sift

Q2 2020 DIGITAL TRUST & SAFETY INDEX

Content Abuse and the Fraud Economy

Contents

E-Commerce and Fraud Economics

Content abuse has played a role in e-commerce for as long as digital businesses have existed. Scams, spam, fake reviews, and misinformation have evolved in tandem with online shopping, discussion forums, and social networks, with the methods behind them becoming smarter and more malicious as fraudsters adapt their techniques to bypass advanced security measures and invade new channels.

As a result, many companies have found out the hard way how damaging spam and scams can be for business. With limited expertise and few tools available to help them deal with such a pervasive problem, merchants—and their customers—are left to navigate an internet where anyone can promote and sell just about anything.

But as inescapable as content abuse might seem, it's not a standalone threat. What happens *after* someone interacts with malicious content is what's truly disruptive; and to understand fake content's vital role in payment fraud and account takeover, it's necessary to dig deep into the behaviors and assumptions that shape it.

The data in this report is derived from Sift's global network of customers, representing over 34,000 sites and apps using Sift across all of e-commerce, as well as a survey of over 1,000 consumers* conducted in June 2020. By reviewing and analyzing signals surrounding content abuse from January through May of 2020, and coupling them with these buyer insights, this report gives online merchants visibility into the covert economics that impact business—along with industry expertise to help you protect customers without losing money or momentum.

**On behalf of Sift, Dynata polled 1,000 adults across the United States via online survey, age 18+, between June 1 and June 8, 2020.*

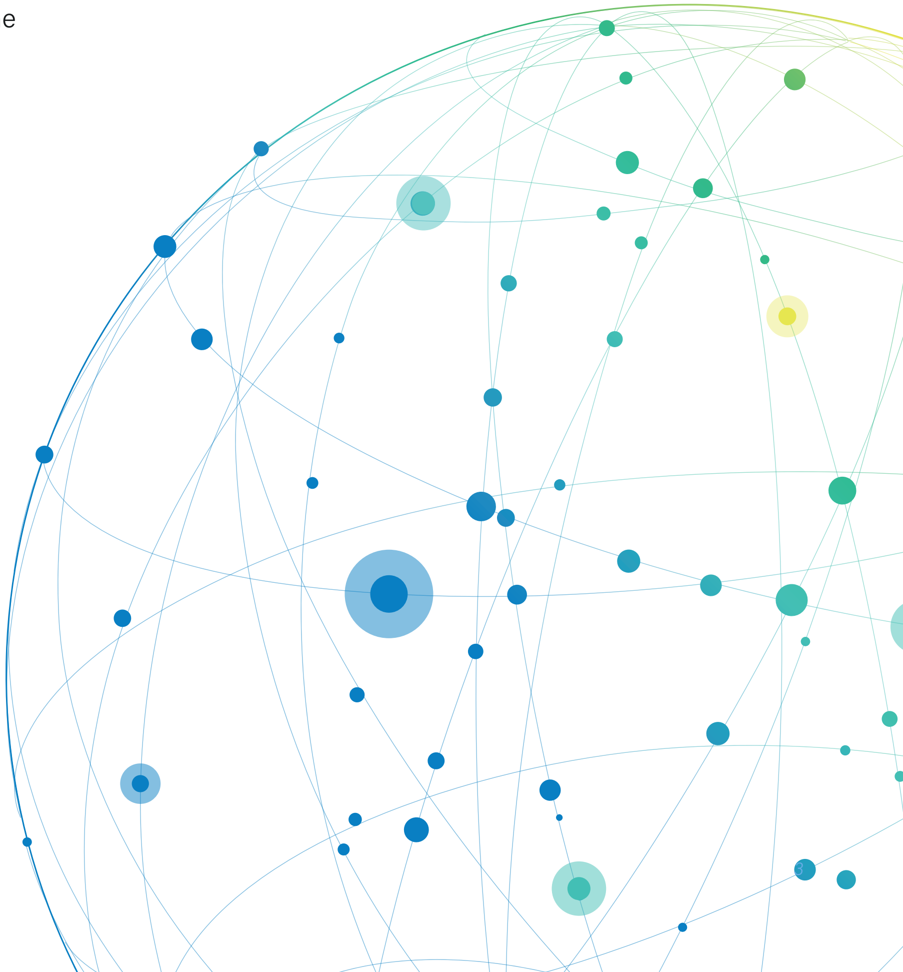

Making Sense of the Fraud Supply Chain

In many cases, content abuse is a means to an end—a stepping stone used by fraudsters to commit payment fraud. When they create a post, comment, email, or text message to disguise a malicious link or drive consumers to unsecured sites and media, it's not enough that someone sees it. The attack only works when people engage with that content and link, either spreading it further across the internet by sharing it, or by clicking on it themselves. One action widens the pool of potential victims, while the other directly impacts the person who clicked.

In order for cybercriminals to make money off of their attacks—outside of using stolen data themselves to directly commit account takeover and payment fraud—there has to be a market for the information they steal. And there is: a very large corner of the internet commonly known as the dark web. This fraudster flea market is essentially an illicit mirror image of digital e-commerce. For the right price, anyone can purchase valuable data from compromised accounts or sell information they've hacked or heisted.

Accessing this internet underbelly isn't difficult to do, but the items and supplies sold there are often hard to detect. The typical gateway between here and there is The Onion Router (a.k.a. "Tor"), an anonymous browser that uses multiple layers of encrypted connections to mask the identity of its users. Once inside, fraudsters on a buying spree can purchase virtually anything they might need to organize and successfully hit vulnerable businesses and excavate user accounts for payment details or other data. Dark web shoppers can even find sellers that offer organized fraud attacks as a service. Do-it-yourself criminals might prefer to grab up templates that enable the creation of fake online storefronts that look exactly like the real thing—the major difference being that the forgeries include built-in malware designed to capture personally identifiable information (PII). Finally, for fraudsters on a budget, there's free, customizable commodity malware that can be tailored to any site they'd like to target.

