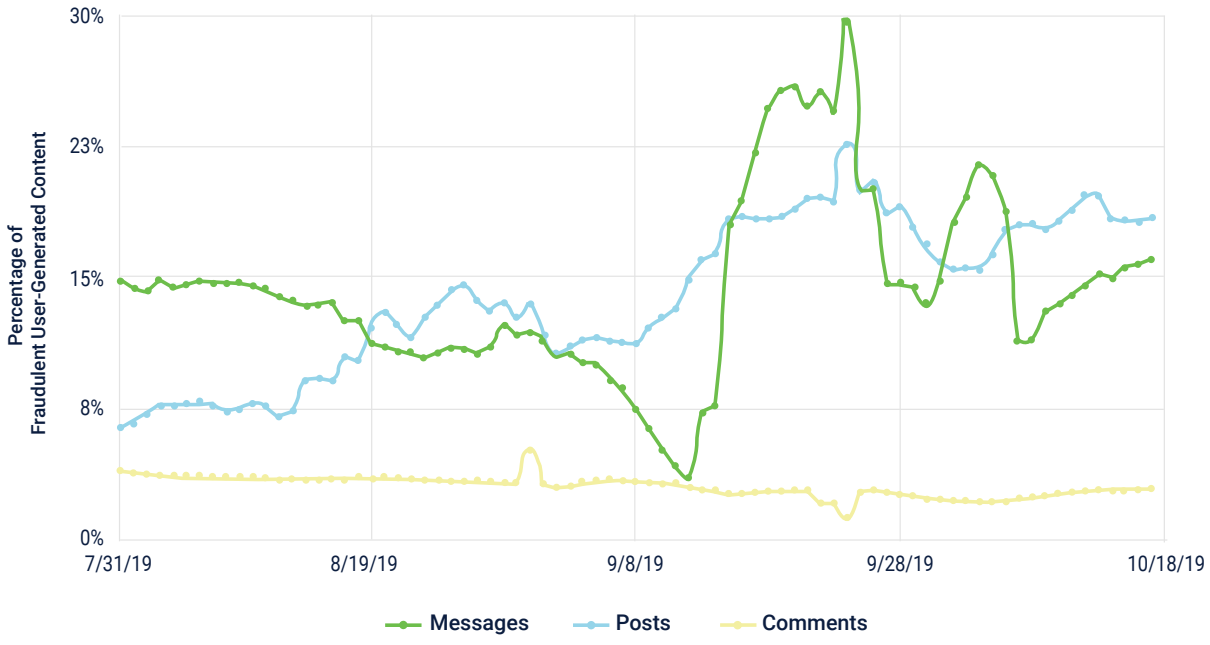

Figure 4: Daily rate of abuse for the three most common types of user-generated content: posts, comments, and direct user-to-user messages and chats. The rate of content abuse fluctuates over time, but posts and comments are consistently preferred by fraudsters over messages and chats. The fraud rate of the former is 13%-15%; 4x-5x higher than the latter.

SECTION 3

How Fraudsters Commit Content Abuse

FAKE NAMES AND IDENTITY COMPONENTS

Impersonation is critical to the success of content abuse efforts. In order to successfully infiltrate, permeate, and drain value from online platforms, fraudsters need accounts that appear legitimate. When these accounts fail to convince victims of their legitimacy, suspicions increase, and illicit plans are foiled. So, fraudsters do everything in their power to create and build authentic-seeming accounts and profiles.

