DIGITAL AD FRAUD 2019

Mobile and Video Remain Riskiest Channels

FEBRUARY 2019

Nicole Perrin

Contributors: Ross Benes, Lauren Fisher, Jillian Ryan, Tracy Tang

eMarketer®

DIGITAL AD FRAUD 2019: MOBILE AND VIDEO REMAIN RISKIEST CHANNELS

Digital advertising is a big business, with hundreds of billions of dollars spent each year. That money has long lured fraudsters to the space, and deceptive practices cost the advertising industry billions annually.

- **How big a problem is digital ad fraud?** Overall estimates of fraud vary widely, but even the most conservative estimates put the money involved worldwide well into the billions annually. Recent estimates vary from $6.5 billion to as high as $19 billion, a range that points to the difficulty in measuring fraud's true impact.
- **What types of digital advertising are most vulnerable to fraud?** One rule of thumb is that fraudsters go where the money is. That's true, but it's not the whole story. They're also attracted to new markets and technologies, such as connected TV, where verification firms haven't yet developed robust detection algorithms and where the ad tech market is relatively immature.
- **What are marketers doing to protect themselves?** Using a verification vendor is table stakes at this point, for advertisers and publishers alike. More companies are getting Trustworthy Accountability Group (TAG) certifications, implementing ads.txt, and evaluating their ad supply chains to eliminate avenues for fraud. But more can be done, and marketers will need to keep their eyes on the ball.

WHAT'S IN THIS REPORT? This report summarizes recent trends in digital ad fraud and its mitigation, and it advises marketers on what they need to understand to keep their budgets reasonably safe.

Share of Programmatically Enabled* Websites Worldwide that Have Implemented Ads.txt, Q4 2017-Q3 2018
% of total

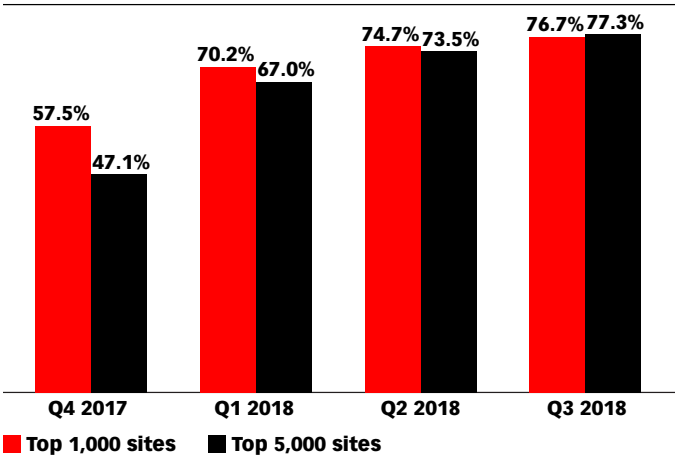

*Note: represents activities on Pixalate's platform, broader industry metrics may vary; *top 5,000*
Source: Pixalate, "Q3 2018 Ads.txt Trends Report," Dec 7, 2018

243921 www.eMarketer.com

KEY STAT: The adoption of ads.txt has been one of the most significant recent developments in programmatic ad fraud mitigation. According to measurement firm Pixalate, about three-quarters of leading programmatically enabled publishers have implemented ads.txt.

CONTENTS

THE PROBLEM

"How much of the internet is fake? Turns out, a lot of it, actually," asked and answered a New York Magazine headline in December 2018. Describing the recently broken-up 3ve ad fraud ring as "fake people with fake cookies and fake social media accounts, fake-moving their fake cursors, fake-clicking on fake websites," the article went on to note that around half of web traffic has long been nonhuman.

It also called out examples of everything from falsified metrics to nonexistent people to surreal, computer-generated "fake" content. On the internet, no one knows you're a bot.

This is the world that US advertisers must navigate as they spend more than $132.32 billion on digital placements this year, including $67.87 billion on display ads, according to our forecasts.

HOW BIG A PROBLEM IS DIGITAL AD FRAUD?

The two most recent publicly available estimates of overall losses to digital ad fraud are both from 2017—which is to say, not very recent. In May of that year, the Association of National Advertisers (ANA) and ad verification firm White Ops published their third annual joint "Bot Baseline" report, which estimated advertisers worldwide would lose $6.5 billion to ad fraud that year, down from $7.2 billion in 2016.

A few months later, in September, Juniper Research released the estimate that's more likely to be cited in the trade press: $19 billion in losses forecast for 2018, rising to $44 billion by 2022.

These estimates don't have a lot in common. Juniper's forecast for 2018 is almost three times as high as the ANA/White Ops estimate for 2017. The ANA/White Ops research emphasized that absolute losses to fraud were shrinking even as digital ad spending was rising, and that if more advertisers followed the lead of the savviest ones, digital ad fraud could be all but eliminated in a few years. The Juniper forecast instead predicted massive growth in fraud. The lower figure is considered by industry experts to have been too sanguine, but no one seems to think the problem is shrinking.

It's difficult to compare either figure with our estimate of total digital ad spending. Not only is limited methodological information available about the fraud estimates, but because our ad spending estimates are built in a bottom-up model based on publisher revenues, those estimates inherently exclude a large share of digital ad fraud. Put another way: We begin our estimates of the digital ad market by looking at actual publisher revenues, and that means the ad spend that gets siphoned off by fraudsters doesn't get counted in our figures.

What Is Ad Fraud?

This report includes information on a variety of fraudulent practices relating to digital advertising. Most of the fraud discussed involves various types of invalid traffic, including both general and sophisticated invalid traffic (GIVT and SIVT). General invalid traffic can include crawlers and bots of the type that make the internet work and aren't malicious at all (though they must still be removed from ad serving and measurement), as well as intentionally fraudulent bots. Sophisticated invalid traffic includes a number of ways fraudsters can make their invalid traffic appear more legitimate. Most of the fraud discussed in this report involves some type of invalid traffic, possibly in combination with other activity.

