UNDERSTANDING FRAUD AND INVENTORY QUALITY IN PROGRAMMATIC

BrightRoll
from YAHOO!

AN EVERGREEN ISSUE

Ad fraud and inventory quality are always-relevant topics in the world of programmatic advertising. If you hear otherwise from an adtech partner, our advice is simple: run for the hills.

The Association of National Advertisers estimates that bot fraud could cost the advertising industry *$7.2 billion* globally in 2016. And while that number is staggering, nonhuman traffic isn't the only challenge in digital advertising. Beyond fraud, low-quality inventory impacts negative brand associations, perception, inaccurate viewer data, and wasted human effort.

Every day, consumer-facing information technology grows and improves, providing people with exciting new ways to interact with the world and offering brands new ways to capture the hearts and minds of connected consumers. These developments also open the door for fraudsters to become more sophisticated, devising new methods for defrauding advertisers of both meaningful views and funds.

Adtech providers, publishers, and exchanges must work continuously to combat the threat of ad fraud and ensure a baseline inventory quality across buys.

THE TYPES OF FRAUD

Fraud and quality are two sides of the same coin. Ensuring high-quality inventory is a natural part of fighting fraud, and vice versa.

Fraud in digital advertising typically falls into the following categories:

Fraudulent Traffic

Non-human traffic

When people hear "online ad fraud," they typically think of non-human traffic, or bot fraud. This refers to traffic generated by automated programs—bots, spiders, or content scrapers—designed to mimic human behavior on the web. These bots act as fake viewers, generating clicks and impressions that are not real. If a page sees unusually high visitor traffic or suspicious traffic patterns, it can be the result of traffic fraud. Fraudsters use botnets—a network of infected computers that, unbeknown to their owners, follow a set of instructions from a central program—to generate huge numbers of fraudulent views on websites.

Human-Generated Traffic

Bot fraud isn't the only type of fraudulent traffic. Some fraud is perpetrated by humans. Click farms—groups of people incentivized to visit websites to rapidly generate impressions and clicks—are a commonly used human source for driving fraudulent traffic.

Not All Bots are Bad

Not all non-human traffic is malicious. For instance, publisher-operated bots are used to crawl sites in order to ensure the integrity of links, flag commentary, and other tasks relating to website maintenance. It's important to understand what's happening on a website to ensure that the right actions are taken when encountering non-human traffic. However, it is important to filter out all impressions generated by bots even if the intent is not malicious.

THE TYPES OF FRAUD

Fraudulent Supply

Non-Transparent Supply

When websites make non-transparent ad calls obscuring their domain names or other information, it can be a signal that the websites have reason to hide their identity. These non-transparent websites are frequently low-quality sites that seek to be anonymous or masquerade as high-quality, premium supply.

Impression fraud

Impression fraud takes place when an ad pixel fires without a legitimate ad view taking place. This can be achieved by stacking ads in a single frame, loading ads into 1x1 hidden pixels, or playing video ads in the background. This means the ads you pay for might never have the opportunity to be seen by your audience.

Low-Quality Supply

Brand Safety Risks

You can find everything under the sun—and more—online. There are entire categories of websites that a brand might not want to be associated with: sites with adult materials, gambling, hate speech, sites involved with consumer fraud, identity theft, or illegal distribution of copyrighted content and software. Having an ad run on any of these kinds of sites poses serious risks to any brand, even if the sites are not technically committing ad fraud.

Fake Sites

These are sites created with the sole purpose of driving traffic. They feature no original content, use high ad density, non-viewable impressions, and redirection to aggressively monetize all traffic that arrives at their sites. While these sites generate impressions, they are not valuable to advertisers.

What about mobile?

For the most part, mobile fraud is conducted the same way as desktop fraud, with a few variations.

Fraudulent Traffic

Incentivized Clicking Games

These are apps that incentivize users to click as many ad experiences as possible, gamifying the process. These clicks generate revenue for the app owner, but do not provide quality views to advertisers.

Background App Activity

Our smartphones allow most apps to function in the background. It is possible for an app designer to continue running ads while the app is not visible to the user, generating unseen impressions.

Low-Quality Supply

Fake Sites

Even though the various app stores have their own filtration and inventory quality processes, some fake apps inevitably get through. Fake apps, like fake websites, obscure their ad calls and masquerade as premium app experiences.

Fraud is a Many-Headed Beast

While these categories do not tend to change over time, the methods for executing fraud are constantly in flux. Imagine a game of whack-a-mole: when one method or site is put down, another pops up in its place. Because of the mutable nature of fraud, it's important to build rigorous quality standards into the industry, and evolve on proprietary methods for fighting fraud.

THE BRIGHTROLL APPROACH TO QUALITY

The best way to minimize fraud is to have a rigorous process for maintaining high-quality supply.