To do so, they use a wide array of techniques, including:

- Posting legitimate-looking profile images (often, these are pictures scraped off of other public sites)
- Building large friend networks (usually comprised of other fraudulent accounts)
- Ensuring that each account originates from a different IP address and device
- Relying on common names and nicknames

Registration Time	Name	Email	Username	IP	Device
2017/6/15 9:10pm	Irma H.	irmaH512@gmail.com	Ih512_djs	67.198.236.72	iPhone 5 OS 9
2017/6/15 9:16pm	Carolyne F.	CarolyneF119@gmail.com	Cf119_wjd	67.198.236.74	iPhone 5s OS 9
2017/6/17 9:11pm	Celina P.	CelinaP130@gmail.com	Cp130_fue	108.171.209.68	iPhone 5s OS 9
2017/6/17 9:06pm	Ned A.	NedA91@gmail.com	Na91_euw	23.27.13.71	iPhone 5s OS 9
2017/6/18 9:15pm	Hilda G.	HildaG823@gmail.com	Hg823_ues	184.170.253.77	iPhone 5 OS 9
2017/6/18 9:21pm	Buford N.	BufordN42@gmail.com	Bn42_duw	67.198.236.174	iPhone 5 OS 9
2017/6/20 9:12pm	Paulita H.	PaulitaH617@gmail.com	Ph617_djf	108.171.209.92	iPhone 5s OS 9
2017/6/20 9:13pm	Earnestine G.	EarnestineG12@gmail.com	Eg12_edu	23.27.13.71.219	iPhone 5s OS 9

Table 1: Anonymized examples of coordinated fake accounts registered on an e-commerce platform. Each row corresponds to one account. All of the accounts appear legitimate individually, but when viewed together, they clearly show names and email addresses created from the same pattern.

Malicious accounts can be made to appear legitimate when viewed individually, but when viewed together, patterns emerge that clearly indicate coordinated attacks.

One of the key challenges when it comes to determining whether accounts are fake or legitimate is that, when viewed in isolation, the accounts appear authentic. However, when viewed together, patterns emerge that clearly indicate the machinations of fraudsters behind the scenes.

This is why contextual detection and holistic data analysis are critical capabilities for any advanced fraud management solution.

Mix of Characters and Numbers	Special Characters / Email Tags	"Numbered" Nicknames
alice123miller	alice+123m@gmail.com	alicetwentyfive
bob25smith	alice+456n@gmail.com	alicethirty
charlie597jones	alice+789p@gmail.com	alicethirtyfive
Different Domains	**Special characters / Dot**	alicefourty
alice123miller@google.com	alice123.miller@gmail.com	alicefourtyfive
alice123miller@hotmail.com	alice12.3miller@gmail.com	
alice123miller@aol.com	alice1.23miller@gmail.com	
	alice.123miller@gmail.com	

Table 2: More anonymized examples of scripted nicknames or email addresses used by fraudulent accounts.

FAKE PURCHASES

Purchase actions are a primary driver of revenue, but when manipulated or hijacked for malicious purposes, these actions can produce extremely negative impacts. Promotion abuse, for example, is an increasingly alarming concern with significant financial ramifications. In these attacks, fraudsters take advantage of sale items or promotion codes by directing massive numbers of fake accounts or bot requests to the target platform, artificially limiting availability of products and driving up prices.

Even though there is no obvious "content" involved in these type of attacks, fraudsters still need to provide information to fulfill the purchases; typically, the full name of the product recipient, an email address (for confirmation), and a shipping address. Just as with the fake names described in the preceding section, fraudsters can be very creative when it comes to generating fake shipping addresses. In a large attack DataVisor detected on an e-commerce site, thousands of fake accounts launched from mobile emulators attempted to make purchases for the same promotional product.

All of the bot purchases were "shipped" to addresses generated from the same template; a template comprised of several components:

- Random house or apartment number
- Common road name (e.g. Oak, Park, Washington)
- Direction (i.e.. North, South, East, West)
- Name of a large city/state
- (Optional) Name of furniture shop, bar, or restaurant

Each of these may well be a legitimate shipping address. However, it is highly unlikely for thousands of correlated users with scripted naming patterns to coincidentally ship to extremely similar locations as well, not to mention that none of the login locations match the shipping addresses.