This report does not include discussions of brand safety or of viewability, except as they relate to invalid traffic.

Even without being able to pin a dollar figure on the cost of fraud, advertisers—especially programmatic advertisers—consider it a significant problem. A November 2018 survey by ad verification firm Integral Ad Science (IAS) found that 36.8% of brand marketers and 45.3% of agency respondents thought addressing ad fraud would be a priority in 2019.

When US media buyers were asked by Digiday about their biggest concerns for programmatic advertising in November 2018, more of those surveyed pointed to ad fraud than anything else, except for lack of transparency, which is the underlying condition making most ad fraud possible.

Ad Fraud Is US Media Buyers' No. 2 Concern About Programmatic Advertising

% of respondents, Nov 2018

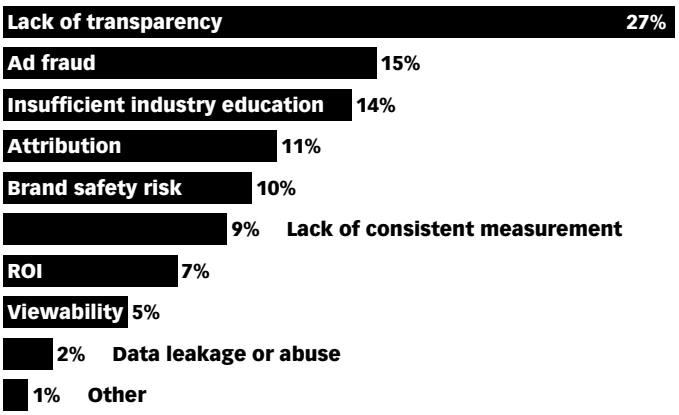

Note: n=282 media buyers; numbers may not add up to 100% due to rounding
Source: Digiday Research, "The state of programmatic media buying," Jan 2, 2019

244436 www.eMarketer.com

In a July 2018 Advertiser Perceptions survey of US agency and marketing professionals, fraud, including both viewability and nonhuman traffic, narrowly beat out brand safety concerns as the worst part of programmatic ad buying. In that survey, the "transparency" response referred specifically to pricing, better differentiating it from the problem of fraud.

What Are the Worst Aspects of Programmatic Ad Buying for US Agency and Marketing Professionals?

% of respondents, July 2018

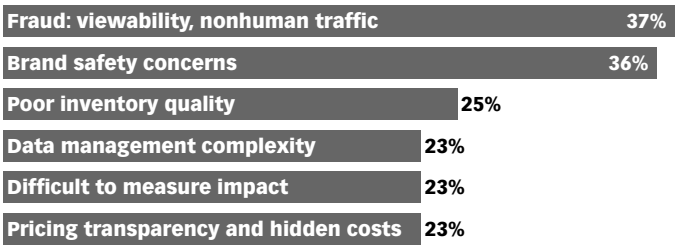

Source: Advertiser Perceptions, "DSP Report Q3 2018," Nov 5, 2018

243155 www.eMarketer.com

Fraud didn't top the list of challenges of programmatic buying in a June 2018 ExchangeWire survey of media agencies worldwide, but it was still a concern for a large share of respondents. Among those who did not own their own programmatic buying technology, 39% considered ad fraud a big problem. Almost half of owners of programmatic buying tech said the same.

Major Challenges of Programmatic Buying According to Media Agencies Worldwide, by Programmatic Buying Technology Ownership, June 2018

% of respondents

	Programmatic buying tech owners	Non-owners
Lack of transparency around media buys	78%	45%
Brand safety	59%	35%
Technological capabilities of third-party technology	54%	27%
Access to quality data	52%	55%
Lack of education/understanding	52%	49%
Programmatic ad fraud	48%	39%
Aligning existing KPIs with desired business outcomes	48%	35%
Measuring the incremental impact of media buys	46%	49%
Cost of third-party technology	46%	35%
Scale—lack of high-quality inventory	39%	45%

Source: ExchangeWire, "Agents of Change: The Rise of the Programmatic Media Agency" in association with Iponweb, July 12, 2018

239608 www.eMarketer.com

About six months earlier, in December 2017, 78% of US senior ad buyers surveyed by Cowen and Company said fraud was a problem with programmatic advertising. That was the top response and beat out transparency by 12 points.

Those surveyed were asked specifically about programmatic, but that's how the vast majority of digital display ad dollars are spent. We estimate that US advertisers spent about 82.5% of digital display dollars programmatically in 2018.

"Over the past 18 months, more and more, this really is a top-of-mind issue for the brands and publishers that we work with," said Mark Kopera, head of product for Oracle Data Cloud's mobile verification service Moat Analytics. "They're investing tons of resources in making sure, on the publisher side, that IVT [invalid traffic] is not a part of the traffic that they're selling advertisers. And on the brand or on the CMO side, they're enacting zero tolerance, or as close to it as is possible, for paying for impressions that are IVT."

PROGRAMMATIC DISPLAY

Programmatic digital display represents a huge chunk of digital ad budgets, both in the US and around the world, and it's highly vulnerable to fraud thanks in large part to long, complex and opaque supply chains. Scammers are able to take advantage of the distance between buyer and seller and the automated, intermediated nature of trading in a variety of creative and increasingly sophisticated ways. The result is significant percentages of invalid traffic to programmatic display campaigns—in other words, a lot of programmatic display impressions that never had any possibility of being seen by a human.

According to Q3 2018 data from measurement and analytics firm Pixalate, 17% of all programmatic digital display advertising impressions in the US during that time went to invalid traffic. That was the fourth-highest rate of programmatic ad fraud in the world, after India (34%), Indonesia (30%) and Australia (20%), but it wasn't evenly spread across inventory types. For traditional display formats, desktop web and smartphone apps were the two channels most vulnerable to fraud.