It's important to note that buying up compromised data—especially in bulk—is itself an investment for money-motivated fraudsters. Gaining access to legitimate credentials and payment details can enable high-value purchases of goods and services, which fraudsters can then use, trade, or resell as necessary. In some cases, these attacks are especially sophisticated and scalable, and can lead to repeat and diversified fraud attempts against a single business or consumer.

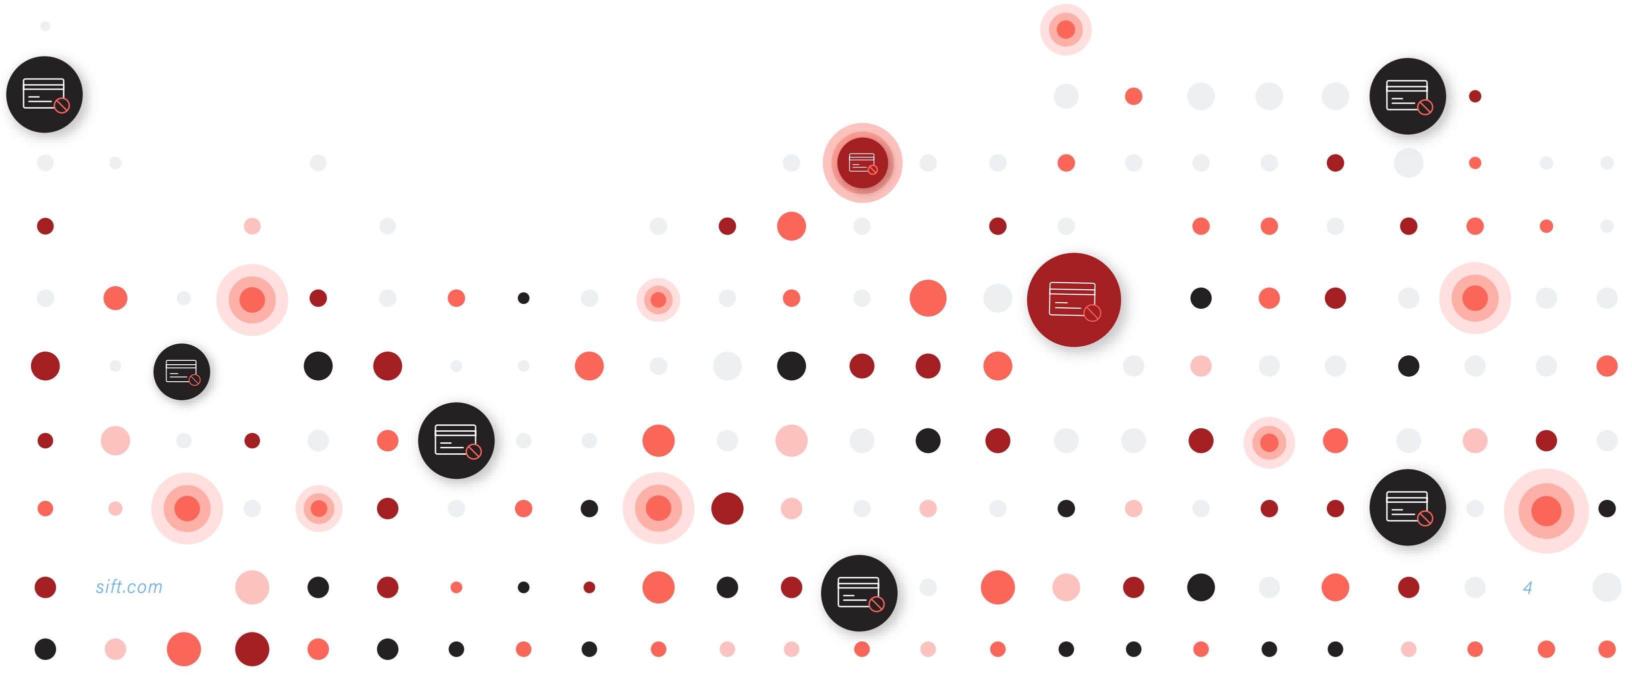

The Fraud Supply Chain: A Network of Content Abuse, Account Takeover (ATO) , and Payment Fraud

The dark web is subject to the same supply-and-demand philosophy that governs any marketplace, and uses a type of supply chain model that makes content abuse an ideal vector for committing financial fraud. It looks like this:

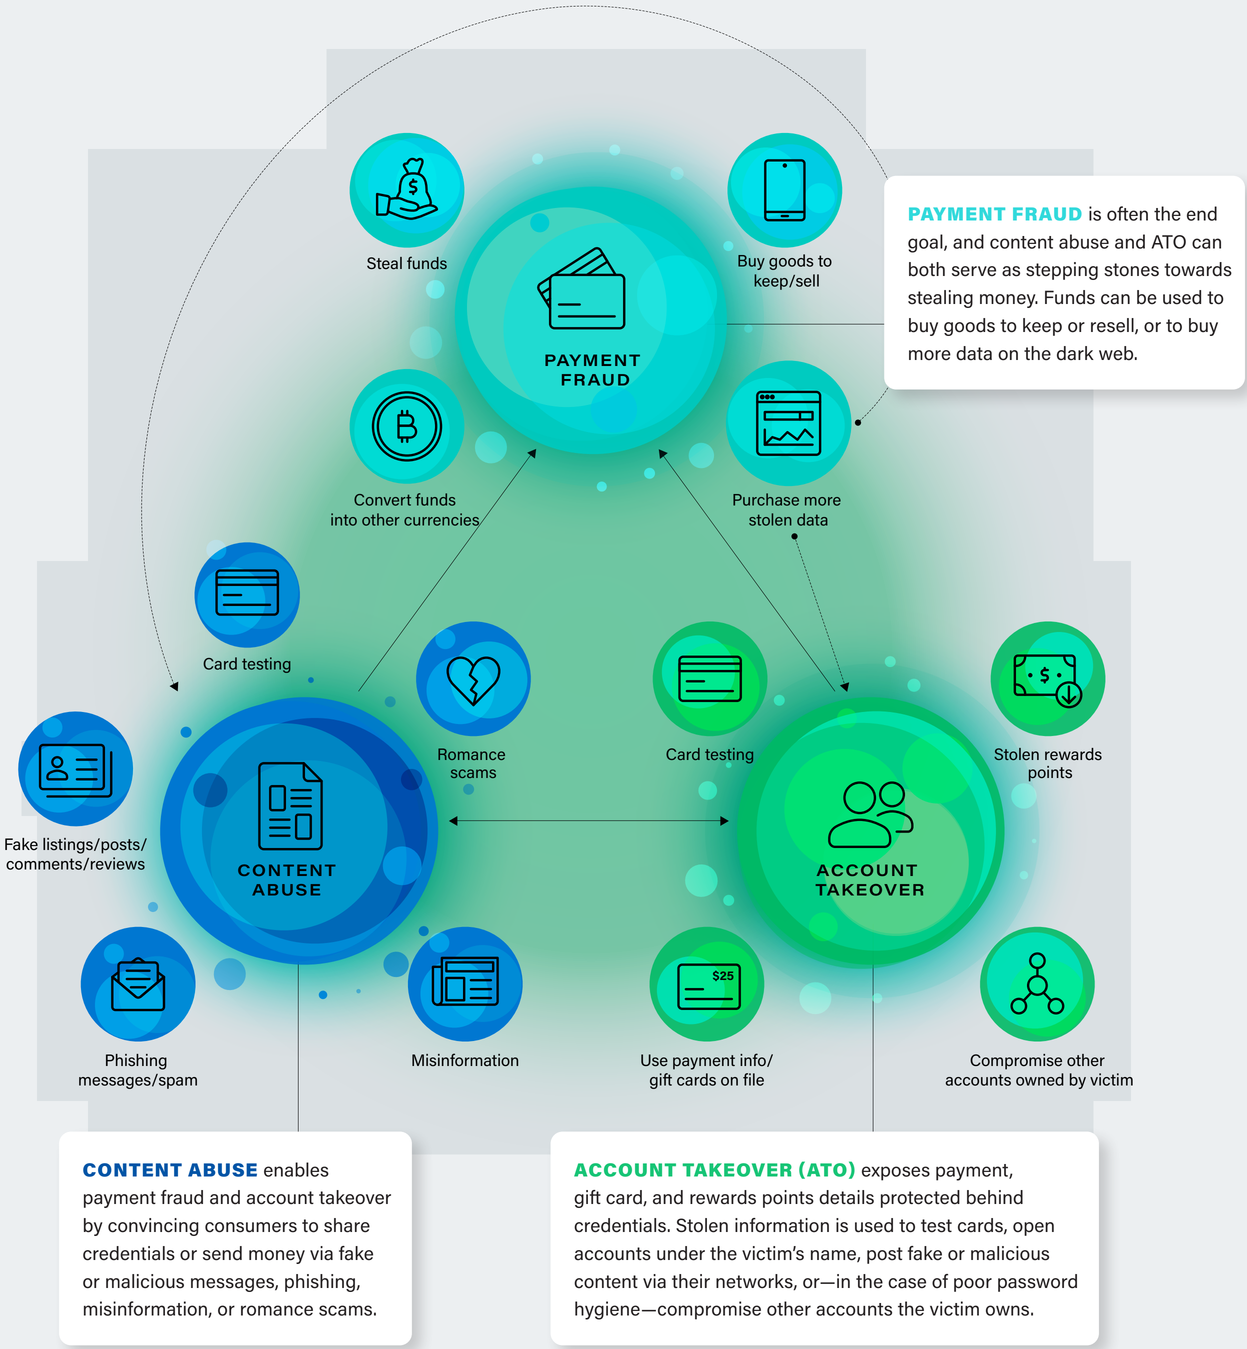
Steal funds
Buy goods to keep/sell
PAYMENT FRAUD
Convert funds into other currencies
Purchase more stolen data
Card testing
Romance scams
Card testing
Stolen rewards points
Fake listings/posts/ comments/reviews
CONTENT ABUSE
ACCOUNT TAKEOVER
Phishing messages/spam
Misinformation
$25
Use payment info/ gift cards on file
Compromise other accounts owned by victim

PAYMENT FRAUD is often the end goal, and content abuse and ATO can both serve as stepping stones towards stealing money. Funds can be used to buy goods to keep or resell, or to buy more data on the dark web.

CONTENT ABUSE enables payment fraud and account takeover by convincing consumers to share credentials or send money via fake or malicious messages, phishing, misinformation, or romance scams.

ACCOUNT TAKEOVER (ATO) exposes payment, gift card, and rewards points details protected behind credentials. Stolen information is used to test cards, open accounts under the victim's name, post fake or malicious content via their networks, or—in the case of poor password hygiene—compromise other accounts the victim owns.

In this fraud supply chain, profit-driven content abuse acts as both a springboard for, and a bridge between, account takeover and payment fraud. The information procured via malicious content can be used in a number of ways, depending on how detailed it is:

To orchestrate account takeover attacks on a person or group of people and access additional data

To supplement data that's been stolen elsewhere via a different attack or after being purchased on the dark web

To directly steal money, gift cards, rewards points, or other currencies. Sometimes, data merely gets sold on the dark web for a healthy profit

One of the more covert ways content abuse fits back into the fraud supply chain involves card testing, which takes place after login information, gift card details, or payment data has been stolen via phishing or bought. Using accounts they've taken over or created, fraudsters will attempt to submit hijacked payment details to a vendor and purchase goods or services—typically low-value, low-effort items, like a single pair of shoes or a pizza. If the purchase is successful, they know the payment information is valid—meaning they can now start upping the cost of their illicit purchases to determine how much value they can extract from the looted data, or to buy expensive items they can then resell for 100% profit.

This doesn't usually happen one card at a time, either. Often, a large group of fraudsters (a fraud ring) will work together to test multiple, sometimes thousands, of associated payment details by posting bogus content listings on digital marketplaces. Using these fake listings, they'll "sell" items to each other in order to vet stolen data, even "negotiating" the costs of those items down so that the exchanges appear more legitimate, while still allowing them to test payment info using minimal transaction amounts.

In fact, Sift's data team identified precisely this type of operation in early June 2020.