This example shows that businesses need to think beyond traditional forms of user-generated content—the problem is no longer confined to social platforms. As noted earlier in this report, there are now vastly more "entry points" that enable fraudsters to introduce fraudulent user-generated content into online ecosystems. Effective fraud strategies must cover the full scope of potential attack points as well as address the problems from multiple angles.

Approximately **13%** *of the posts uploaded to marketplaces by fraudulent accounts contain spam or phishing URLs, though this number can be as high as* **87%** *on social network platforms.*

PHISHING AND MASKING

No discussion of content abuse is complete without addressing phishing and spam. Many platforms offer means of online communication—often conducted between complete strangers. Direct messages, forums, comments, and feedback forms, all provide fraudsters with cost-free ways of injecting phishing and spam content into conversations across wide audiences. An innocuous message such as: "Good afternoon! Funds have arrived in your name. Click on the link to get it." could lead victims to cloned websites that trick them into giving up their financial login information.

Approximately 13% of the posts uploaded to marketplaces by fraudulent accounts contain spam or phishing URLs, though this number can be as high as 87% on social network platforms. To successfully post these high volumes of malicious content, fraudsters need to be able to get around detection systems. One common approach fraudsters use to evade detection and blacklisting is to use URL shorteners or other URL redirection mechanisms. This method hides the actual landing page of a given malicious URL. Through redirection, fraudsters can also serve different versions of a site, (e.g., based on parameters in the URL, to support multiple attack campaigns simultaneously).

Another approach to avoid detection is to host malicious pages on newly registered domains. Since new domains have a "neutral" reputation (i.e., there is no history associated with them)—and given that high numbers of new and legitimate domains continue to be registered—they are less likely to be blocked.

A closer look at spam and phishing URLs showed that 78% were registered within the last two years, with the most popular top-level domains (TLDs) being .site, .com, .ru, .tk, and .ga. During the time period analyzed for this report, 18% of malicious URLs used HTTPS, a secure extension of HTTP. Using SSL certificates available at no cost from services like Let's Encrypt, fraudsters can make their URLs appear more legitimate to trick victims into clicking on the links.