On an industry-wide scale, we work actively with TAG, the Trustworthy Accountability Group, to combat fraud. On our own platforms, we vet inventory before we make it available, institute pre-bid solutions to attempt to block fraud at the impression level, and monitor and validate post-impression to fuel our algorithms and manual auditing efforts. We use a combination of proprietary technology, third-party integrations, human monitoring, and industry partnerships to make sure that inventory across our platforms remains high-quality and minimally fraudulent.

OUR APPROACH TO QUALITY

Our internal inventory vetting process follows three steps in a continuous cycle:

The Cycle – Step 1: Supply Verification

We take a two-pronged approach to supply verification to ensure inventory quality: a comprehensive onboarding process, and human monitoring. First, we internally audit all publisher inventory prior to onboarding it. The multi-stage onboarding process begins with reviewing the publisher's application, URL, ranking and volume.

This is followed by an in-depth review of several conditions, including publisher company type, content, industry code, and several other parameters during the bid request. Provided the publisher meets our criteria, we then review technical requirements: for example, video publishers must accept VAST/VPAID-compliant creative, have an ad server and player, and confirm that all URLs reviewed contain content that is safe for work, i.e. not violent or pornographic.

The second piece of our supply verification process relies heavily on people, one of the most important aspects of programmatic technology: continuous human monitoring. BrightRoll has dedicated teams responsible for proactively analyzing, reviewing, and monitoring the quality of BrightRoll's traffic and investigating any claims of illegitimate activity.

The Cycle– Step 2: Pre-Bid Filtering

To help prevent fraud, specifically non-human traffic, we use proprietary technology and third-party integrations to evaluate each available impression. We look at domain, geography, IP address, past performance, and other key signals that indicate if an impression is generated by a source that might be perpetrating fraud. Each impression that fails to meet our standards is filtered before an advertiser has the opportunity to deliver an impression. This thorough vetting process eliminates millions of potentially fraudulent impressions from the platform per month.

The filters we apply to supply lead to insights which are applied to machine learning models that, in turn, determine supply quality in real time:

- **Domain and URL Review** - Every publisher is pre-reviewed and profiled, then, based on its qualities, added to a whitelist or blacklist. The data from domain profiles is fed to machine-learning models to generate reputation scores for new domains. The DSP bids only on supply that is in the whitelist.
- **Content Classifications** – Content classifiers can be used to flag prohibited content, such as nudity, weapons, or gambling.
- **Traffic Rules** – All incoming traffic is analyzed against the l latest blacklists that capture ongoing patterns.
- **Real-Time Scoring for Bid Events** – A machine-learning scoring model uses behavioral patterns in concert with real-time traffic to predict quality scores for incoming bid requests in real time.

OUR APPROACH TO QUALITY

The Cycle – Step 3: Post-Serve Monitoring and Measurement

There's always the possibility that, despite our best efforts, we miss something. If fraudulent supply slips through, we work to catch it the next time. BrightRoll's ad-serving systems scan and tag any ad serving events that appear to be non-human or invalid. These systems use several methodologies to learn and refine the detection rules continuously. They use real-time and historical data to detect anomalies, constantly improving the quality of the algorithm as they work.

Just as a referee shouldn't also be playing the game, digital publishers and ad-serving providers shouldn't be the ones measuring viewability and fraud. Being able to use neutral third-party measurement partners enables advertisers to better compare performance across publishers and ensures appropriate return on advertising spend while maintaining brand safety.

We offer independent viewability and fraud measurement for all display and video advertising through established third-party measurement partners. Our server-to-server integrations allow for seamless measurement by minimizing errors and data loss, while increasing coverage.

When we're done with step three, we take the most important step: **Return to Step One**

There are natural fluctuations in the amount of fraud measured during a calendar year. It makes sense that we see increased rates of fraud that correspond with spikes in spending across the board. During the holiday season, for instance, when buys climb dramatically, fraud follows suit. The important thing to know about fighting fraud is that it's a continuous battle, one in which technology providers, publishers, and third-party organizations must remain vigilant. We strongly believe that a continuous approach is the best way to combat fraud, as our methodology continues to evolve, even while the fraudsters' techniques do.

WHAT DOES THE FUTURE HOLD?

As we've mentioned before, fighting ad fraud and maintaining inventory quality are cyclical processes that require constant attention and reexamination. While we can't say for certain what the fraud of tomorrow will look like, we know what it'll take to combat it: strong industry partnerships, constant vigilance, and pioneering new technologies.

While the internet will never be 100% fraud-free, building a safer experience for users and brands is a goal the industry as a whole can get behind. As digital ad budgets grow and more money is put into programmatic campaigns, adtech providers will play a role in creating as brand-safe and confidence-instilling an environment as possible. Working diligently together to reduce fraud can help programmatic advertising succeed, which can ultimately help advertisers reach audiences, publishers monetize their platforms, and give users a relevant and engaging online experience.

BrightRoll from YAHOO!

brightroll.com/contact