Programmatic Display Ad Fraud Rates Worldwide, by Device, Q3 2018
among impressions analyzed by Pixalate

Note: represents activity on Pixalate's platform, broader industry metrics may vary; read chart as 16.1% of desktop web display ad impressions were measured as invalid
Source: Pixalate, "Q3 2018 Ad Fraud Update," Jan 17, 2019

244696 www.eMarketer.com

For video formats, smartphone and tablet apps were the riskiest places to buy.

Programmatic Video Ad Fraud Rates Worldwide, by Device, Q3 2018
among impressions analyzed by Pixalate

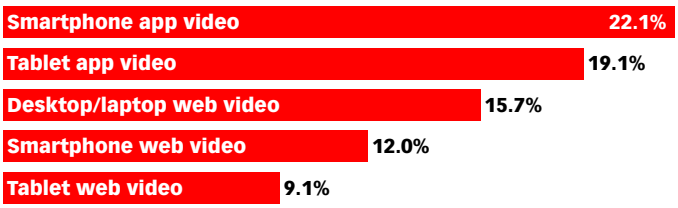

Note: represents activity on Pixalate's platform, broader industry metrics may vary; read chart as 22.1% of smartphone app videos impressions were measured as invalid
Source: Pixalate, "Q3 2018 Ad Fraud Update," Jan 17, 2019

244697 www.eMarketer.com

In November 2018, the Trustworthy Accountability Group (TAG) released a report evaluating the success of its program in the US. That report included, for comparison, a "blended rate" of fraud detected by DoubleVerify, Forensiq, IAS and White Ops, calculated by TAG at 10.43% across formats and channels.

According to IAS, 14.3% of worldwide impressions on desktop programmatic display campaigns that don't use its anti-fraud optimization service were fraudulent in H1 2018. Almost 10% of desktop programmatic video impressions were also fraudulent. Rates of invalid traffic to mobile web formats were lower. IAS did not report US fraud rates in its 2018 report on media quality due to volatility.

Programmatic Digital Display Ad Fraud Rates for Nonoptimized Impressions* Worldwide, by Device and Format, H1 2018
among impressions analyzed by Integral Ad Science

*Note: *impressions not using anti-fraud verification tools*
Source: Integral Ad Science, "Media Quality Report H1 2018," Sep 25, 2018

244620 www.eMarketer.com

The above fraud rates are based on impressions actually bought and sold. Ad fraud can also be detected and blocked further up the supply chain, through practices such as pre-bid verification. Digital advertising solutions provider RhythmOne reported that in Q4 2017, it blocked more than half of all desktop bid requests due to their riskiness—because they fit patterns associated with either invalid traffic or brand-inappropriateness. Across desktop and mobile devices, nearly half of all bid requests for video were suspect.

It's important to note that rates of fraudulent traffic don't translate in any clear way to rates of wasted spending. Riskier inventory is often associated with lower, too-good-to-be-true prices. On the other hand, fraudsters are attracted to inventory types with premium prices (like video), and bots often collect cookie profiles that make them look like attractive targets. This is just one reason why it's difficult to estimate how much money advertisers lose to fake impressions.

Spotlight on Domain Spoofing

Since our last report about digital ad fraud, major ad-supported publishers have adopted ads.txt to protect themselves from domain spoofing, a tactic in which fraudsters falsify the URL or site from which an impression is originating. The tool has also given them the opportunity to study how big a problem domain spoofing was—and remains. In September 2017, Digiday reported on an investigation the Financial Times conducted of its own liability to domain spoofing. The publisher found fraudulent display impressions being offered on 10 exchanges, and fraudulent video impressions on 15—especially noteworthy since FT.com wasn't selling any video inventory programmatically at the time. Hundreds of accounts were involved in selling an estimated £1 million ($1.3 million) in fraudulent impressions per month.

Later, in December 2017, a similar study of video inventory allegedly at 16 major publishers found that advertisers were spending $3.5 million per day on fraudulent impressions. Digiday noted that would add up to $1.3 billion per year in fraudulent video ads just across those sites (or, rather, not). For a sense of the scale, we estimate that US programmatic digital video ad spending was less than $17 billion in 2017.

Then, in July 2018, programmatic consultancy MightyHive announced the results of a test it ran with publisher Guardian US and Google, investigating how much money The Guardian might be losing to domain spoofing. MightyHive conducted two programmatic display buys, one that accepted only ads.txt-compliant inventory and another without that restriction. The study partners reconciled all campaign info to track exactly which ad purchases had resulted in revenues to Guardian US. For traditional display units bought without the ads.txt restriction, the risk was fairly low: 1% of campaign spending was diverted to fraud. But for video units, the share of spending lost to fraud was 72%.

RISKY INVENTORY

Certain types of inventory are more vulnerable than others—recall Pixalate's findings about varying fraud rates across devices and formats. What makes formats or channels particularly attractive to fraudsters? For one, a scarcity of legitimate inventory.

One of the biggest clichés about ad fraud is that fraudsters follow the money. That's true, of course—there's no point to their activities other than making money. But they also look for opportunities where fraud is easier to perpetrate. That could be for technical reasons, a few of which are outlined below. But it can also mean fraudsters see opportunity in channels where they know marketers are desperate for more inventory than legitimately exists.

"New, niche inventory formats or types, like native audio, OTT [over-the-top] or connected TV, usually give rise to new supply sources that generally pique the interest of not only advertisers, but also fraudsters," said Nich Seo, platforms lead for innovation at MightyHive.