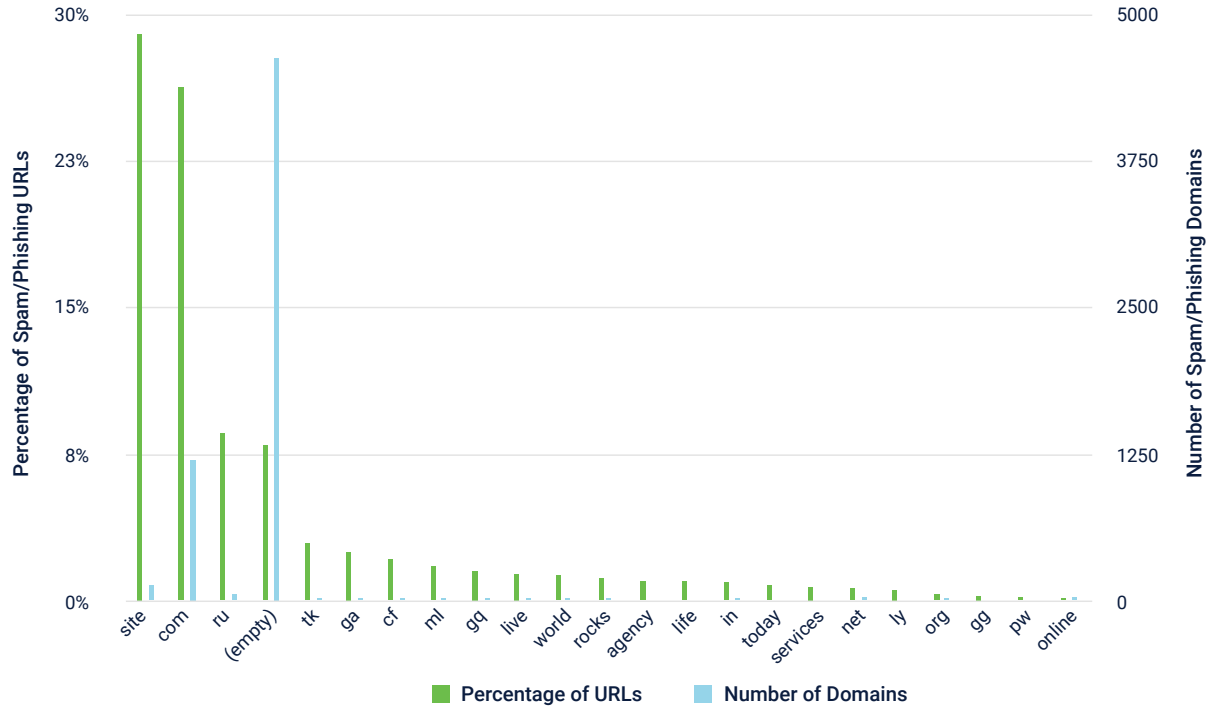

Figure 5: The most popular top-level domains (TLDs) used in phishing and spam URLs. The figure shows the percentage of phishing and spam URLs and the number of second-level domains associated with each TLD.

Among the malicious URLs, some attempt to piggyback off the reputation of established sites (e.g., pages created on blogspot.com) and some perform "typosquatting" attacks (i.e., URL hijacking) based on common typos or misspellings of popular sites (often registered on different TLDs, e.g., qoogle.site, bisney.live, outube.com).

Interestingly, some fraudsters intentionally omit the top-level domain or otherwise construct invalid URLs to avoid detection. Much like how security-conscious users replace the '@' symbol with "at" when sharing their email addresses publicly, fraudsters also replace the '.' symbol with "dot" when advertising their spam or phishing URLs (e.g., qoogle dot com). This indicates that tech-savvy fraudsters are knowledgeable about the limitations of machine learning solutions, and that they craft their attacks accordingly.

FAKE LISTINGS

For marketplaces and e-commerce sites, content abuse can have serious consequences—both for businesses and their customers—and the negative impacts can be both financial and reputational. Fraudsters can post fake listings for counterfeit products, launch scams involving advanced fees, or lure buyers off platforms to conduct under-the-table (and unsafe) transactions.

One of the common attack techniques is to use bots to control fraudulent accounts. This enables fraudsters to launch massive waves of attacks that comprise tens of thousands of accounts. In a fraud ring discovered by DataVisor, 5,000+ accounts were used to post suspicious luxury watch listings. All of the accounts had similar behaviors (e.g., 2-3 login events before each post, where the time interval between consecutive logins followed a very narrow distribution).