The Fellowship of the (Fraud) Ring: Exploiting a Two-Way Marketplace

Sift's Data Science team discovered something sinister unfolding on an e-commerce marketplace: a fraud ring using fake content listings to execute a classic card-testing scheme. Several users with identical IP addresses created multiple listings at a price point of $99.00 USD. To test dozens of stolen cards, they "sold" the items to each other, after "haggling" those prices down to $1.00 USD—a typical price used to test hijacked payment details. Each listing was uncharacteristic for this marketplace, purchased on the same day, and included several fake reviews to strengthen the appearance of authenticity. Interestingly, this fraud ring—nicknamed Bargain Bear by Sift—was based out of Russia, which is not one of top five regions known for content-based fraud. Sift worked with the impacted customer to refine scoring thresholds and block the ring from their site, as well as our entire global network—a crucial move, given that it is extremely likely this fraud ring was setting up similar attacks across the internet and well beyond this specific marketplace. Finally, our machine learning models were updated with these findings in real-time, enhancing their efficacy for all Sift customers.

Content Fraud in Action

In June 2020, Sift identified and stopped a fraud ring that was using fake content listings to test dozens of stolen debit and credit cards in order to see (1) if they worked at all, and (2) how much they were worth. Exactly how those payment details were compromised is unknown, but the data was likely purchased on the dark web.