Video, which isn't exactly new or niche, is an area where marketers are known to demand more impressions than can be supplied by premium publishers. This makes it highly susceptible to fraud, as found by MightyHive's research with Guardian US. Almost half of respondents to YouAppi's 2018 survey of digital marketers worldwide said they were concerned about fraud on the mobile ad networks that deliver their video, making it the leading challenge for the format.

In 2019, while video generally may still be vulnerable, all eyes are on mobile in-app advertising and OTT and connected TV (CTV) advertising as the riskiest for buyers.

Mobile Apps

The problems in the mobile space go beyond impression fraud, which is the most common if not the only problem with desktop and mobile web display ads. Because significant mobile advertising budgets hinge on actions like app installs, fraudsters also falsify that activity and engage in other forms of attribution fraud.

According to mobile analytics and attribution platform AppsFlyer, more than a quarter of mobile app installs worldwide in October 2018 were fraudulent. The rate in the US was lower, but still significant, at 14.7%. Fraud rates fluctuated from month to month but were generally on an upward trajectory throughout 2018.

Mobile App Install Fraud Rate in the US and Worldwide, Jan 2018-Oct 2018
% of total paid app installs

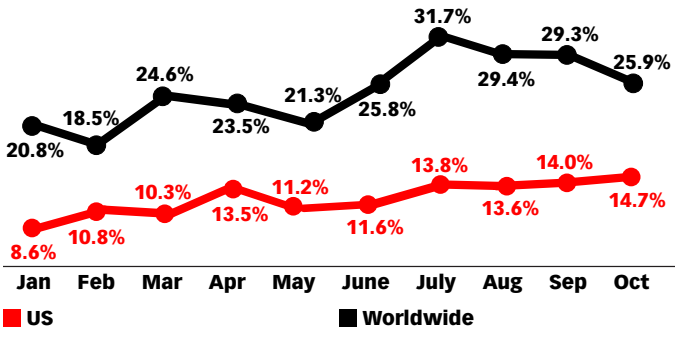

Note: represents activity tracked by AppsFlyer, broader industry metrics may vary
Source: AppsFlyer, Nov 2018

243153 www.eMarketer.com

"The environment is significantly different from desktop," Roy Rosenfeld, vice president of product management at verification provider DoubleVerify, said. "The telemetry that is available for the ad seller is significantly different. The methods that fraudsters employ are significantly different."

And those methods vary. According to AppsFlyer's analysis of the mobile app install fraud that took place in February 2018, for example, bots were the biggest source of suspicious activity, followed by device farms. The previous month, the leading culprits were device farms and click flood attacks.

Share of Mobile App Install Fraud Worldwide, by Type, Aug 2017-Feb 2018
% of total fraudulent paid app installs

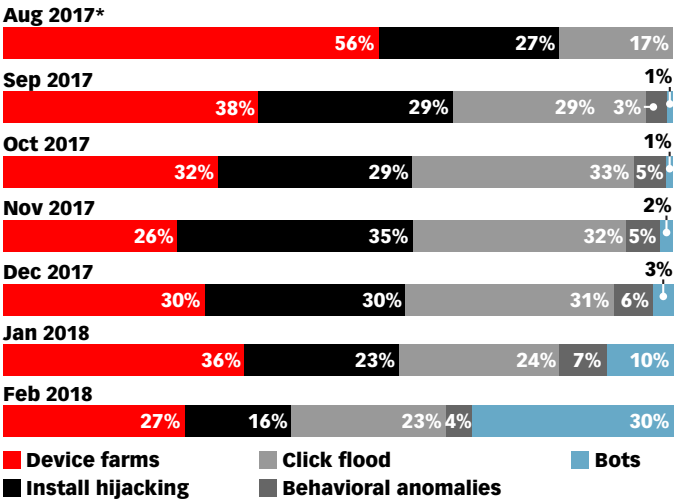

*Note: represents activity tracked by AppsFlyer, broader industry metrics may vary; *total for period of June-August 2017*
Source: AppsFlyer, "The State of Mobile Fraud: Q1 2018," April 2, 2018

236651 www.eMarketer.com

The retail industry was most likely to be affected by app install fraud in Q1 2018, AppsFlyer reported, followed by the gaming industry.

"The mobile ad industry doesn't look much cleaner now than it did a year ago, but the mix of that fraud has changed," Grant Simmons, head of client analytics at mobile analytics and attribution firm Kochava, said. "The appetite for networks to flood signal, like 'Let's get as many clicks introduced into the ecosystem as possible,' that has historically been the biggest standalone fraud type. It still is, but that's been down-trending. What has been on the uptrend has been manufactured installs, so fake installs. Some of it's really lazy, like device IDs that are impossible and SDK [software development kit] versions that were never in market."

"Overall, IVT rates or fraud rates Moat-wide or across a full brand media buy may not look super alarming depending on the brand at first glance," Oracle's Kopera said. "They could be in the low single digits, for example, but you really have to dig a lot deeper than that to figure out where the pockets of fraud are going. These sophisticated fraudsters are moving away from the stuff where it's hard and where they can't make money anymore, and they're finding the pocket of apps or sites that can be exploited."

Those pockets might have extremely high rates of invalid traffic. Kopera cautioned against looking at only top-level figures and emphasized the need to optimize campaigns away from those types of locations.

In many cases, mobile in-app ad fraud relies on hijacked devices, which can request thousands of ads as apps run in the background or overnight while the phone screen isn't even on. A couple of the experts interviewed for this report noted that there could theoretically be a consumer interest in fighting ad fraud from this perspective as well—but the average mobile user likely neither knows nor cares.

"In-app fraud, both in attribution and just ad impressions, is a big problem," said Shailin Dhar, CEO and co-founder at verification provider Method Media Intelligence. "One, because not many people pay attention to how much data each app is using. And the other thing is people download absolutely crazy things on their phone. We see it quite prevalent with these utility apps being on everybody's phones, and nobody's ever looked to see 'Why is my QR code reader app using 1.5 GB of data every month?' It's just not a general human technology hygiene process. Engineers and ad tech nerds like us look at that stuff, because nobody else really cares—especially if you're on Wi-Fi most of the time."