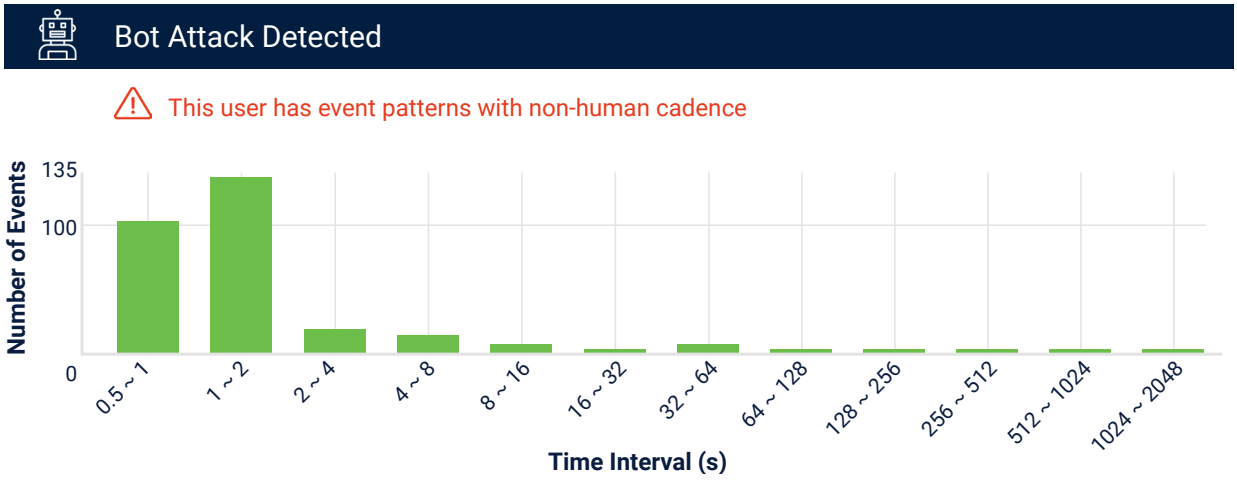

Figure 6: The distribution of inter-event timings for fraudulent accounts selling fake goods.

In addition to leveraging bots for scalability and automation, fraudsters also obfuscate their content so that each listing contains slightly "different" text or images.

Often random characters are added to the post, which does not affect readability for a human but can confuse ML models and negatively impact their effectiveness.

Figure 7: Anonymized examples of fake listings posted by fraudulent accounts. Each listing has very similar—but slightly different—content, obfuscated using random characters appended to the end of the listing description.

Fraudsters posting fake listings move quickly: 60% of fraudulent accounts generated malicious content within two hours of registration, and 76% did so within 24 hours of account registration. This is much faster than fraudulent accounts in general, where only 54% launch attacks within 24 hours of registration.

Accounts used for fake listings on marketplace sites tend to place more emphasis on the posting content (images, descriptions, and pricing) to attract users. The listings are often associated with trending products (e.g., the latest iPhone) and so the fraudsters' actions need to transpire in a timely manner. Because of these reasons, these fake accounts do not rely as much on account-level reputation and have less need to incubate.

Fraudsters posting fake listings move quickly: ***60%*** *of fraudulent accounts generated malicious content within two hours of registration, and* ***76%*** *did so within* ***24 hours*** *of account registration.*

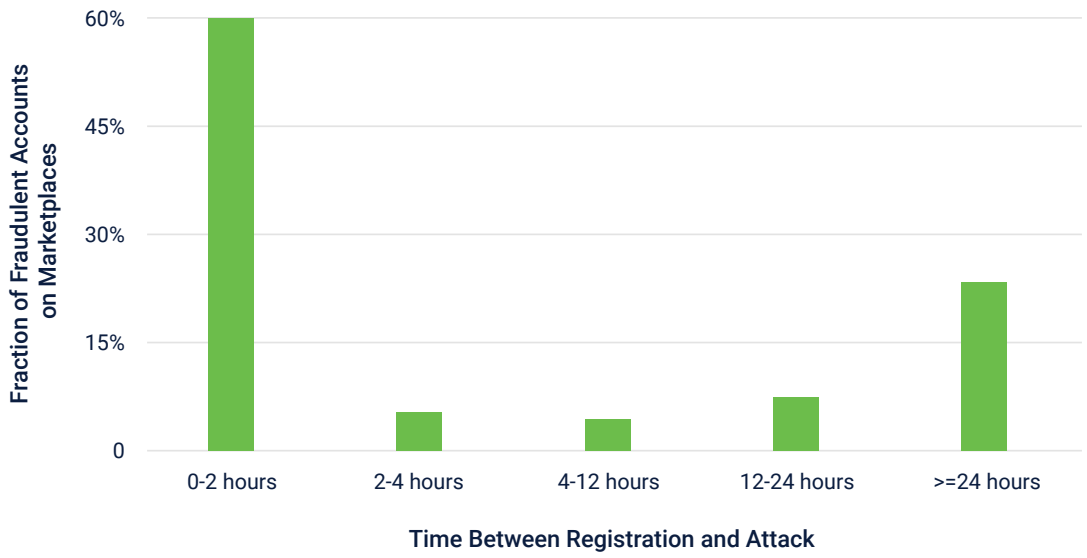

Figure 8: The "sleep" time distribution of fraudulent accounts used for attacks on marketplaces. 60% of fraudulent accounts posted or sent malicious content within two hours of registration, and 76% did so within 24 hours of account registration.