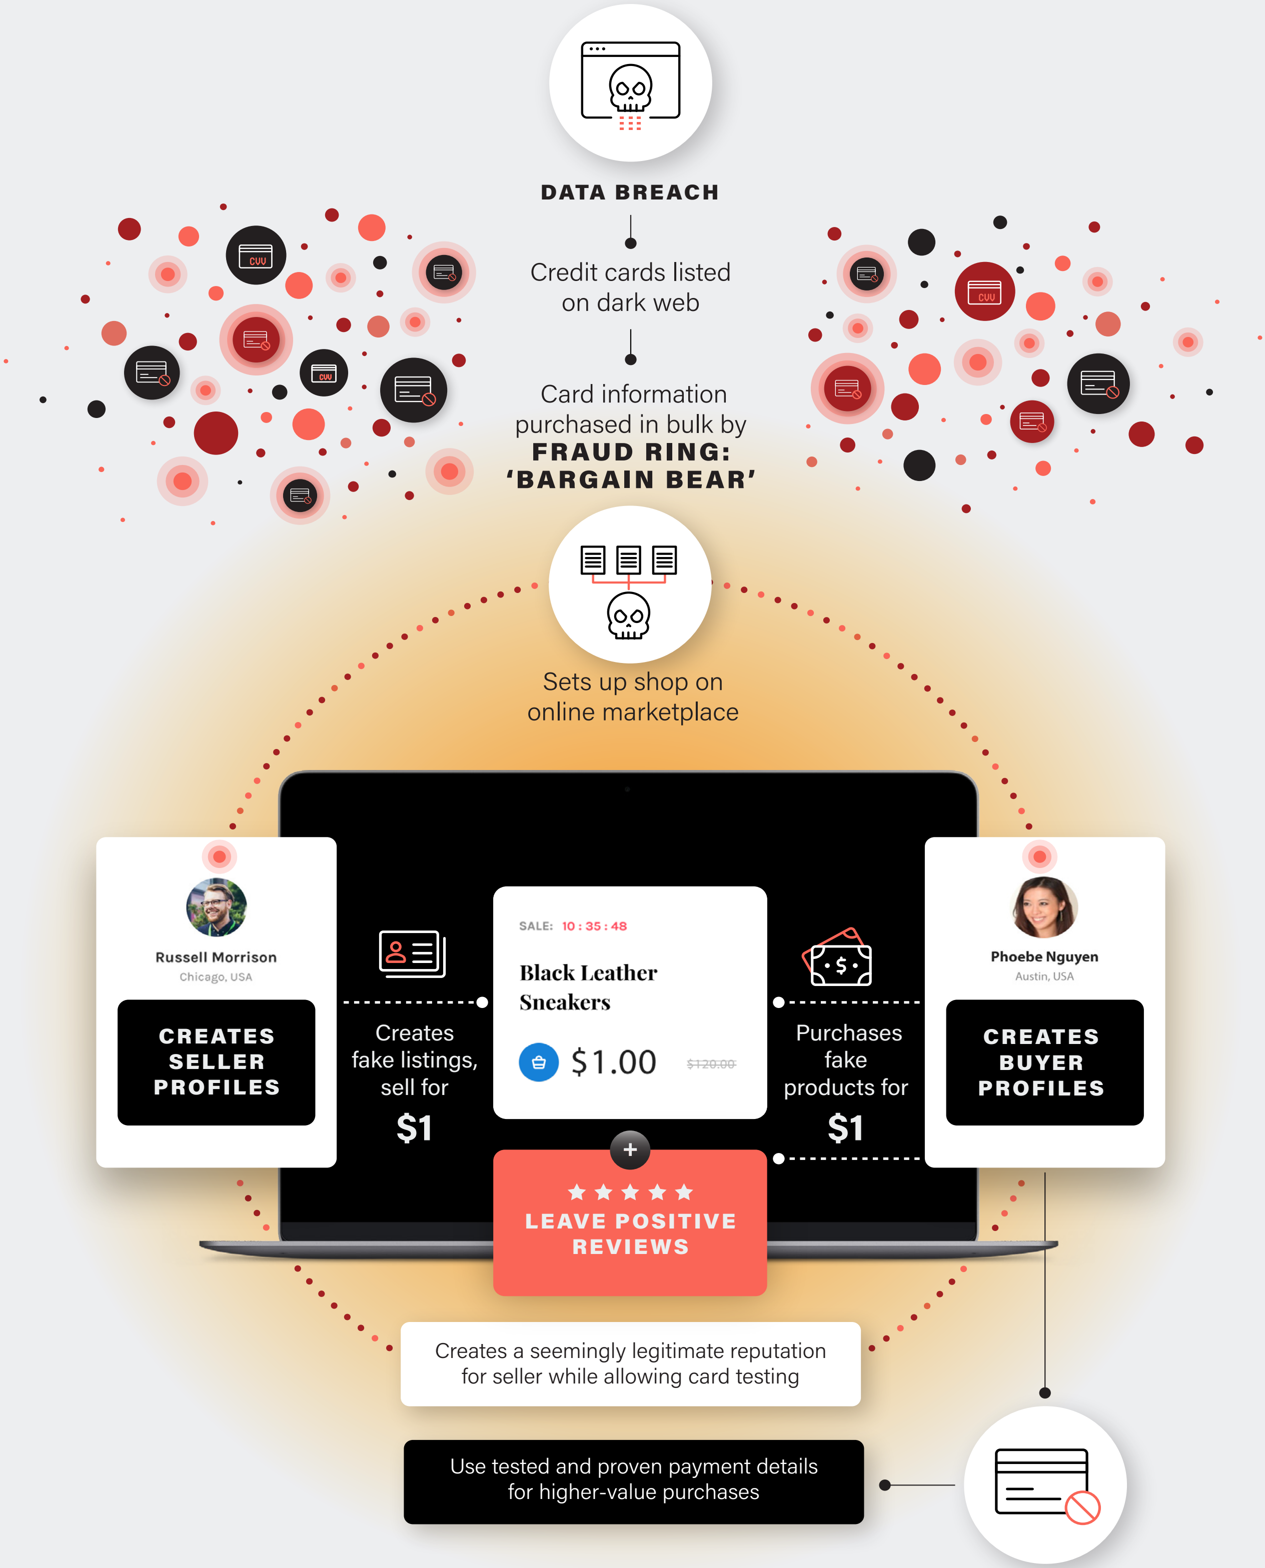

DATA BREACH

Credit cards listed on dark web

Card information purchased in bulk by **FRAUD RING: 'BARGAIN BEAR'**

Sets up shop on online marketplace

Russell Morrison
Chicago, USA

CREATES SELLER PROFILES

Creates fake listings, sell for **$1**

SALE: 10 : 35 : 48

Black Leather Sneakers

$1.00 ~~$120.00~~

Purchases fake products for **$1**

Phoebe Nguyen
Austin, USA

CREATES BUYER PROFILES

★★★★★ **LEAVE POSITIVE REVIEWS**

Creates a seemingly legitimate reputation for seller while allowing card testing

Use tested and proven payment details for higher-value purchases

sift

There are a couple of reasons this illicit supply chain is detrimental to e-commerce merchants, rather than just the individuals whose personal details and funds are being bootlegged. Theoretically, compromised customer information can enable hackers to bypass security checks, putting other merchants' platforms and proprietary data at risk, and potentially ushering in system-wide data breaches. For example, a fraudster could determine where a victim works with a quick Google or LinkedIn search. Suppose that the victim is using the same login information for both personal and professional profiles (as 62% of employees do)—the fraudster could then take steps to gain entry to those business accounts, and subsequently access a company's protected information. In fact, poor password hygiene and weak credentials are a primary cause of large-scale breaches.

And when a breach happens, it often results in multiple types of customer data being unveiled—sometimes including the account, payment, and personal information of thousands or even millions of people. If any of those customers' details and credentials are being used on multiple sites or platforms, there's even a possibility that other merchants could be impacted.

But malicious content can have major negative repercussions for a business even if it doesn't lead to something as dramatic and messy as a data breach. Consumers reported that if they discovered that their personal information had been exposed as a result of a scam on a website, **56% of them would stop using the site or service and choose a different provider.**

The True Cost of Content Fraud: Brand Abandonment

Content abuse can damage brand loyalty, thereby stunting profits and long-term growth. When we asked consumers how they'd react if their information was compromised through content abuse, over half of respondents said they'd stop using the impacted site or service and seek out a competitor.

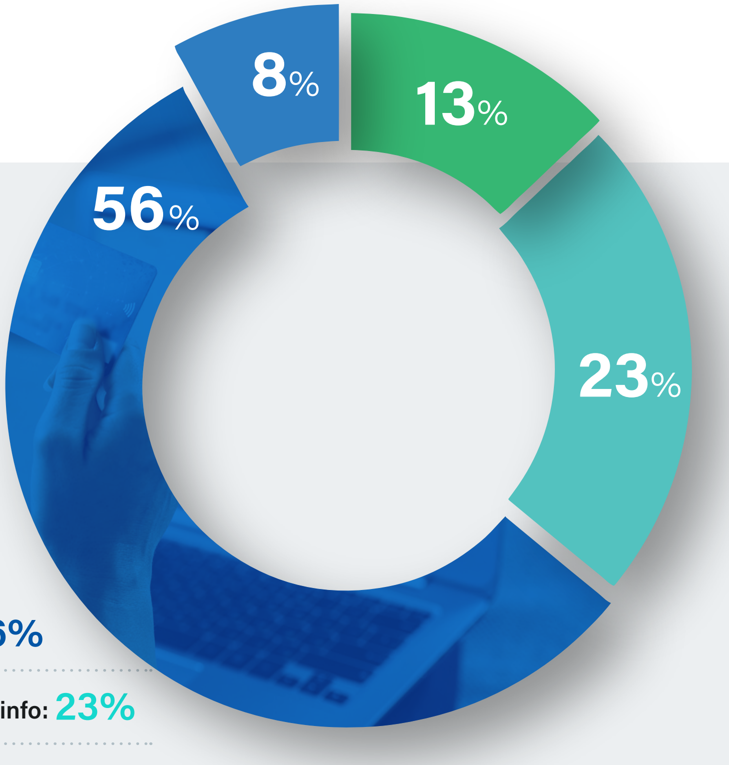

- **Stop using the site or service and select another provider:** 56%
- **Keep using the site/service, but change credentials/personal info:** 23%
- **Keep using the site/service, and contact support:** 13%
- **No change in behavior:** 8%

Our consumer survey results clearly indicate that content abuse can do a great deal of long-lasting damage to a brand, well beyond undermining the trust merchants work so hard to build with their customers. When such a large percentage of a company's customers would readily—and permanently—spend their money elsewhere after being impacted by content abuse, it's not just existing revenue that gets sliced in half; customer acquisition costs (CAC) will go up, and the lifetime value (LTV) of each of those lost customers, as well as the advocacy they could have provided your brand, is forfeited, too. Finally, these consequences could be compounded by multiple attacks or other types of fraud that might result, spelling disaster for merchants whose platforms have been infiltrated by spam, scams, fake reviews, and other fraudulent content.