OTT

Over-the-top (OTT) TV is another area where a confluence of high advertiser demand, low legitimate supply and novel tech infrastructure has led to a high level of fraud. According to Pixalate, 19% of worldwide OTT impressions and 18% of US OTT impressions were invalid in Q3 2018. (That's not due only to malicious fraud—the tech infrastructure of OTT is not yet seamless and can lead to incorrect ad serving or measurement unintentionally.)

"Roku has its own store. Amazon has its own store. Xbox has its own store. A lot of the media companies have their own store," said Amy King, vice president of product marketing at Pixalate. "There is no standardization across all of those apps and stores. There's no universal directory of what those apps should be called. So if you decide to make up fake apps and call them something very similar, it makes it very difficult to tell if that's invalid or fraudulent or not." On mobile, these types of IDs have been standardized. "Even if there's no fraudulent intent, it still makes it very difficult for a buyer or seller to know what they're buying on, because of the fragmentation and lack of standardization across those apps in all of those stores," King said.

"The demand is very, very, very strong," DoubleVerify's Rosenfeld said. "But the supply is very, very limited, even more so than on mobile apps. And what we've created is an exceptional opportunity for fraudsters to essentially generate more inventory, even if it doesn't exist, which is conjuring it from thin air, and it's selling."

"OTT is really going to be scrutinized this year," Mike Juhas, executive vice president of client services at media solutions firm Digital Remedy, said. "If you take a step back and think about it, yes, OTT is a premium channel for advertisers—but how many people are actually consuming the content from a lot of these smaller providers that are claiming they have so much scale?"

Experts cautioned marketers that special fraud detection methods are required in OTT or CTV. It's often not a browser environment, and measures like clicks aren't meaningful. So marketers need a specific strategy to fight fraud on this channel, including working with verification partners that have begun to tackle the channel and vetting supply partners for trustworthiness.

Social and Influencer Fraud

When digital advertisers think about fraud, they probably think more of the open programmatic exchanges than the walled social gardens. But we would be remiss not to mention the wide world of social fraud, especially when one of the biggest national news stories for the past two years has been fake accounts spreading disinformation.

We have long published estimates of social media users that were lower than what the networks reported publicly, because we've attempted to estimate individual human users and not double count users with multiple accounts—or count bots. And the bot population has been substantial in the past. In May 2018, Facebook reported that it had removed 1.3 billion fake accounts over the prior six months. In September, the social platform announced the latest six-month rolling tally of takedowns: 1.5 billion.

In its Q4 2018 earnings update, Facebook announced that it had 2.32 billion monthly active users as of December 2018. eMarketer's estimate of 1.64 billion worldwide Facebook users in 2018 is about 70% of that figure.

Twitter users regularly accuse each other of being bots, and Instagram ostensibly bans follower-buying services but still ends up accepting ads from those same services.

Brands that work with influencers are generally familiar with the fact that some buy fake followers and work together in various schemes to boost their engagement metrics. They've begun using tech tools to protect themselves, as well as working more with microinfluencers, where they feel more confident in an authentic connection between social personality and followers. But they may not think of how the broader picture of fake social engagement might be affecting them.

For one thing, there's measurement. An underappreciated aspect of fraud is that it doesn't mean just that advertisers spend money on impressions or actions that never happen. It also means advertisers incorporate data about those impressions or actions into their overall models. If bots, or human-farmed devices, are interacting with social content, what's popular or trending may not have much to do with what real people care about. Plus, brands are vulnerable to being dragged into any of the social media controversies foreign and domestic political actors continually gin up with the help of bots and fraudulent ad buys.

For more information on how brands are using social media personalities to market authentically, see our forthcoming report on influencer marketing, and this report from 2018: **"Global Influencer Marketing: What Platforms to Use, Policies to Follow and the Paths to Purchase Around the World."**

THE SOLUTIONS

Though none of them completely fix advertisers' fraud problem yet, marketers do have a number of tools available to help fight it. At this point, several of them can be considered industry best practices, like using verification providers and implementing ads.txt.

AD VERIFICATION SERVICES

For major digital advertisers, it's considered table stakes at this point to work with an ad verification provider. There's no question they weed out many fraudulent impressions, and they can help advertisers (along with publishers and others in the programmatic ecosystem) keep tabs on how much invalid traffic is getting through. This allows them to raise issues with their agency, demand-side platform (DSP) or publisher—and claw back spending on ad impressions that had no chance of being seen by a human.

According to Integral Ad Science, 1.8% of desktop display ad impressions on campaigns using its anti-fraud verification service were fraudulent in H1 2018, along with 0.8% of mobile impressions. IAS did not release a directly comparable figure for fraud rates among impressions overall, but it reported that worldwide rates of fraud for programmatic desktop and mobile display campaigns that did not use protection were 14.3% and 9.9%, respectively. By its own detection and measures of fraud, at least, the service eliminates the vast majority of invalid traffic.

US Desktop and Mobile Display Ad Performance Metrics: Fraud* Rate Among Ads Optimized Against Fraud, by Purchase Method, H1 2018

among impressions analyzed by Integral Ad Science

	Desktop	Mobile
Publisher direct	2.4%	0.8%
Programmatic	1.7%	0.7%
Total	**1.8%**	**0.8%**

*Note: represents activity on the Integral Ad Science platform, broader industry metrics may vary; read as 1.7% of programmatic desktop digital display ads from campaigns that used ad fraud prevention tools were fraudulent; *inventory with no possibility of being viewed by a human, including general and sophisticated invalid traffic defined by the MRC*
Source: Integral Ad Science, "Media Quality Report H1 2018," Sep 25, 2018

243927 www.eMarketer.com

Not all verification services work the same way or specialize in the same channels. Some, for example, sample a portion of impressions to determine an overall fraud rate, without being able to point to specific impressions, clicks or app installs that were fake, while others check each transaction. Marketers should feel comfortable asking questions about their verification service, how it works and what it reports. Vendors won't reveal proprietary information, of course, but should be able to provide basic explanations of their practices—and many are actively working to educate advertisers through content marketing like plain-English blog posts explaining how they uncovered the latest scheme.