SECTION 4

Detecting Auto-Generated Texts with Deep Learning

> *Deep learning models are able to automatically transform input data into high-dimensional representation, and the technology uses multiple "layers" to learn complex concepts and patterns in large amounts of data.*

User-generated content provides fraudsters with an easy point-of-entry to inject malicious content into online platforms. Fraudsters are able to create and post massive amounts of content quickly, and adjust their techniques with equal rapidity. This makes it exceedingly difficult for fraud solutions that rely on hand-crafted features, rules, and blacklists to keep pace. More often than not, by the time new rules are written and new features are pushed out, fraudsters will have long since moved on.

Deep learning has emerged as a promising alternative, and represents a viable and innovative approach to analyzing user-generated content. The advantage of deep learning models is that they are able to automatically transform the input data into high-dimensional representation, and the technology uses multiple "layers" to learn complex concepts and patterns in large amounts of data.

At DataVisor, we observe that fraudsters typically automate the generation of content for fake accounts under their control. This creates patterns that are discernible with the right technologies and solutions in place. When components such as introductions, messages, names, or nicknames are more similar across accounts than what is generally to be expected for unrelated users, this is indicative of coordinated activity. Utilizing a combination of UML technology and deep learning, we train models to recognize when a group of user accounts share suspiciously "similar" content—without explicitly defining what "similar" means.

As an example, a deep learning model can identify texts generated from the same script/pattern without knowing beforehand what that pattern is. Consider the following account nicknames:

- alice826n42302
- bob4400a42284
- charlie39b44720

All these nicknames have the same pattern: name + 2-4 digits + one letter + 5 digits. If we knew this pattern, we could write a rule to detect these fake accounts, but it would be difficult to continue keeping the rule up-to-date as new patterns emerged. With a deep learning network, however, the model transforms "alice826n42302," "bob4400a42284," and "charlie39b44720" into high-dimensional representations. Through the multiple layers of the neural network, the model quantifies how "close" these strings are, and outputs a score indicating the similarity of strings for this group of users.

At DataVisor, we have designed a novel deep learning architecture to train models able to identify previously unknown suspicious patterns across alphabets (e.g., Latin, Chinese), languages (e.g., English, Spanish, Turkish), and content types (e.g., short messages, but also full names and emails) across online platforms spanning multiple industries (e.g., e-commerce, social networks, financial sites, and online marketplaces). Using this approach, we are able to address content abuse at the source by flagging and neutralizing suspicious accounts before they're ever used to generate malicious content.

FRAUDULENT ACCOUNTS WITH SUSPICIOUS PATTERNS

On average, 63% of fraudulent accounts exhibit email or name/nickname patterns that are suspiciously "similar" to other fraudulent accounts. This number can be as high as 87% in some cases, especially on online platforms experiencing large waves of coordinated fraud attacks. Those attacks are typically carried out by bots controlled via scripts or other automation software, where the fake accounts' profile information (as well as their attack activities) are generated programmatically.

The majority of users with suspicious email or nickname patterns are detected in large, coordinated fraud rings. Typically, the fraud rings are bimodal—they are either made up entirely of users exhibiting suspicious email or nickname patterns, or none at all (likely fake accounts orchestrated manually, or compromised accounts that belong to real users).