Fake Content: A Digital Epidemic

Between January and May of 2020, digital marketplaces were hit hard by content fraud. Sift global network data shows that, during this time period, attempted content abuse rose by **109%** as compared to that same time period in 2019. An analysis of abuse types shows, unsurprisingly, that much of this fraud is financially motivated, with scams making up **46.8%** of the content abuse blocked by Sift. And while scams are the type of content abuse that most directly enables payment fraud and ATO, all content-based fraud can serve as a jumping off point for both.

Digging deeper, it's clear that content abuse is becoming disproportionately problematic for digital goods merchants. This includes online learning, streaming entertainment, and even online donation platforms—which, despite their best intentions, are seeing scammers attempt to steal money from unwitting donors.

Content Abuse by Vertical: Jan-May 2020

The below data is segmented by vertical, and illustrates the percentage of user-generated content posted across our customers' websites that turned out to be fraudulent—all of which was blocked by Sift.

11.2%	Ticketing & Events
8.9%	Digital E-Commerce
1.34%	Local Services
0.79%	Ads & Marketing
0.68%	Business Services & Tools
0.5%	Online Donations
0.27%	Physical E-Commerce

Consumers Find Fake Content Everywhere They Go

Digital consumers are wary of being scammed on the internet. When asked where they're most watchful, respondents said that social media channels, classifieds, and digital marketplaces are the most risky—but it's clear that people are increasingly cautious about interacting with user-generated content no matter where it shows up.

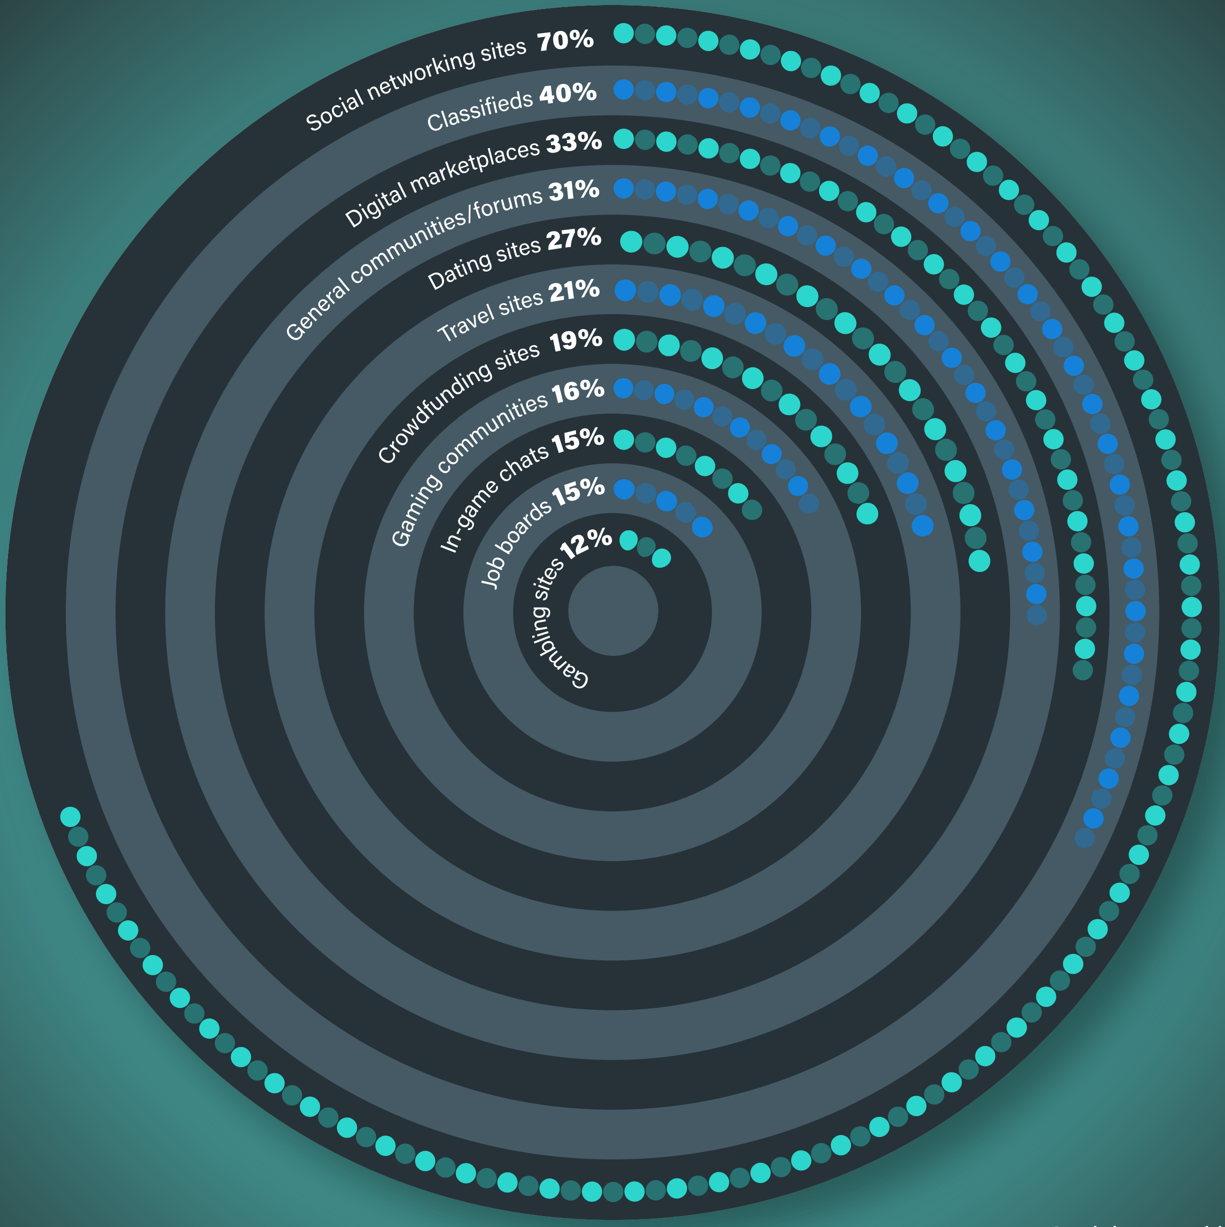

70%	Social networking sites
40%	Classifieds
33%	Digital marketplaces
31%	General communities/forums
27%	Dating sites
21%	Travel sites
19%	Crowdfunding sites
16%	Gaming communities
15%	In-game chats
15%	Job boards
12%	Gambling sites