"I would encourage marketers to take a proactive approach to fraud, to ask questions of the vendor," said Stefano Vegnaduzzo, senior vice president of data science at IAS. "Don't just turn it on and forget it. Ask questions; try to understand not the technical details maybe, but the types of fraud. Ask questions about flag rate vs. accuracy. And in general, try to engage with the vendor if you are a CMO."

Marketers cannot set it and forget it—and they should also keep in mind that merely using a verification service isn't enough to protect themselves.

"A lot of times advertisers think if you work with one of the verification partners, you're doing enough," MightyHive's Seo said. "But there's also a lot more diligence to be done on the buy side relative to how you're actually running your campaigns, optimizing, etc. It's a great first step to work with a verification company or partner, but that's ultimately not enough. And a lot of marketers confuse that with being enough."

"It's important, obviously, to measure your media buys and keep a monitor on fraud rate, not just at the aggregate level, but at the detailed domain or app or seller level," Oracle's Kopera said. "It's important to have KPIs [key performance indicators] beyond reach, clicks or conversion, so that you're not optimizing toward some of the KPIs that we know sometimes fall prey to these schemes, especially when budgets are tighter and quarters are ending."

Other guardrails for ad campaigns might include pre-bid targeting, as well as making sure contracts with partners provide for clawbacks. "There's a whole bunch of different tools in the toolkit, and depending on the media strategy, there's four or five of the seven or eight that you should apply," Kopera said. "Some of them are technical, some of them are policy or contract, and some of them are just business process. Marketers need to have a good handle on all of those."

Another reason verification services aren't a silver bullet is many can't point to specific sources of fraudulent impressions—especially if they're using sampling methods that don't consider all impressions. Add to that the different methodologies of the different services and the fact that any number of parties to a programmatic transaction may be using different providers with different results, and it can be impossible for advertisers to be sure they're getting the right refund on fake ads.

"There is no standardization right now in what is considered fraud or what is considered IVT," Dhar of Method Media Intelligence said. "The MRC [Media Ratings Council] guidelines are basically only housekeeping guidelines of: Do you have the data stored in the appropriate way? Do you have security protections on your database? Do you have a distribution of access for all your engineers? Do you have a feedback loop properly set up to make sure that any new fraud you catch is caught in the future by including it in your algorithm going forward?"

So different providers have different methodologies, and if they're using sampling, their machine learning algorithms can actually learn different things because they're looking at different samples of traffic.

A further problem with verification services: Fraudsters are constantly trying to reverse-engineer their filtering algorithms. Internet searches reveal numerous results for traffic sellers offering "visitors" that will pass specific vendors' verification filters.

TAG CERTIFICATION

Using a verification service is just one component of the Certified Against Fraud program run by the Trustworthy Accountability Group (TAG) since 2016. Depending on the role a company plays in the digital ad ecosystem, TAG Certified Against Fraud compliance has various requirements, like having a compliance officer, attending regular trainings and using ads.txt. As of February 2019, TAG listed 119 companies as being eligible for the "Certified Against Fraud" seal—which signals to partners that transacting with these companies should be less risky than average.

And, according to TAG, that is the case. In November 2018, it released the results of research it conducted the previous July with 614 Group, which examined 75 billion US impressions—display and video units across desktop, mobile web and mobile app. Just 1.68% of those impressions were fraudulent. The report compared this figure with an average fraud rate of 10.43%, meaning TAG-certified channels were 84% less prone to fraud. In January 2019, TAG released a report analyzing ads in Europe and finding similar effects.

It's important to note that, here and above, verification firms are measuring their own success at weeding out fraud. The ever-present problem with measuring fraud is that no one is ever sure they've found it all—and it's more reasonable to assume they haven't. Verification firms are successful at finding and mitigating the types of fraud they look for, but there's no guarantee they're looking for all the latest things scammers are doing.

Opinions on TAG and other trade group activity aren't uniformly positive, however. For example, TAG sells IP blacklists to help fight fraud—a tactic some experts say is actively a bad idea. Proxy IPs are readily available, and money spent on blacklists means less available for more effective anti-fraud methods.

ADS.TXT

Implementing ads.txt, short for "authorized digital sellers," is also a component of TAG certification. Ads.txt is a text file that publishers can use to add to their sites a list of authorized sellers of their inventory. Its implementation has probably been the single most important development in fraud mitigation efforts since our last report was published in mid-2017. Since that time, more than three in four leading publishers with programmatic advertising have implemented it, according to Pixalate data from Q3 2018.