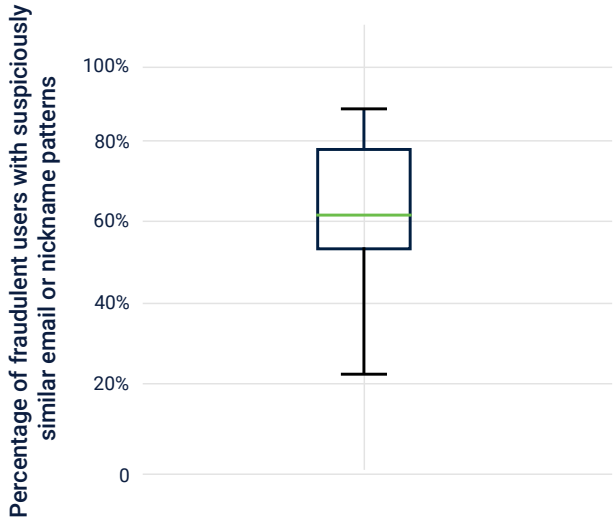

Figure 9: The distribution of the percentage of fraudulent users with suspiciously similar email or nickname patterns on each online service and/or platform.

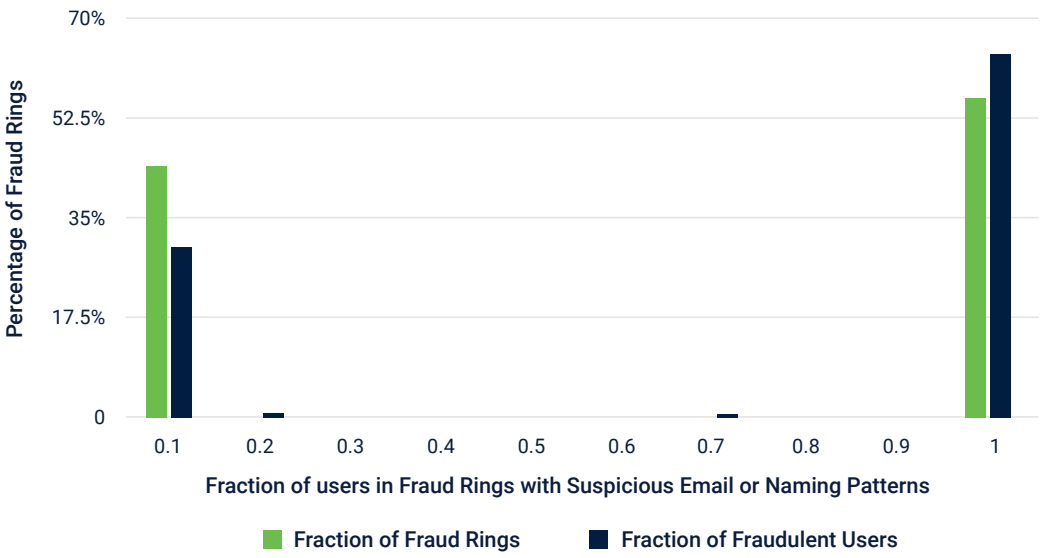

Figure 10: The distribution of the fraction of users with suspicious email or naming patterns in each fraud ring. The distribution of largely bimodal—the fraud rings either consist entirely of users with suspicious patterns, or none at all.

A closer look at the sizes of these fraud rings also shows an interesting phenomenon. Fraud rings made up of users exhibiting suspicious email or nickname patterns are on average 1.8x larger than those without the patterns.

This shows that with the help of scripts and automation software, fraudsters are able to launch bigger attacks with the potential to do much greater damage.