Consumers aren't ignorant about the existence of fake content or its consequences, though. In fact, **67%** of those surveyed* believe they come across some type of fraudulent content or false information on a daily, weekly, or monthly basis, and **94%** of them deem content to be suspicious based on conspicuous factors: pie-in-the-sky promises, multiple typos or grammatical errors, outlandish claims, or a lack of identity information from the person posting it.

But fraudsters are paying attention, too. They know that people are increasingly on the lookout for fake content, which is why successfully tricking people into sharing personal details, payment data, banking information, or credentials requires some thought. Take, for example, one of the most sophisticated social media content scams ever perpetrated: the subscription trap.

The subscription trap is a lucrative scam, and works by convincing people to purchase what they believe is a single, free trial of a product or service. And while they do, in many cases, receive the product or service, buying that free trial means the unknowing customer has signed up for an expensive monthly subscription—one that is designed to be incredibly difficult to cancel, and that includes no way to recoup money spent. And, even if the buyer manages to stop recurring charges, some level of continuing damage is done: the fraudster behind the fake-out now has access to personal information and payment details they can use or sell elsewhere.

**On behalf of Sift, Dynata polled 1,000 adults across the United States via online survey, age 18+, between June 1 and June 8, 2020.*

Content Fraud Comes in All Shapes, Sizes, and Scams

Despite widespread media coverage of fake news as a primary source of content abuse, consumers are more concerned about spam and scams. Misinformation was third on the list of top methods they're watching out for, followed not-so-closely by phishing attempts and fake reviews.

69% Spam

Unrelated, irrelevant, inappropriate, or aggressively repetitive messages and materials

63% Scams

Content designed to trick users into sharing personal information and/or purchasing non-existent goods/services

50% Misinformation

Sometimes called "fake news," this fraudulent content is written, recorded, or designed to look as though it comes from a legitimate media source

42% Phishing

When fraudsters pose as legitimate users to trick their victims into giving up personal information

32% Fake listings

Posts of fraudulent listings or counterfeit goods on online marketplaces

28% Fake reviews

Often used for phishing, malware distribution, spam, and other harmful ends

15% Catfishing

When a scammer impersonates someone to gain a victim's trust

Where Does Content Abuse Originate?

The below data illustrates the top five countries by block rate where attempted content-based attacks originate, according to Sift global network data analyzed between January and May 2020. Clearly, fraudsters who use malicious content quite literally span the globe, with over 7,000 miles separating the top two: Canada and India.

① **24**% Canada

② **19**% India

③ **4**% United States

④ **1.5**% United Kingdom

⑤ **.73**% Israel

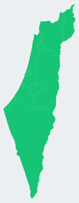

But the subscription trap is only one example, and isn't content fraud in and of itself. It simply uses everyday advertising copy as a gateway for scamming consumers, demonstrating how easy it can be for fraudsters to use content like a fisherman uses a hook. Despite consumer awareness of spam and scams, fake content still frequently worms its way into the average person's life—and in many ways, that's the whole point. Scams and spam attacks done well are effective because they aren't obvious, and when someone's data or information is compromised through fraudulent comments or posts, it's not necessarily because the victim made a mistake. It's because the perpetrator made an effort.

Content Abuse and COVID-19

The aforementioned **109%** year-over-year increase in content fraud is very likely connected to the global disruption caused by the coronavirus pandemic, and because fake content is so often designed to enable payment fraud, it's possible that the increased content abuse has influenced fraud rates and instances of payment fraud in the first half of 2020.

There's additional evidence to support that connection as well. Just a few weeks into the global lockdown, Sift blocked the highest number of fraudulent content attempts executed across all verticals between January and May. The spike occurred the week of April 4th to April 11th 2020—with April 9th being the fraudiest day of them all. Interestingly, the ticketing and events space was hit the hardest by attempted content abuse since the start of the year, while also experiencing record drops in event volume (down 84% from April 2019) as large gatherings of any kind became impossible. Sift's in-house fraud experts say that this spike was likely due to fraudsters attempting to exploit customer fear and unrest. People struggling with impacted incomes and strict lockdown orders could be seen as easy targets for travel scams offering free getaways, ticketing scams surrounding non-existent streaming concerts, or fake virtual events claiming to fund philanthropic causes.

Though ticketing merchants appear to have been especially hurt by pandemic-driven economic disruption, digital e-commerce merchants across the board have been, and continue to be, heavily targeted, demonstrating that fraudsters are always ready to take advantage of uncertainty.

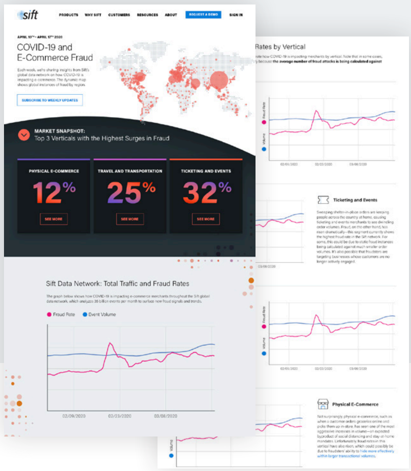

PRODUCTS WHY SIFT CUSTOMERS RESOURCES ABOUT REQUEST A DEMO SIGN IN

APRIL 10TH - APRIL 17TH 2020

COVID-19 and E-Commerce Fraud

Each week, we're sharing insights from Sift's global data network on how COVID-19 is impacting e-commerce. The dynamic map shows global instances of fraud by region.