Share of Programmatically Enabled* Websites Worldwide that Have Implemented Ads.txt, Q4 2017-Q3 2018
% of total

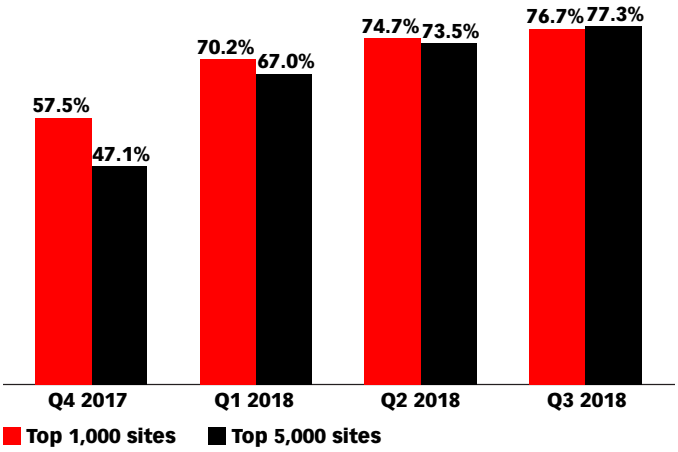

*Note: represents activities on Pixalate's platform, broader industry metrics may vary; *top 5,000*
Source: Pixalate, "Q3 2018 Ads.txt Trends Report," Dec 7, 2018

243921 www.eMarketer.com

Ads.txt doesn't fix or fight all types of fraud—it's limited to web environments, for one thing, and for another, it aims only to weed out domain spoofing and arbitrage. But those are both huge problems in the programmatic display world. And, as outlined in the "Spotlight on Domain Spoofing" section of this report, ads.txt can make a big difference in terms of how much fraudulent inventory buyers are exposed to.

The ollout of the ads.txt file isn't just about publishers, however. DSPs and other ad tech partners have had to change how they work so they can allow buyers to bid only on compliant inventory. In 2018, there was indication that some buyers would want to stick with noncompliant inventory because it was cheaper. But compliance continues to grow, and MightyHive's Seo said that shortly after the test they ran on Guardian inventory, the DSP used to test the noncompliant side of inventory switched to offering only ads.txt-compliant inventory.

CLAWBACKS

Another shift in the industry since around the time of our last report is a greater expectation among advertisers that they'll be able to claw back spending on fake ad activity, often if fraud rates exceed an agreed-upon threshold. It's far from a perfect system. It means, in many cases, multiple parties to a given transaction are all paying verification services. They reconcile their findings in a time-consuming process at the end of campaigns. And no one is really sure if they're getting all their money back, or if it's from all the right people. Programmatic supply chains are too long and complex for advertisers or their partners to be completely satisfied with this system, but it's considered a major improvement on earlier attitudes. Advertisers should make sure their media buyers have such deals in place and understand their terms.

Sahil Gupta, director of global partnerships at Adobe Ad Cloud, recalled that the DSP had started its fraud refund program about two years ago. "It seemed audacious at that point. But fast forward to now, and we've seen a lot of our competitors launch similar programs," he said.

"The biggest issue we're still seeing is that there is generally no way to take action against any detected fraud numbers," Method Media Intelligence's Dhar said. "There is not much that I can do after the fact. Finding out where in the supply chain it came from is close to impossible, because these are all aggregate numbers. Without log files knowing which actual seller ID—not the name of the domain, but the name of the company with that seller ID and the exchanges that sold it—I don't know who to withhold money from. If I go back to my DSP and say, 'Hey, I've detected this fraud with this vendor, here's the data,' I still don't have enough information for the DSP to go back to its supply partners and figure out who it came from."

LOOKING AHEAD

Experts expect the fraud fighting ground to keep shifting, as anti-fraud efforts in spaces like mobile apps and OTT video make progress and innovation among fraudsters continue. But few seem to be putting their faith in purely technical solutions, when they know all of those also require human solutions—working together to find problems, making sure new standards are implemented industrywide, and educating advertisers about how they might be vulnerable.

Verification firms and ad tech providers have already been working together on issues like 3ve and other big fraud ring takedowns, but more in way of information sharing could be helpful.

"We are collaborating a lot more, and I think in 2019 I see a bunch of that still to come," Oracle's Kopera said. "A lot of the focus on what trade groups can do together is on the fraud or the sophisticated invalid traffic. It's also important to note the general invalid traffic, which is the stuff that's not crime; it's just that sites get crawled by various spiders and things like that. It's part of how the internet works, and brands shouldn't have to pay for it. But also, we're seeing a lot of false positives and inconsistencies on what is and isn't GIVT. And we see things like corporate proxies and people out on university campuses being classified as invalid traffic. Another opportunity for these industry groups is potentially to come up with some better standards on that, so we don't use our precious resources on sorting out those inconsistencies. And instead, we're focused on sophisticated invalid traffic."

The industry is also looking to updated standards, like the forthcoming OpenRTB 3.0 with ads.cert, which will extend the capabilities of ads.txt, as well as app-ads.txt, to eliminate more fraud. But those interviewed for this report were not optimistic about the length of time it would take for these new standards to be adopted.

Many experts also noted the need for law enforcement to stay involved for fraud levels to actually fall. Fraudsters view their activities as low-risk, especially when they're using data centers rather than hijacked devices to create invalid traffic. Opinions as to whether the indictments in 3ve were a one-off or part of a new trend were divided.

For more on the rollout of app-ads.txt and continued federal involvement in fighting fraud, see our January 2019 report **"Digital Display Advertising 2019: Nine Trends to Know for This Year's Media Plan."**

The industry is also looking to improved measurement in areas like mobile apps and OTT TV to help make fraud more difficult. "Mobile in-app is a dominant form of advertising across the board right now, whether its US or across the world," Pixalate's King said. "That will remain a very appealing area for fraudsters, especially since the measurement capabilities haven't quite caught up with the volumes yet."

But improved measurement will also require improved KPIs from marketers. Some performance-type measures may appear helpful, but actions other than impressions can still be gamed. Ultimately, tying measurement back to real business results is the best way to make sure digital measurements are representative of real activity.

"The conversation oftentimes changes when the CMO or the CFO is in the room," Kochava's Simmons said. "At that point, we can talk about what the overall corporate goal is here. We need to market with some intent toward causality, not just how many clicks and impressions can get sprinkled around so that we can attribute installs that were going to happen anyway."

Advertiser education will also be key, and it's improving.