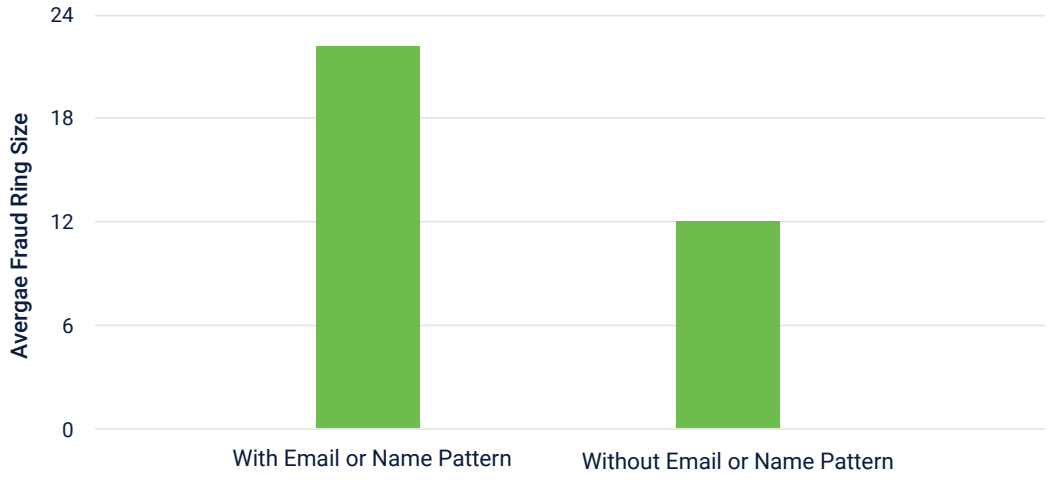

Figure 11: The average size of fraud rings that contain users with suspicious email or name patterns, versus those without. Fraud rings with suspicious patterns are 1.8 times larger than those without the patterns.

Conclusion

The challenges facing businesses and platforms that rely on, and incorporate, user-generated content are, in many ways, unique. At the same time, however, there are commonalities that transcend the boundaries of industry or economic sector. For example, financial services providers don't, for the most part, rely on user-generated content, and they don't accordingly have to contend with content abuse in the way a social platform does. However, they do have to finesse a similar balance of customer experience and risk management. In the same way reviews platforms, for example, have to make user convenience a top priority to preserve and promote business growth and to remain competitive, financial services providers are under increasing pressure to deliver products and services that optimize for ease, efficiency, and access. And in the same way that social platforms inadvertently open the door to malicious exploitation when they make their platforms open and accessible, financial services providers expose themselves to greater degrees of risk when they roll out easy-access mobile banking services and rapid review of online loan applications.

If there is one thing certain about online business, it's that innovation is the key to success. Staying competitive means constantly finding new ways to meet user demand, solve user concerns, and deliver seamless experiences. Because user experience is defined in great degree by convenience, it's virtually inevitable that online platforms are only going to become more porous, not less. Already in Europe, with PSD2, this is a matter of regulation, in that European banks must now open their data and infrastructure to fulfill regulatory requirements.

In a climate like this, businesses today can't even know the innovations they're going to be introducing in the future. What they can do, however, is prepare their organizations to protect against the threats they're inevitably going to invite as a byproduct of their user-centric innovations. Supervised machine learning and rules-based systems are already outmoded, and while they retain value when integrated into a fully comprehensive fraud management solution, they are on the edge of extinction as standalone strategies. For those platforms that depend on user-generated content, and which are dealing with major content abuse concerns, AI and unsupervised machine learning-powered solutions hold the key to a safe and secure future in which customers can contribute without fear of abuse.

About DataVisor

DataVisor is the leading fraud detection platform powered by transformational AI technology. Using proprietary unsupervised machine learning algorithms, DataVisor restores trust in digital commerce by enabling organizations to proactively detect and act on fast-evolving fraud patterns, and prevent future attacks before they happen. Combining advanced analytics and an intelligence network of more than 4B global user accounts, DataVisor protects against financial and reputational damage across a variety of industries, including financial services, marketplaces, ecommerce, and social platforms.

For more information on DataVisor:

info@datavisor.com

www.datavisor.com

967 N. Shoreline Blvd. | Mountain View | CA 94043

DATAVISOR