SUBSCRIBE TO WEEKLY UPDATES

MARKET SNAPSHOT: Top 3 Verticals with the Highest Surges in Fraud

PHYSICAL E-COMMERCE	TRAVEL AND TRANSPORTATION	TICKETING AND EVENTS
12%	25%	32%
SEE MORE	SEE MORE	SEE MORE

Sift Data Network: Total Traffic and Fraud Rates

The graph below shows how COVID-19 is impacting e-commerce merchants throughout the Sift global data network, which analyzes 35 billion events per month to surface new fraud signals and trends.

Rates by Vertical

...rate how COVID-19 is impacting merchants by vertical. Note that in some cases, ...y because the average number of fraud attacks is being calculated against

Ticketing and Events

Sweeping shelter-in-place orders are keeping people across the country at home, causing ticketing and events merchants to see dwindling order volumes. Fraud, on the other hand, has risen dramatically—this segment currently shows the highest fraud rate in the Sift network. For some, this could be due to static fraud instances being calculated against much smaller order volumes. It's also possible that fraudsters are targeting businesses whose customers are no longer actively engaged.

Physical E-Commerce

Not surprisingly, physical e-commerce, such as when a customer orders groceries online and picks them up in-store, has seen one of the most aggressive increases in volume—an expected byproduct of social distancing and stay-at-home mandates. Unfortunately, fraud rates in this vertical have also risen, which could possibly be due to fraudsters' ability to hide more effectively within larger transactional volumes.

How COVID-19 Has Changed E-Commerce

Since March of 2020, Sift has been tracking how fraud rates and event volumes are changing each week across multiple e-commerce verticals in response to the pandemic. The data shown represents a 7-day moving average to illustrate the acute effects COVID-19 is having on online merchants. In some cases, slowdowns in traffic are driving fraud rates higher because the average number of fraud attacks is being calculated against declining event volumes.

Under the coronavirus umbrella alone, there have been some alarming examples of how fraudsters use content to deceive and exploit. They've used text messages to encourage stockpiling and sow fear about quarantine restrictions. They've sent emails to trick people into believing a vaccine exists and is being withheld, to proffer fake treatments, and to offer refunds to jetsetters whose plans were disrupted by travel bans. They've even used social media posts to pose as medical representatives with access to tests and antiviral medications—all for a nominal fee, of course.

On its own, this type of spammy content might seem relatively harmless, or clearly suspicious. But combined with fear, community unrest, and economic uncertainty, it can become its own type of virus, disseminating throughout the internet and greatly expanding the pool of potential victims from which fraudsters can extract valuable data. And when it comes to any financially-motivated type of content abuse, that's the ultimate goal: cracking open credentials, accounts, and other personally identifiable information that enables criminals to steal money and assets, and profit from their digital plunder.

Digital Trust & Safety in the Fraud Economy

Trust and safety teams must continuously adapt their strategies to effectively combat fraud, but when those efforts aren't focused on holistic prevention and mitigation, fraudsters have a far better chance of sidestepping even the most sophisticated online security measures. The findings in this report demonstrate exactly how pervasive, detrimental, and interconnected fraud can be throughout online marketplaces, as well as why it's crucial that merchants consider the entire ecosystem at play and address vulnerabilities across all areas of business.

Adopting a Digital Trust & Safety approach enables fraud fighters to better understand, internalize, and proactively prevent all types of fraud, including content abuse, account takeover, and payment fraud.

Look out for our next Digital Trust & Safety Index to stay informed about the evolution of fraud across e-commerce and how online merchants can protect customers, preserve revenue, and drive exponential growth. You can also read our Q1 2020 report here.

About Sift

Sift is the leader in Digital Trust & Safety, empowering businesses of all sizes, from digital disruptors to Fortune 500 companies, to unlock new revenue without risk. Sift dynamically prevents fraud and abuse through industry-leading technology and expertise, an unrivaled global data network of 35 billion events per month, and a commitment to long-term customer partnerships. Global brands such as Twitter, Airbnb, and Twilio rely on Sift to gain competitive advantage in their markets. Visit us at sift.com and follow us on Twitter @GetSift.

Sources

1. Marketplace Risk, "Fraud Economics and the Dark Web's Multi-Billion Dollar Marketplace." https://www.marketplacerisk.com/post/fraud-economics-and-the-dark-web-s-multi-billion-dollar-marketplace
2. Okta, "3 Common Mistakes That Lead to a Security Breach." https://www.okta.com/identity-101/mistakes-that-lead-to-security-breach/
3. Bank Info Security, "3 Key Risks with Employee Passwords in the Financial Services Industry." https://www.bankinfosecurity.com/blogs/enzoic-blog-1-6x2-p-2801
4. Cybint Solutions, "Data Breaches 101: Why They Happen and What Data Gets Stolen." https://www.cybintsolutions.com/data-breaches-101-why-they-happen-and-what-data-gets-stolen/
5. Buzzfeed News, "How A Massive Facebook Scam Siphoned Millions Of Dollars From Unsuspecting Boomers." https://www.buzzfeednews.com/article/craigsilverman/facebook-subscription-trap-free-trial-scam-ads-inc
6. Sift, "COVID-19 and E-Commerce Fraud." https://sift.com/covid-19
7. Self, "A Complete List of Coronavirus (COVID-19) Scams." https://www.self.inc/info/coronavirus-scams/
8. Sift, "Digital Trust & Safety Assessment." https://pages.sift.com/digital-trust-and-safety-assessment-request.html
9. Sift, "Digital Trust & Safety Index: A Rapidly-Changing Fraud Landscape." https://resources.sift.com/ebook/digital-trust-safety-index-a-rapidly-changing-fraud-landscape/
10. Consumer survey data: On behalf of Sift, Dynata polled 1,000 adults across the United States via online survey, age 18+, between June 1 and June 8, 2020.