"We've seen that increasing over the last 10 years, and even more over the past two or three years," DoubleVerify's Rosenfeld said. "People have become aware, not only about the impact this makes on them from a business perspective, but also from the pure stance of campaign performance, and how it is wise or unwise to spend media budget."

"One of the areas for me that's very important for marketers to understand is the distinction between flag rate, detection rate and accuracy," Vegnaduzzo of IAS said. "So, what I mean by detection rate: the percentage of fraud that a vendor will report to you in your advertising specs. So, 5%, 10%, whatever it is. That doesn't mean that all the inventory that was flagged as fraud is actually fraud. That's the second concept of accuracy. What we have now is this flag rate war about 'Oh, my flag rate is higher than yours, so I'm catching more fraud.' Is that true? Well, you cannot tell, because you don't know whether all that inventory is actually correctly identified as fraud."

Spotlight on 3ve

In November 2018, the FBI and a number of ad tech firms, including verification provider White Ops, revealed that a massive digital ad fraud ring had been disrupted—and the alleged perpetrators indicted. 3ve was interesting for a few reasons, including the sophistication of the fraud operation. But what really caught the industry's interest was the fact that a number of entities came together to identify and solve the problem. That law enforcement was involved was also unusual.

Marketers interested in understanding fraud better would do well to read the **report on the takedown of 3ve** co-authored by Google and White Ops. The body of the report is about 12 pages of plain English about how 3ve worked and how the scheme was discovered.

Once marketers are educated, they'll have to be ready to change how they buy media, because they currently depend on risky inventory and practices to meet their campaign targets.

EMARKETER INTERVIEWS

Shailin Dhar
CEO and Co-Founder
Method Media Intelligence
Interviewed November 19, 2018

Sahil Gupta
Director, Global Partnerships
Adobe Ad Cloud
Interviewed January 17, 2019

Tamer Hassan
CTO
White Ops
Interviewed November 30, 2018

Mike Juhas
Executive Vice President, Client Services
Digital Remedy
Interviewed January 22, 2019

Amy King
Vice President, Product Marketing
Pixalate
Interviewed January 8, 2019

Mark Kopera
Head of Product, Moat Analytics
Oracle Data Cloud
Interviewed January 10, 2019

Roy Rosenfeld
Vice President, Product Management
DoubleVerify
Interviewed December 4, 2018

Shani Rosenfelder
Head of Content and Mobile Insights
AppsFlyer
Interviewed January 23, 2019

Maor Sadra
CEO
Applift
Interviewed February 1, 2019

Grant Simmons
Head of Client Analytics
Kochava
Interviewed January 15, 2019

Stefano Vegnaduzzo
Senior Vice President, Data Science
Integral Ad Science
Interviewed January 9, 2019

Dorian Kim
Vice President, Retargeting and Programmatic Operations
Applift
Interviewed November 30, 2018

Nich Seo
Platforms Lead, Innovation
MightyHive
Interviewed January 16, 2019

READ NEXT

Blockchain's Promise: How Blockchain Could Increase Transparency, Reduce Friction and Solve Audience Identity Challenges

Cleaning Up the Digital Media Supply Chain: What Will Move the Industry Toward Transparency and Optimization?

Digital Display Advertising 2019: Nine Trends to Know for This Year's Media Plan

Global Influencer Marketing: What Platforms to Use, Policies to Follow and the Paths to Purchase Around the World

Mobile App Installs: What You Need to Know About User Acquisition

Mobile Video Advertising 2019: Leveraging Rewarded Video, 6-Second Ads and Vertical Video

Programmatic Advertising Beyond Display: Automation Moves to Audio, Out-of-Home and Television

Q4 2018 Digital Video Trends: As Consumers Increase Time Spent with Video, More Ad Spending Goes Through Programmatic

Television Update, Fall 2018: Advanced TV's Progress in Addressable, Programmatic and OTT

US Digital Video and TV StatPack 2018: Ad Spending Data and Audience Metrics for Digital and Traditional Platforms

US Programmatic Ad Spending Forecast Update 2018: Video Powers Significant Growth Through 2020

SOURCES

614 Group

Advertiser Perceptions

AppsFlyer

Association of National Advertisers (ANA)

Cowen and Company

DoubleVerify

Digiday

Google

ExchangeWire

Financial Times

Forensiq

Guardian US

Integral Ad Science (IAS)

Iponweb

Juniper Research

MightyHive

Pixalate

RhythmOne

Trustworthy Accountability Group (TAG)

White Ops

YouAppi

EDITORIAL AND PRODUCTION CONTRIBUTORS

Anam Baig	Senior Editor
Joanne DiCamillo	Senior Production Artist
Katie Hamblin	Chart Editorial Manager
Dana Hill	Director of Production
Erika Huber	Copy Editor
Ann Marie Kerwin	Executive Editor, Content Strategy
Stephanie Meyer	Senior Production Artist
Heather Price	Managing Editor, Content
Magenta Ranero	Senior Chart Editor
Amanda Silvestri	Senior Copy Editor

eMarketer The leading research firm for marketing in a digital world.

Coverage of a Digital World

eMarketer data and insights address how consumers spend time and money, and what marketers are doing to reach them in today's digital world. Get a deeper look at eMarketer coverage, including our reports, benchmarks and forecasts, and charts.

Confidence in the Numbers

Our unique approach of analyzing data from multiple research sources provides our customers with the most definitive answers available about the marketplace. Learn why.

Customer Stories

The world's top companies across every industry look to eMarketer first for information on digital marketing, media and commerce. Read more about how our clients use eMarketer to make smarter decisions.

Your account team is here to help:

Email research_requests@emarketer.com to submit a request for research support, or contact accounts@emarketer.com or 866-345-3864 to discuss any details related to your account.

To learn more about eMarketer advertising and sponsorship opportunities, contact advertising@emarketer.com.